Investment Opportunities and the Sources of Lifetime Inequality*

Kartik Athreya[†] FRB Richmond

Felicia Ionescu[‡] Federal Reserve Board

Urvi Neelakantan[§] FRB Richmond

Ivan Vidangos[¶] Federal Reserve Board

February 15, 2016

Very preliminary and incomplete. Please do not cite.

Abstract

How much of the dispersion in lifetime earnings, wealth, consumption, and ultimately, utility or well-being is resolved early in life (prior to working life) vs. later? This critical question has received much attention, with influential recent contributions coming from Keane and Wolpin (1997), Storesletten et al. (2005), and Huggett et al. (2011). The goal of this paper is to provide quantitative measures of how the full range of households' investment (and financing) opportunities matters for the fraction of lifetime inequality determined relatively early in life relative to later on. We focus on the role played by three specific investment opportunities: risky and lumpy college education, risky equity, and costly borrowing. To our knowledge, our work is the first to provide quantitative measures of the importance of each in the temporal resolution of lifetime inequality and the importance of the interaction between them. We find, first, that nearly all income inequality is attributable to human

*The views expressed in this paper are those of the authors and do not necessarily reflect the views of the Federal Reserve Bank of Richmond or the Federal Reserve System.

[†]Federal Reserve Bank of Richmond, P.O. Box 27622, Richmond, VA 23261, Kartik.Athreya@rich.frb.org, Ph:804-697-8225

[‡]Board of Governors of the Federal Reserve System, felicia.ionescu@frb.gov

[§]Federal Reserve Bank of Richmond, P.O. Box 27622, Richmond, VA 23261, Urvi.Neelakantan@rich.frb.org, Ph:804-697-8913

[¶]Board of Governors of the Federal Reserve System, ivan.vidangos@frb.gov

capital variation as of age 23. In other words, for individuals who have likely completed major educational investments, our results suggest that financial opportunities play only a minor role in altering inequality. Second, we find that the option to invest in high-return, high-risk assets meaningfully increases the importance of initial inequality, whereas the ability to borrow lowers the importance of initial inequality.

JEL Codes: D31; E21; G11; I24; J24;

Keywords: Inequality; Human Capital; Higher Education; Financial Investment

1 Introduction

How much of the dispersion in lifetime earnings, wealth, and ultimately, utility or well-being is resolved early in life (prior to an individual's working life) vs. later? The goal of this paper is to understand *quantitatively* how an empirically plausible representation of the full range of households' investment (and financing) opportunities matters for the portion of inequality (in income, wealth, consumption, and utility) that is determined relatively early in life (e.g. at age 18 or 23) relative to that determined afterward. We will focus on what we view as the three primary investment opportunities available to individuals: human capital in the form of college education, risky equity, and costly borrowing.

To our knowledge, our work is the first to measure the roles played by three potentially critical types of investment opportunities for the temporal resolution of lifetime inequality. We find, first, that for a population that begins making financial investment decisions at age 23, nearly all income inequality is predetermined and attributable to human capital variation as of age 23. In other words, for individuals who have likely completed major educational investments, our results suggest that financial opportunities play only a minor role in altering lifetime earnings and wealth inequality. Second, we find that the option to invest in a high-return, risky financial asset (stocks) in addition to low-return, risk-free assets meaningfully increases the importance of initial conditions. In particular, the amount of inequality attributable to human capital as of age 23 increases by 4 percentage points once households are assumed to have access to stocks. Individuals with high initial human capital have a high opportunity cost of investing in human capital and choose to invest in stocks instead. Since these individuals have high earnings and wealth to begin with, access to stocks in effect enables the rich to get richer. The absence of stocks as an investment option shuts down this channel, and therefore dampens the effect of initial human capital on inequality.

When it comes to wealth inequality and income, a natural question is the extent to which borrowing constraints matter. Early in life, borrowing constraints may significantly affect human capital investment. Later in life, borrowing constraints may affect the ability of households to achieve desired financial investment positions. Moreover, because households who vary in their human capital attainment will also vary in their earnings, households' willingness to bear risk will differ in the population. Thus borrowing constraints may systematically change the distribution of risk-bearing capacity in the economy. Therefore, if access to credit is unequal, so too will be lifetime outcomes, not just for earnings but also for wealth. We study the role of credit in a model where households may invest in both stocks and bonds. We find that access to credit lowers the contribution of (age-23) initial conditions by 5 percentage points.

While age 23 is a natural starting place for studying the role of financial investments on the resolution of lifetime inequality, the results above must be interpreted carefully. In particular, the attribution of substantial resolution of uncertainty to "initial conditions" cannot be taken as a statement about the role played by wholly exogenous factors. Most obviously, initial conditions at age 23 likely reflect the majority of formal education acquisition. Indeed, the need for more and better educational investment is the usually drawn implication in almost any policy discussion of how to deal with inequality (see, e.g., Goldin and Katz, 2007; Heckman and Krueger, 2005). To account for the role played by formal educational investment, we next allow households the opportunity to invest in four-year college education, where that investment is risky and lumpy. [Results from college model to be added.]

The broader question of the temporal resolution of inequality has received significant attention, with recent important contributions coming from influential papers by Keane and Wolpin (1997), Storesletten et al. (2005), and Huggett et al. (2011). Keane and Wolpin (1997) argue, using a model in which individuals choose occupations, that the overwhelming proportion (90 percent) of inequality is determined by forces that place different individuals on different trajectories very early in life (by age 16). Storesletten et al. (2005) examine this question using a heterogeneous-agents model with uninsurable earnings shocks, where earnings are exogenous, and find that nearly one-half of the dispersion in lifetime earnings uncertainty is resolved before individuals enter the labor market (by age 23). Lastly, Huggett et al. (2011) find, in a Ben-Porath model with consumption and savings decisions, that pre-working-life (by age 23) dispersion accounts for about two-thirds of lifetime inequality. The findings of all three papers support the idea that inequality has deep roots in early life. This assessment is consistent with a large stream of work pioneered by James Heckman and coauthors (see, e.g., Heckman and Carneiro, 2003, for a review) and is also the subject of recent work by Caucutt et al. (2015).

We lay out a model in which agents are heterogeneous with respect to their innate ability and their initial human capital and financial wealth (measured at age 18, upon high school completion), and they are subject to earnings shocks during their working life. In addition, agents have the option to invest in a project resembling risky college education, and have the ability to finance this investment via costly borrowing, and to invest in risky, but high-return, financial assets. We then calibrate this model, given that it contains the essential choices we view as available to individuals in practice. As part of the calibration, we estimate the joint distribution of ability, initial human capital, and initial wealth at age 18 to match (1) key earnings statistics and (2) college enrollment and completion outcomes by groups of initial wealth. We then use the calibrated model to provide a decomposition of the contribution of initial conditions (in this case, pre-college) and human capital shocks over the working life to dispersion in lifetime earnings, wealth, consumption, and utility. Given this benchmark measurement, we then conduct a series of counterfactuals in which we systematically remove household access to the three investment/financial-market tools under examination. Comparison of these cases to the baseline model will uncover the role being played by each.

We now describe our approach to modeling each investment opportunity in more detail. In practice, a critical part of human capital accumulation arises through college education, and disparities in educational attainment are, at least empirically, tightly associated with disparities in life prospects. Indeed, a large empirical literature in labor economics has demonstrated that college education (and the lack thereof among many) is a central factor behind the large degree of economic inequality in the United States, as well as its substantial growth in recent decades.[1]

College, in turn, contains several features with significant potential for altering the view of how initial conditions matter. In particular, prior work (e.g., Caucutt and Kumar, 2003; Akyol and Athreya, 2005; Abbott et al., 2013; Garriga and Keightley, 2007; Hendricks and Leukhina, 2014) almost uniformly approaches college as a risky, lumpy investment. This work also recognizes the irreversibility inherent in all human capital acquisition. Making explicit allowance for the features of college as a risky investment is also supported by the following three observations. First, there is evidence that the returns to education are not a smooth function of the number of years of education, with the return to a year of education that yields a diploma being higher than previous

[1]See, for example, Lemieux (2010) and Goldin and Katz (2007), who show that most of the growth in the variance of wage rates (at least for men) is accounted for by changes in the between-group variance of wage rates induced by growing top-end education wage differentials. Furthermore Altonji et al. (2013) use a rich statistical model of earnings, wage rates, hours, employment, and job changes to decompose the variance of male earnings over the life cycle at different points in the life cycle. Although this is not the main focus of their paper, they use their model to decompose the variance in lifetime earnings into the contribution of factors that are known when an individual enters the labor market (observed and unobserved heterogeneity) vs. various shocks that occur in their model over the working life. They find that about 55 percent of the variance in lifetime earnings is due to factors representing heterogeneity, and that differences in observed educational attainment account for more than one-half of that contribution.

ones (Hungerford and Solon (1987)). Second, noncompletion rates in the U.S. suggest substantial risk: Around 50 percent (Bowen et al., 2009) of college entrants fail to obtain *any* degree within eight years of enrolling. Moreover, recent work by Ozdagli and Trachter (2011) suggests that the chances of exit are highest not immediately but at the two- and three-year mark. Lastly, and obviously, college is costly, not only in explicit costs but also in forgone earnings. We therefore model college as an investment opportunity that includes the following features: college is a time consuming, lumpy and risky project, which individuals may use personal assets or costly borrowing to finance. This richness allows us to discipline the model to match the array of facts on both college enrollment and completion.

With respect to capturing the effect of a richer array of financial investment choices beyond the traditional risk-free claim to the capital stock (or a bond), we note first that households frequently finance college (and their consumption while attending) by reducing assets or taking on nondefaultable debt. These factors are likely to put downward pressure on household investment in stocks, and hence, future wealth. They are, however, partially offset by higher earnings (Card, 1999) and lower risk of unemployment (Mincer, 1991) for those who obtain a college degree. As a whole, these features of college can each be expected to affect the impact of any pre-college distribution of "innate" ability and "initial" human capital.

The possibility that investment in a college education may fail to yield a college degree also makes the presence of financial investment alternatives more relevant, all else equal. Consider an individual with a low likelihood of success in college. For such a person, the availability of high return (though risky) financial assets may well be attractive. For her, it may be optimal to forgo college, and instead work hard and save in risky assets (or simply borrow less). Thus, the influence exerted by high-yield financial assets on the dispersion of lifetime income depends on the distribution of college failure risk. The question is: By how much?

We note that the work cited above, without exception, simplifies by providing households with investment options that do not vary in their risk-return characteristics. In principle, this may well matter, especially once college is modeled. The lumpiness of college will situate many individuals with large debts early in life that will, all else equal, dissuade investment in risky assets and, in turn, diminish wealth accumulation. As a result, the evolution of the distribution of financial wealth hinges on the decisions individuals make with respect to higher education. We therefore evaluate the effect of a richer menu of financial assets on the assessment of early and late resolution of inequality. Conversely, the inability to take advantage of the stock market will have effects on overall savings, and the decision over when, if ever, to invest in financial wealth.[2]

[2]Understanding the role played by the specifics of the college investment process is also clearly relevant for policy. At an immediate level, the explicit accommodation of college is necessary to even begin consideration of

With respect to the role played by the risky, lumpy human capital investment in college, there are two sets of forces that work in opposition to each other. On the one hand, the ability of households to greatly enhance their human capital in one step via college will, all else equal, mute the role of initial (pre-college) conditions. On the other hand, the riskiness of college will act as a disincentive to invest in college, and as such, may amplify the power of initial conditions. In particular, initial human capital and financial wealth may interact particularly strongly. For instance, wealth-poor households will have a relatively lower willingness to take the risk of college than their better-heeled counterparts–even if both have the same initial ability and human capital. In sum, college and its specifics have the potential to significantly alter the measurement of the relative importance of learning ability, initial human capital, and financial wealth. But the question is ultimately a quantitative one.

In broad terms, our aim is to open the "black box" of what would, in a model without explicit allowance for the particulars of college and access to financial instruments, be reflected in initial conditions (of ability and human capital) *as of entry into the workforce*, especially at age 23.

We turn now to the model (Section 2) and data (Section 3). The calibration is laid out in Section 4, and results are provided in Section 5.

2 Model

2.1 Overview

Time is discrete and indexed by $t = 1, ..., T$ where $t = 1$ represents the first year after high school graduation. Agents enter the model (as high-school graduates) endowed with a level of human capital, h^{HS}. Central to our model is their decision whether or not to attend college. College enables agents to make a substantial investment in human capital, which is productive in the labor market once they leave college. However, to capture an important source of risk to human capital, we assume that agents may fail to complete college.[3] At the end of four years in college,

major policies such as student loan subsidies, need-based aid, and the choice of regime governing the consequences of student loan default. When combined with what will be a more accurate assignment to pre-college heterogeneity in ability and human capital, our approach has promise for delivering more accurate implications not only for the temporal resolution of uncertainty, but also for the implications of the policies noted above. While not a focus of our paper, we note that our approach is capable of helping shed light on the value of policies directed at modifying or providing insurance for initial conditions (e.g. public education) relative to policies directed at alleviating the effects of shocks over the working lifetime (e.g. unemployment insurance). Maybe more importantly, they also suggest that policies that make human capital more equal at an early age are more likely to be effective than policies that make financial wealth more equally distributed (e.g. estate taxation).

[3] For example, Bound et al. (2010) report, using NLS72 data, that only slightly over half of all college enrollees graduated within 8 years of enrollment.

the probability of completion—which depends on the agent's innate ability and human capital accumulation while in college—is realized. Those who complete college start their working life with human capital h^{CG} while those who fail to complete start life with human capital h^{SC}, where SC denotes "some college." Those who choose not to go to college start their working life with human capital h^{HS}.

Agents who choose to attend college can finance their education using any wealth that they have. They can also take on student debt, d, and unsecured debt, b, both of which are non-defaultable. Unsecured debt is subject to a borrowing limit, $-\underline{b}$, where $\underline{b} > 0$. Once agents leave college, they become working adults who divide their time between work and human capital accumulation, as in the model of Ben-Porath (1967). They also consume and save. Agents can save in both a risky asset, s, and a risk-free asset, b. As before, the instrument b may also be used to borrow, in which case $b < 0$.

We allow for three potential sources of heterogeneity across agents in the first model period — their immutable learning ability, a, their initial stock of human capital, h, and their initial assets, x. The set of these characteristics in the first model period are drawn jointly according to a distribution $F(a, h, x)$ on $A \times H \times X$, and we allow for and estimate correlations between returns to stocks, bonds, and human capital.

Agents work and accumulate human capital until $t = J - 1$ and start retirement in period $t = J$. After retirement, they face a simple consumption-savings problem.

2.2 Preferences

The general problem of an individual is to choose consumption over the life-cycle, $\{c_t\}_{t=1}^T$ to maximize the expected present value of utility over the life cycle

$$\max_{(\{c_t\} \in \Pi(\Psi_0))} E_0 \sum_{t=1}^{T} \beta^{t-1} u(c_t), \tag{1}$$

where $u(.)$ is strictly concave and increasing. $\Pi(\Psi_0)$ denotes the space of all feasible combinations $\{c_t\}_{t=1}^T$, given initial state Ψ_0. Agents have a common discount factor, β. Preferences are represented by a standard time-separable CRRA utility function over consumption. Agents do not value leisure.

2.3 Human Capital

Agents can invest in their human capital in two ways — by investing in a college education when young and by apportioning some of their available time to acquiring human capital as adults through on-the-job training. Both types of investment in human capital are risky. Agents who invest in college face a significant risk of noncompletion. Furthermore, all human capital, included that accumulated through on-the-job training, is subject to shocks. The stock of human capital represents an agent's "earning ability" and can be accumulated over the life cycle, in contrast to learning ability, which is fixed at birth and does not change over time. We assume that the technology for human capital accumulation is the same during and after college and that human capital is not productive until after graduation.

2.3.1 College investment

While in college, agents optimally choose to divide time between work and human capital accumulation. Human capital is not productive until agents leave college. Working during college diverts time from human capital accumulation and therefore increases the probability of non-completion.[4] In addition, college students have jobs that pay a low wage and do not necessarily value students' human capital stock or contribute to human capital accumulation (Autor et al., 2003). However, students of high ability may be hired in better paid jobs than students of low ability. Thus, we model a wage rate per time units worked in college, $w_{col}(a)$, instead of per efficiency units; this rate increases with the ability level of the student. This assumption prevents low-ability students from enrolling in college only to enjoy earnings during college that are much higher than the earnings they would have earned had they not enrolled in college. We assume that the growth rate in earnings during college is 0.

College investment is risky. The risk of noncompletion depends on the human capital stock at the end of college, which in turn is determined by the agent's decision to allocate time to human capital accumulation during college. If a student with initial human capital h_1 decides to acquire a college degree, the probability with which she succeeds is given by $\pi(h_5(h_1, a, l^*_{1,\ldots,4}))$. This is a continuous, increasing function of the human capital stock after college years, h_5, which in turn increases with the initial human capital stock, h_1, the ability of the individual, a, and her choice of time devoted to human capital investment during college years, $l^*_{1,\ldots,4}$. This formulation captures

[4] Our modeling is motivated by several important observations. Numerous studies show that working while in school can adversely affect academic performance (Stinebrickner and Stinebrickner, 2003) and increase the likelihood of dropping out (Braxton et al., 2003). Thus, unlike in an environment where the possibility of failing to complete college is not modeled (or is completely exogenous), allowing for the choice to allocate time to work versus human capital accumulation during college is a key ingredient when accounting for the risk of investing in human capital.

the idea that college preparedness, embodied in h_1, student's learning capacity, captured in a, and effort to invest in human capital during college are important determinants of college completion. If the student completes college, she will walk into period 5 as a college graduate. If not, she will walk into period 5 as a college dropout with some college education.[5]

There are several possible sources of college financing: family contributions, work during college, need-based student loans and non-defaultable unsecured debt. Students may also use their labor income and savings (if any) during college to finance their college education. Agents are allowed to take out student loans up to $d(x) = \max[d_{max}, (\overline{d} - x)]$, which represents the full college cost, $\overline{d}$, minus the expected family contribution, x, up to a student loan limit d_{max}. They choose the loan amount, d, at the beginning of college, and they receive equal fractions of the loan each period in college. After college they will repay this loan in equal payments, p which are determined by the loan size, $d(x)$, interest rate on student loans, R_g and the duration of the loan, P. Consistent with the data, the interest rate on student loans is $R_f < R_g < R_b$, where R_f is the risk-free savings rate and R_b the borrowing rate on unsecured debt.

2.3.2 Human capital investment as on-the-job training

College graduates, college dropouts, and high school graduates who do not enroll in college optimally allocate time between market work and human capital accumulation as on-the-job training during the adult phase of their life (as in the classic Ben-Porath, 1967, model).

Human capital evolves according to the human capital production function, $H(a, h_t, l_t)$, which is increasing in shocks to human capital, z_{t+1}, the agent's current stock of human capital, h_t, immutable learning ability, a, and the fraction of available time put into human capital production, l_t. The law of motion for human capital is given by

$$h_{t+1} = \exp(z_{t+1}) H(h_t, l_t, a) \tag{2}$$

Following Ben-Porath (1967), the human capital production function is given by $H(h, l, a) = h + a(hl)^{\alpha}$ with $\alpha \in (0, 1)$. As in Huggett et al (2011), we assume that shocks to human capital are i.i.d. over time. However, these shocks lead to persistent differences in human capital and therefore produce persistent differences in earnings over the life cycle.

[5]Modeling investment in four-year college and the risk of dropping out at the end of the fourth period in the model are justified by data: according to the BPS 1996/2001, 68.5% of students enroll in four-year colleges. Our findings also show that 89% of college dropouts are enrolled in college at least for three full years.

2.4 Labor Income

An agent's working life corresponds to $t = t_w, ..., J-1$, with $t_w = 1$ for those who choose not to enroll in college and $t_w = 5$ for college graduates and dropouts. In this phase, the human capital stock is valued in each period in the labor market. Earnings are given by the product of the rental rate of human capital, w_t, the agent's human capital stock, h_t, and the time spent in market work, $(1 - l_t)$:

$$y_t = w_t h_t (1 - l_t) \tag{3}$$

The rental rate of human capital evolves over time according to $w_t = (1 + g_i)^{t-1}$, where the growth rate, g_i, depends on whether or not agents have a college degree. Specifically, the growth rate for college graduates, g_{cg}, is greater than the growth rate for those with no college, g_{nc}, which includes college dropouts and high school graduates.[6]

2.5 Means-Tested Transfer and Retirement Income

Because we want our model to reflect an empirically realistic representation of the risk environment faced by households, we allow agents to receive, in addition to labor income, means-tested transfers, τ_t, which depend on age, t, income, y_t, and net assets, x_t. These transfers capture the net effect of the various U.S. social insurance programs, which are aimed at providing a floor on consumption. Following Hubbard et al. (1994) we specify these transfers by

$$\tau_t(t, y_t, x_t) = \max\{0, \underline{\tau} - (\max(0, x_t) + y_t)\} \tag{4}$$

Total pre-transfer resources are given by $\max(0, x_t) + y_t$ and the means-testing restriction is represented by the term $\underline{\tau} - \max((0, x_t) + y_t)$. These resources are deducted to provide a minimal income level $\underline{\tau}$. For example, if $x_t + y_t > \underline{\tau}$ and $x_t > 0$, then the agent gets no public transfer. By contrast, if $x_t + y_t < \underline{\tau}$ and $x_t > 0$, the the agent receives the difference, case in which he has $\underline{\tau}$ units of the consumption good at the beginning of the period. Agents do not receive transfers to cover debts, which requires the term $\max(0, x_t)$. Lastly, transfers are required to be nonnegative, which requires the "outer" max operator.

After period $t = J$, in which agents start retirement, they get a constant fraction $\phi_i(y_J)$ of their income in the last period as working adults, which they allocate between risky and risk-free investments.

[6]The growth rates for wages are estimated from data. Evidence shows that wage growth rates for college dropouts and people with no college are similar and are lower than the growth rate for college graduates.

2.6 Financial Markets

There are two financial assets in which the agent can invest, a risk-free asset, b_t, and a risky asset, s_t.

Risk-free assets

An agent can borrow or save using asset b_t which can be 0, positive, or negative. Savings will earn the risk-free interest rate, R_f. We assume that the borrowing rate, R_b, is higher than the savings rate: $R_b = R_f + \phi$. Debt is non-defaultable and comes with a borrowing limit $\underline{b} > 0$.

Risky assets

We call the risky assets "stocks" and we denote the agent's holdings of equity between period t and $t+1$ by s_{t+1}. This amount entitles the owner to its stochastic return in period $t+1$, $R_{s,t+1}$. This represents a gross real return and its excess return is given by:

$$R_{s,t+1} - R_f = \mu + \eta_{t+1}, \tag{5}$$

where η_{t+1}, the period $t+1$ innovation to excess returns, is assumed to be independently and identically distributed (i.i.d.) over time and distributed as $N(0, \sigma_\eta^2)$. We assume that innovations to excess returns are uncorrelated with innovations to the aggregate component of permanent labor income.

Given asset investments at age t, b_{t+1} and s_{t+1}, financial wealth at age $t+1$ is given by $x_{t+1} = R_i b_{t+1} + R_{s,t+1} s_{t+1}$, with $R_i = R_f$ if $b \geq 0$ and $R_i = R_b$ if $b < 0$.

2.7 Agent's Problem

The agent maximizes lifetime utility by choosing college investment, asset positions in bonds and stocks, time allocated to market work and to human capital, and borrowing.

We formulate the problem in a dynamic programming framework where any period t variable j_t is denoted by j and its period $t+1$ value by j'. The state vector is defined as follows. An individual's feasible set of consumption and savings is determined by his age, t, ability, a, beginning-of-period human capital, h and net worth, $x(b, s)$, and current-period realization of the shock to human capital, z.

We solve the problem backwards starting with the last period of life when agents consume their savings. The value function in the last period of life is set to $V_T^R(a, h, x) = u(x)$. For the retirement phase, the value function is given by

$$V^R(t,a,h,b,s) = \sup_{b',s'}\{\frac{c_t^{1-\sigma}}{1-\sigma} + \beta V^R(t+1,a,h,b',s')\} \tag{6}$$

where

$$c + b' + s' \leq \phi_i(y_J) + R_j b + R_s s$$

In the above, $R_j = R_f$ if $b \geq 0$ and $R_j = R_b$ if $b < 0$.

Retired agents do not accumulate human capital. They face a simple consumption-savings problem but may choose to invest in both risk-free and risky assets.

Given important differences between education groups during the remaining stages of the life cycle, we present the value functions and further details separately for the no college and college paths.

2.7.1 No College

We use $V_J^{R,i}(t,a,h,b,s)$ from Equation 6 as a terminal node for the adult's problem on the no college path. Note that $V_R^{HS}(t,a,h,b,s,z_{nc}) = V_R^R(a,h',b',s')\ \forall z_{nc}$.[7] We solve for the set of choices in the working phase, for which the value function is given by

$$V^{HS}(t,a,h,b,s,z) = \sup_{l,h',b',s'}\{\frac{c_t^{1-\sigma}}{1-\sigma} + \beta E V^{HS}(t+1,a,h',b',s',z')\} \tag{7}$$

where

$$\begin{aligned} c + b' + s' &\leq w(1-l)h + R_b b + R_s s + \tau(t,y,x) \text{ for } t = 1,..,J-1 \\ & l \in [0,1], h' = exp(z'_{nc})[h + a(hl)^\alpha] \end{aligned}$$

The value function $V^{HS}(t,a,h,b,s,z)$ gives the maximum present value of utility at age t from states h, b, and s, when learning ability is a and the realized shock is z. Solutions to this problem are given by optimal decision rules $l_j^*(t,a,h,b,s,z)$, $h^*(t,a,h,b,s,z)$, $b^*(t,a,h,b,s,z)$ and $s^*(t,a,h,b,s,z)$, which describe the optimal choice of the fraction of time spent in human capital production, the level of human capital, and risk-free and risky assets carried to the next period as a function of age, t, human capital, h, ability, a, and current assets, b and s when the realized state is z. The value function, $V^{HS}(1,a,h,x)$, gives the maximum expected present value of utility if the agent chooses not to go to college from initial state h, when learning ability is a and initial assets are x.

[7]The same idea is used whenever we switch from one phase to another without explicitly stating this.

2.7.2 College

As before, the problems for college graduates and dropouts are solved backwards, starting with the retirement phase for which the value function is given by Equation 6.

We use $V_J^{R,i}(t,a,h,b,s)$ from the retirement phase as a terminal node and solve for the set of choices in the working phase $j = 5, .., J-1$ of the life cycle. We further break down this phase into a student loan post-repayment period and a repayment period. For the post-repayment period, $t = P+1, ..., J-1$, the problem is identical to the one for working adults on the no-college path. Recall that during the repayment period after college, $t = 5, ..., P$, agents have to repay their student loans with a per-period payment $p = \frac{d(x)}{\sum_{t=1}^{P-5} R_g^t}$. Here $d(x_1)$ is the size of the loan, which depends on initial assets, x_1, and R_g is the interest rate on student loans.

The value function is given by

$$V^i(t,a,h,b,s,z) = \sup_{l,h',b',s'} \{\frac{c_t^{1-\sigma}}{1-\sigma} + \beta E V^i(t+1,a,h',b',s',z')\} \tag{8}$$

where

$$\begin{aligned} c + b' + s' &\leq w(1-l)h + R_j b + R_s s + \tau(t,y,x) \text{ for } t = P+1, .., J-1 \\ c + b' + s' &\leq w(1-l)h + R_j b + R_s s + \tau(t,y,x) - p(x_1) \text{ for } t = 5, .., P \\ & l \in [0,1], h' = exp(z_i)[h + a(hl)^\alpha] \end{aligned}$$

where $i = CG, SC$; $R_j = R_f$ if $b \geq 0$ and $R_j = R_b$ if $b < 0$.

For the college phase $t = 1, .., 4$ of the life cycle we first take into account the risk of dropping out from college and use $V^C(5,a,h,b,s,z) = \pi(h_5)V^{CG}(5,a,h,b,s,z) + (1-\pi(h_5))V^{SC}(5,a,h,b,s,z)$ as the terminal node to solve for the optimal rules. Agents invest in their human capital during college and they may decide to work. Each period in college they pay direct college expenses, $\hat{d}$. In the first period, they also choose the loan amount for college education, d, which will be equally divided in four rounds of loans during college years. The value function is given by

$$\begin{aligned} V^C(t,a,h,b,s,z) &= \max_{l,h',b'',d} \left[\frac{c^{1-\sigma}}{1-\sigma} + \beta V^C(t,a,h,b,s,z)\right] \qquad (9) \\ & c + b' + s' = w_{col}(1-l) + R_b b + R_s s + d/4 - \hat{d} \\ & l \in [0,1],\ h'^{\alpha}] \\ & d \in D = [0, \overline{d}(x)] \text{ for } t = 1. \end{aligned}$$

Solutions to this problem are given by optimal decision rules $l_j^*(t,a,h,b,s,z)$, $h^*(t,a,h,b,s,z)$, $b^*(t,a,h,b,s,z)$ and $s^*(t,a,h,b,s,z)$, which describe the optimal choice of the fraction of time spent in human capital production, the level of human capital, and risk-free and risky assets carried to the next period as a function of age, t, human capital, h, ability, a, and current assets, b and s when the realized state is z. The value function,$V^C(1,a,h,x)$, gives the maximum expected present value of utility if the agent chooses to go to college from initial state h, when learning ability is a and initial assets are x.

Once the college and no-college paths are fully determined, agents then select between going to college or not by solving $\max[V^C(1,a,h,x), V^{HS}(1,a,h,x)]$.

3 Data

In order to map our model to data, we use data on annual earnings from the March Current Population Survey (CPS), and we use data on financial assets (wealth) from the Survey of Consumer Finances (SCF).

3.1 Life cycle earnings

As described in more detail in the next section, we calibrate our model to match the evolution of mean earnings, earnings dispersion, and earnings skewness over the lifecycle. To this end, we first estimate lifecycle profiles, for ages 23 to 60 (i.e. the "working life"), of mean earnings, the earnings Gini coefficient, and the mean/median earnings ratio using data from the March CPS, obtained through IPUMS at the University of Minnesota. We use data on annual wage and salary income for male heads of household with at least a high-school diploma (or equivalent) for calendar years 1963-2013 (corresponding to survey years 1964-2014). We restrict our sample to individuals who worked at least 12 weeks in the reference year and earned at least $1,000 (in constant 2014 prices). We use the CPS weights to ensure that each year's sample is representative of the overall U.S. population; additionally, we renormalize the weights in each year in order to keep the population constant at its 2014 value; this way we abstract from issues related to population growth.

We use these data to construct lifecycle profiles for mean earnings, the earnings Gini coefficient, and the mean/median earnings ratio. Specifically, for each of these statistics, $s_{t,y}$, we compute $s_{t,y}$ in the data for each combination of age t and calendar year y, and regress $s_{t,y}$ against a full set of year and age indicators.[8] We then take the regression coefficients on the age indicators (we

[8]By using a full set of year indicators, this treatment controls for year effects in the construction of the age profiles. We have also computed age profiles controlling for cohort effects, rather than year effects. The behavior

use calendar year 2013 as our base year), and normalize them so that at age 40 the coefficients profile goes through the unconditional average value of $s_{40,y}$ across all years y in our sample. The corresponding normalized age coefficients constitute the lifecycle profiles that we use in the calibration. Figure 1 shows the lifecycle profiles of mean earnings, the earnings Gini, and the mean/median earnings ratio obtained in this fashion.

3.2 Life cycle financial assets

We use data from the SCF to measure wealth and its composition. Our measure of wealth includes all financial assets. To be consistent with assumptions that we make later, we assume that wealth is comprised of one risky and one risk-free asset. Our measure of risky assets corresponds to a broad measure of households' equity holdings in the SCF, which includes directly held stocks as well as stocks held in mutual funds, IRAs/Keoghs, thrift-type retirement accounts, and other managed assets.

As in the case of earnings, we construct lifecycle profiles of asset holdings. We do this following a procedure similar to Cagetti (2003): We pool data from all nine waves of the triennial survey conducted during the 1989–2013 period and calculate age-specific, unconditional, weighted means (by education group). The results for the pooled sample (in 2013 dollars) are reported in Figure 2-4.

Figure 2: Average Life cycle Assets (SCF)

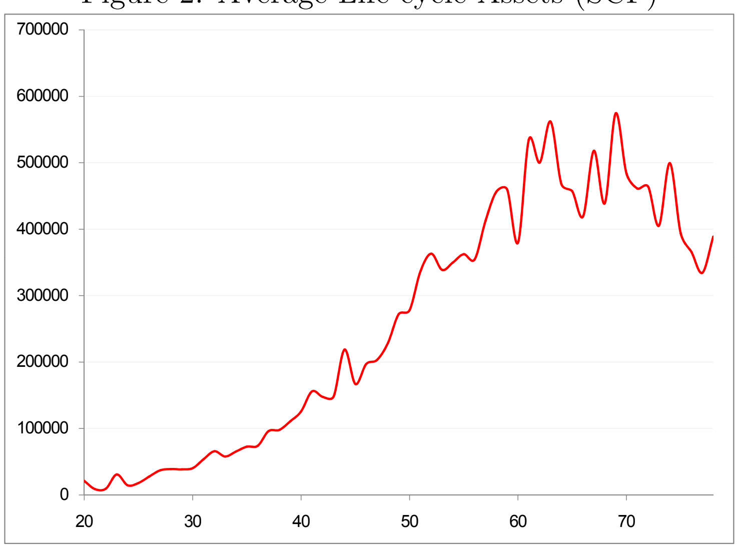

of the lifecycle profiles is qualitatively similar.

Figure 1: Life-cycle earnings statistics

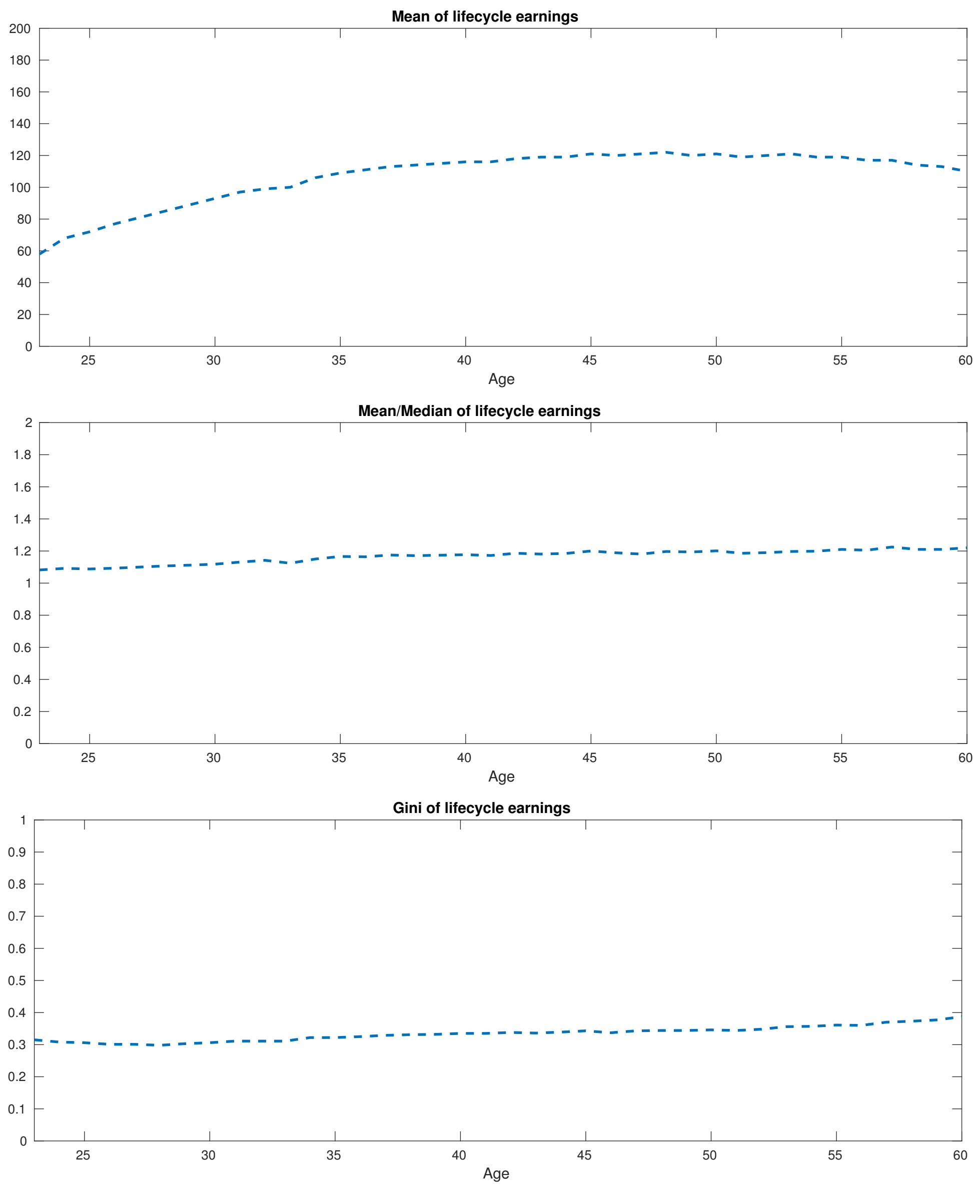

Figure 3: Average Risk-free Assets Conditional on Ownership (SCF)

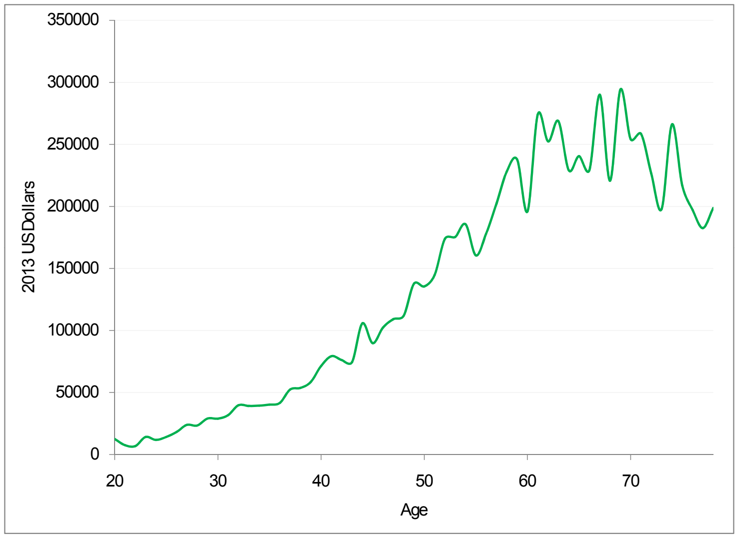

Figure 4: Average Risky Assets Conditional on Ownership (SCF)

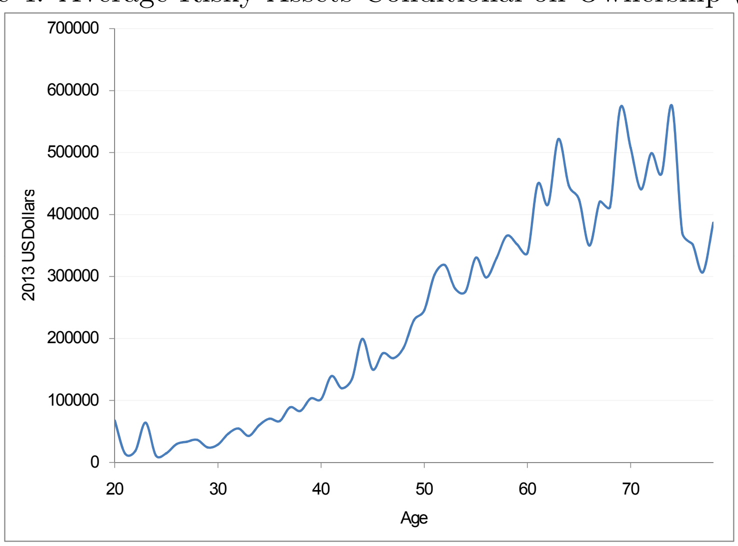

4 Mapping the model to the data

The parameters in our model include: 1) standard parameters such as the discount factor and the coefficient of risk aversion; 2) parameters specific to human capital and to the earnings process; 3) parameters governing the distribution of initial characteristics; 4) parameters specific to college investment and financing; and 5) parameters specific to asset markets. Our approach involves a combination of setting some parameters to values that are standard in the literature, calibrating some parameters directly to data, and jointly estimating the parameters that we do not observe in the data by matching moments using several observable implications of the model. Agents enter the model at age 18 and live for 63 model periods, which correspond to ages 18 to 80.

4.1 Preference Parameters

The per period utility function is CRRA as described in the model section. We set the coefficient of risk aversion, σ, to 5, which is consistent with values chosen in the financial literature. We conduct robustness checks on this parameter by looking at alternative values such the upper bound of $\sigma = 10$ considered reasonable by Mehra and Prescott (1985) as well as lower values such as $\sigma = 3$. The discount factor chosen ($\beta = 0.96$) is also standard in the literature. We set retirement income to be a constant fraction of labor income earned in the last year in the labor market. Following Cocco (2005) we set this fraction to 0.682 both for high school graduates and for those with some college education and to 0.93 for college graduates.

4.2 Human capital

We set the elasticity parameter in the human capital production function, α, to 0.7. Estimates of this parameter are surveyed by Browning et al. (1999) and range from 0.5 to 0.9. As previously noted, the rental rate of human capital in the model evolves according to $w_t = (1 + g_i)^{t-1}$. The growth rate g_i is calibrated to match the average growth rate in mean earnings observed in CPS data. Specifically, we obtain the values 0.0014 for the full sample, 0.0013 for individuals with no college degree, and 0.0065 for college graduates.

4.2.1 Human capital shocks

The stochastic component of earnings, z_{it}, is parameterized following Abbott et al. (2013): $z_{it} = u_{it} + \epsilon_{it}$, where $u_{it} = \rho u_{i,t-1} + \nu_{it}$ is a persistent component with $\nu_{it} \sim N(0, \sigma_\nu^2)$, and $\epsilon_{it} \sim N(0, \sigma_\epsilon^2)$ is a transitory (iid) component. Abbott et al. (2013) estimate the parameters of the z_{it} process using CPS-type wage measures from the National Longitudinal Survey of Youth (NLSY) and obtain, for high-school graduates, $\hat{\rho} = 0.951$, $\hat{\sigma_\nu}^2 = 0.055$, and $\hat{\sigma}_\epsilon^2 = 0.017$.

4.3 Distribution of Initial Characteristics: Assets, Ability and Human Capital

The distribution of initial characteristics (human capital, ability, and assets) is determined by seven parameters. These parameters are estimated to match the evolution of three moments of the earnings distribution over the life cycle: mean earnings, the Gini coefficient of earnings, and the ratio of mean to median earnings. The estimation proceeds as follows. First, for the distribution of assets, we use data from the Survey of Consumer Finances (SCF) to compute the mean and standard deviation of initial assets to be $22,568 and $24,256, respectively (in 2013 dollars). Second, we calibrate the joint distribution of ability and initial human capital to match the key properties of the earnings distribution over the lifecycle reported earlier using the March CPS.

The dynamics of the earnings distribution implied by the model are determined in several steps: i) we compute the optimal decision rules for human capital using the parameters described above for an initial grid of the state variable; ii) we simultaneously compute college and financial investment decisions and compute the lifecycle earnings for any initial pair of ability and human capital; and iii) we choose the joint initial distribution of ability and human capital to best replicate the properties of the CPS data.

We search over the vector of parameters that characterize the initial state distribution to minimize a distance criterion between the model and the data. We restrict the initial distribution to lie on a two-dimensional grid spelling out human capital and learning ability, and we assume that the underlying distribution is jointly log-normal. This class of distributions is characterized by five parameters.[9] We find the vector of parameters $\gamma = (\mu_a, \sigma_a, \mu_h, \sigma_h, \varrho_{ah})$ that characterizes the initial distribution by solving the minimization problem

$$\min_{\gamma} \left(\sum_{j=5}^{J} |log(m_j/m_j(\gamma))|^2 + |log(d_j/d_j(\gamma))|^2 + |log(s_j/s_j(\gamma))|^2 \right)$$

where $m_j, d_j,$ and s_j are the mean, dispersion, and skewness statistics constructed from the CPS data on earnings, and $m_j(\gamma)$, $d_j(\gamma)$, and $s_j(\gamma)$ are the corresponding model statistics.[10]

[9] In practice, the grid is defined by 20 points in human capital and in ability.

[10] For details on the calibration algorithm see Huggett et al. (2006) and Ionescu (2009).

4.4 College Parameters

In terms of parameters specific to college, we set the total college cost, $\overline{d} = \$53,454$ and tuition, $\hat{d} = \$28,320$. The limit and interest rate on student loans are $d_{max} = \$23,000$ and $R_g = 1.07$. We set the wage during college, $w_{col} = \$17,700$ (NCES data). Lastly, the probability of college completion, $\pi(h_5)$ is based on completion rates by cumulative GPA scores in BPS 1996-2001 data set. We define the completion rate as the fraction of students who have earned a bachelor's degree by June 2001. We use the GPA scores cutoffs in the data and the counterpart cutoffs of h_5 as shown below.

Completion rate	Grades
0.07	grades C and D
0.30	mostly Cs
0.45	mostly Bs and Cs
0.56	mostly Bs
0.67	mostly Bs and As
0.70	mostly As

4.5 Financial Markets

We turn now to the parameters in the model related to financial markets. We fix the mean equity premium to $\mu = 0.06$, as is standard (e.g., Mehra and Prescott, 1985). The standard deviation of innovations to the risky asset is set to its historical value, $\sigma_\eta = 0.157$. The risk-free rate is set equal to $R_f = 1.02$, consistent with values in the literature (McGrattan and Prescott, 2000) while the wedge between the borrowing and risk-free rate is $\phi = 0.09$ to match the average borrowing rate of $R_b = 1.11$ (Board of Governors of the Federal Reserve System, 2014).

Borrowing limits in the model will be allowed to vary across households. We introduce heterogeneity in these limits as follows: We first group agents in the model by quartiles of initial human capital, then compute average earnings over the life cycle for each quartile. We then set the borrowing limit for all agents within a quartile to be a given percentage of the average life-cycle earnings for that quartile. We obtain the relevant percentages from the SCF by dividing the sample into income quartiles and calculating the average credit limit as a percentage of the average income within each quartile. The resulting borrowing limits as a percentage of average earnings by quartiles are: 55%, 48%, 35%, and 27%.[11] Lastly, in our baseline model, we assume

[11] We extrapolate the first percentage from the other three rather than calculating it directly because of the large numbers of zeros in the earnings data for the lowest quartile.

that the returns to both risky assets (human capital and financial wealth) are uncorrelated. We assume for the time being that the returns to assets are uncorrelated.

5 Results

As stated at the outset, the aim of this paper is to provide an evaluation of the quantitative importance of key investment opportunities for the temporal resolution of lifetime inequality, with a focus on three specific types of investment: risky and lumpy college; risky, high-return financial assets; and costly borrowing. Our full model, as described above, accommodates the salient features of college education, as well as a rich menu of financial investment opportunities. We will refer to the full model as the "College-Stocks," or "CS", model. We are interested in comparing outcomes from the CS baseline to those of various special cases in which one or more features of the baseline are eliminated in favor of a simpler, more stylized, investment environment.

We begin our analysis with a special case of our model, in which households have access to risky, high-return stocks, but where human capital is accumulated via the much smoother and deterministic process envisioned in the Ben-Porath (1967) model. We refer to this case as the Ben-Porath-with-Stocks, or BPS, case. To measure the role of financial assets and borrowing on the fraction of lifetime inequality that is determined prior to entry into the workforce, we compare BPS outcomes to those of a model with Ben-Porath-style human capital accumulation opportunities, but where financial assets are restricted to nonnegative holdings of risk-free bonds only (we denote the latter the "Ben-Porath-with-Bonds," or BPB, model). The comparison illustrates the role of financial assets in a world with smooth and riskless human capital accumulation.

The next step in our analysis will repeat this exercise for the case in which the process of human capital accumulation reflects the salient features of empirically relevant representations of college: That is, we will compare the CS to the CB ("College-with-Bonds") model. The two previous comparisons should provide measures of the role of financial assets under the two polar structures for human capital accumulation that we consider. [Results from college model to be added.]

To better understand the role of lumpy and risky human capital accumulation (all else equal), we then derive two additional and analogous sets of results. First, we compare the baseline, i.e., CS, model to the BPS model. This holds financial investment opportunities fixed but alters the risk and lumpiness of human capital accumulation. Taken as a whole, these calculations give us measures for the role played by financial investment opportunities under both models of human capital accumulation opportunities. And next, we turn to the question of whether the interaction between financial investment choices and human capital accumulation options is quantitatively important. We do this by comparing the changes obtained from moving from the CS to the BPS setting, against those obtained from moving from the CB to the BPB setting. That is, does the array of financial assets available matter for the impact of risk and lumpiness in human capital accumulation?

5.1 Financial investments and sources of lifetime inequality

5.1.1 Model fit

We start by presenting measures of the goodness of fit for the BPS model. Table 1 presents key properties of the initial distribution of ability and human capital found to best fit the 114 earnings moments in our CPS data (mean, skewness, and Gini of earnings for each age in 23-60). Figure 5 shows the earnings moments for a simulated sample of individuals in the model (for this initial distribution) versus the CPS data. We obtain a fit of 11% (where 0% represents a perfect fit).[12] As the figure shows, the model does a reasonably good job of fitting the evolution of mean earnings over the lifecycle, though the model's profile is a bit less hump-shaped than in the data. The skewness of earnings is a touch lower in the model than in the data. And, for the Gini coefficient, the model matches the data quite well, except perhaps in the last few years of the lifecycle. Since we are interested in explaining the level of economic inequality over the lifecycle, this last statistic is particularly important. The bottom figure suggests that the model does a good job of matching the level of inequality.

Table 1: Properties of the initial distribution

Statistic	BPS model
Mean learning ability a	0.40
Coefficient of variation a	1.72
Mean initial human capital h_1	95
Coefficient of variation h_1	0.42
Correlation (a, h_1)	0.79

We next test the model predictions for key non-targeted data moments. Figure 6 shows the mean wealth accumulation over the lifecycle for both risky and risk-free assets, while Figure 7 shows participation in the stock market. Overall, the model is quite consistent with the observed financial investment behavior. In particular, despite not being targeted, our model produces empirically consistent estimates of lifecycle wealth and its allocation between risky and risk-free

[12] As a measure of goodness of fit, we use $\frac{1}{3J}\sum_{j=5}^{J}|log(m_j/m_j(\gamma))| + |log(d_j/d_j(\gamma))| + |log(s_j/s_j(\gamma))|$. This represents the average (percentage) deviation, in absolute terms, between the model-implied statistics and the data.

Figure 5: Life-cycle earnings statistics

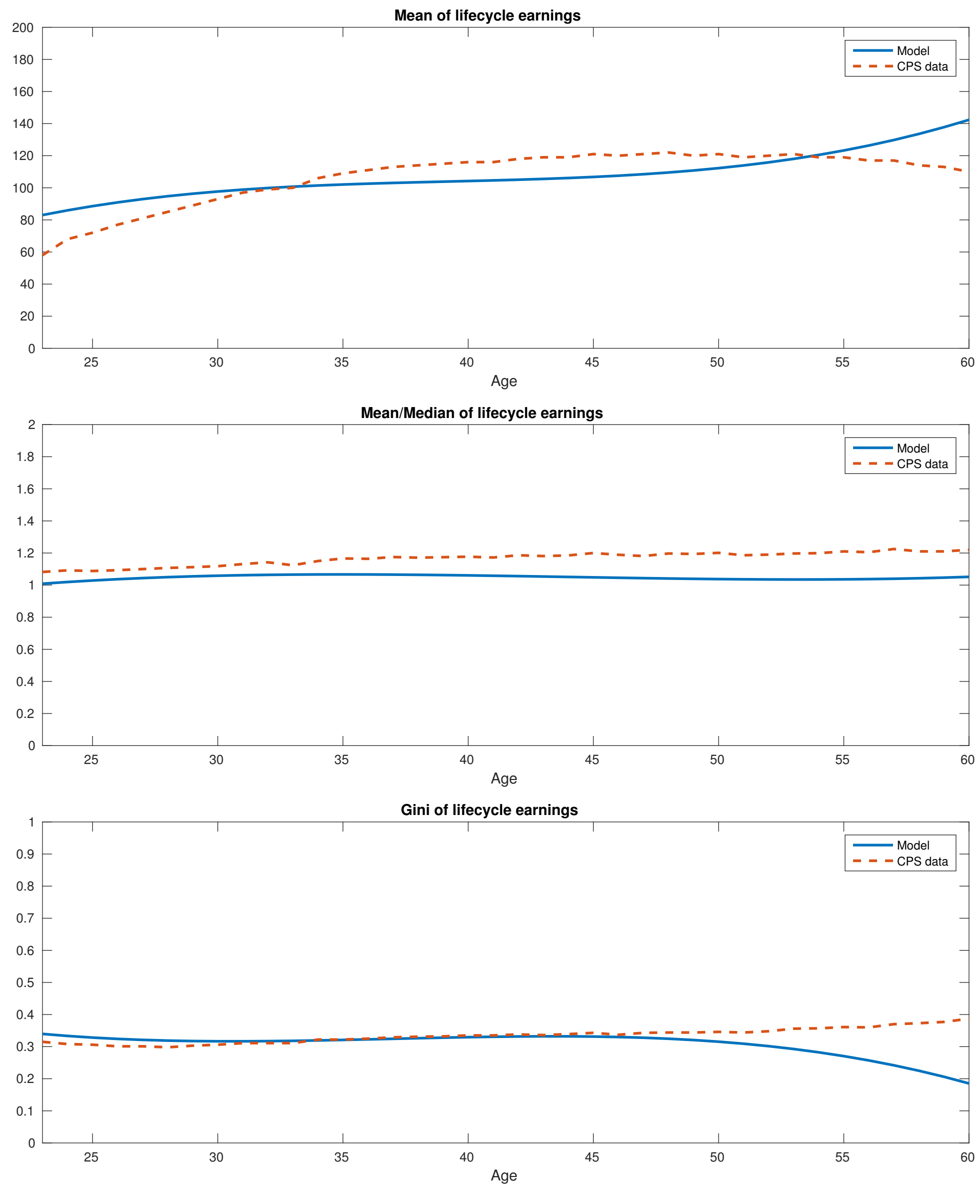

assets. In addition, our model's prediction for the stock-market participation rate is consistent with the data, over the *entire* lifecycle. This result is driven primarily by the presence of human capital investment in our model: time spent on human capital accumulation is at its highest early in life. For instance, at age 23, households spend about a third of their time endowment on human capital accumulation. During the early part of life, we see also that only around 30% of all households participate in the stock market. Diminishing returns, and a shorter horizon to recoup the investment, imply that human capital accumulation falls with age. As this occurs, we see that stock-market participation steadily increases, reaching around 90% at retirement age. As retirement approaches, we see that the fraction of time allocated to human capital falls sharply, reaching less than 0.05 by retirement age.

5.1.2 Initial conditions vs. shocks

We next decompose measures of lifetime inequality (starting with lifetime earnings and lifetime wealth) into the contribution of initial conditions in the model versus that of shocks.[13] Our decomposition follows Huggett et al. (2011). Specifically, we first simulate each of the relevant models for a large number of individuals. For each simulated individual, we compute lifetime earnings (calculated as the present value of earnings realized over ages 23-60) and lifetime wealth (calculated as lifetime earnings plus the value of initial wealth), and we then compute the cross-sectional variance (across the simulated individuals) of lifetime earnings and lifetime wealth. Next, we simulate the model again, but this time we set the variance of the human capital shocks to zero, adjusting parameter μ in order to keep the mean shock level constant (we do not change the distribution of initial conditions), and recompute the variance in lifetime earnings and lifetime wealth. The resulting variance, expressed as a fraction of the variance in the first simulation, or f_s^{var}, is the contribution of the shocks to inequality in lifetime earnings and wealth. The fraction of the variance due to variation in initial conditions, f_{ic}^{var}, is set to $f_{ic}^{var} = 1 - f_s^{var}$.

As shown in Table 2, the initial conditions in our calibrated BPS model account for about 96% of the variance in lifetime earnings and about 97% of the variance in lifetime wealth. In the BPB model, initial conditions account for 93% of the variance in lifetime earnings and 95.5% of the variance in lifetime wealth. Note that for both of these models, the contribution of the initial conditions to lifetime inequality in earnings and wealth is a fair bit larger than that found by

[13] In the current version of the draft, we measure initial conditions as of age 23 only. Ultimately, we will also be interested in the role of initial conditions as of age 18. That way, and as noted in the Introduction, our goal is to open the "black box" of what is, in a model without college, reflected in initial conditions as of age 23. Future versions of the draft will also report results for two other measures of lifetime inequality: inequality in lifetime consumption and lifetime utility.

Figure 6: Life-Cycle Wealth Accumulation

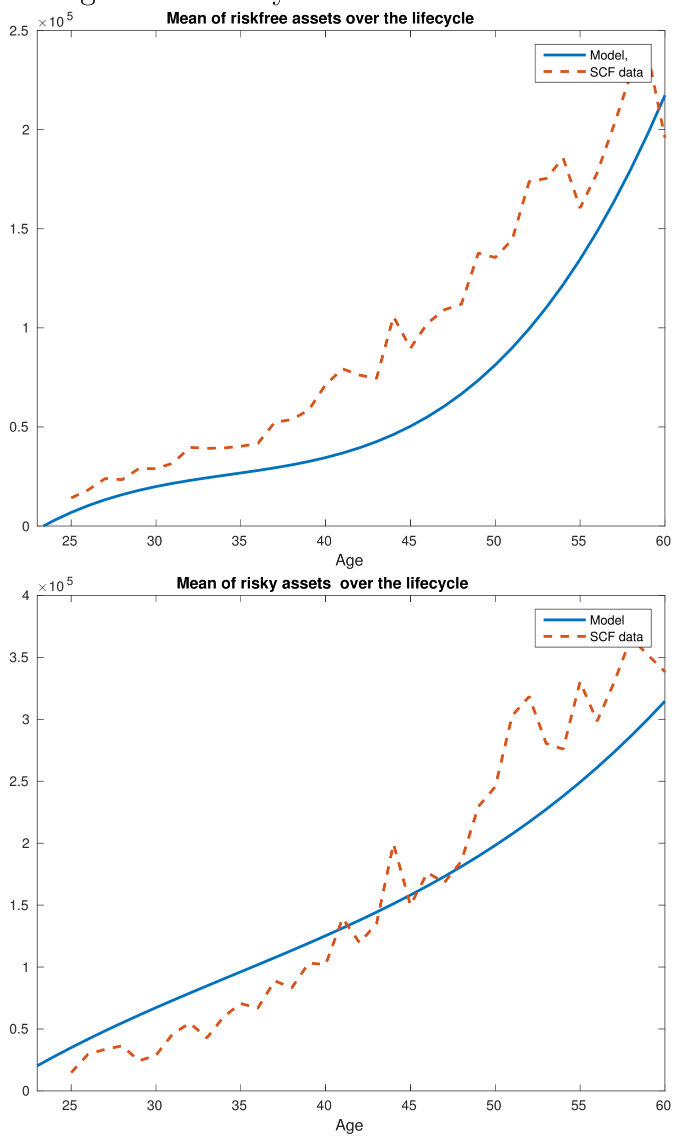

Figure 7: Investment in stocks over the life-cycle

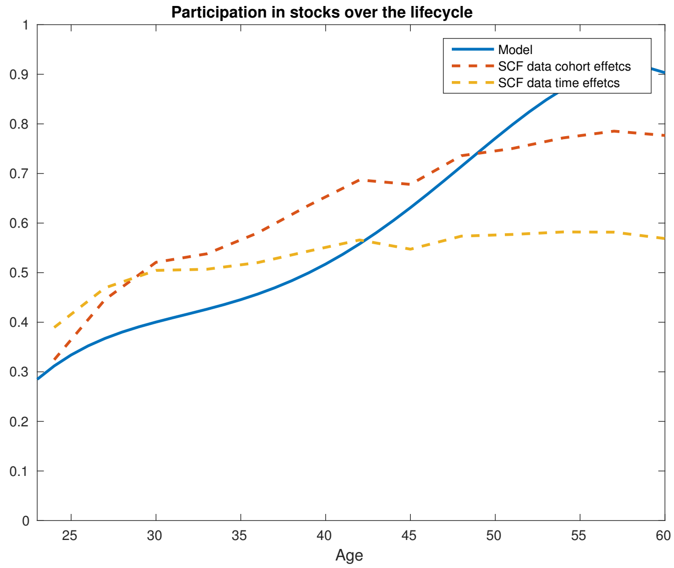

Huggett et al. (2011) or Storesletten et al. (2005), but roughly similar to that found by Keane and Wolpin (1997). One important reason for the large role of initial conditions is that we obtain a very unequal initial distribution of both initial human capital and ability (e.g. the ability distribution we find is substantially more unequal than that found by Huggett et al. (2011)).[14]

Table 2: Sources of lifetime inequality

Fraction due to all initial conditions	BPS Model	BPB Model
of variance in lifetime earnings	95.6	93.2
of variance in lifetime wealth	97.2	95.5

5.1.3 The role of risky financial assets for lifetime inequality

Comparing the BPS and BPB models indicates that the relative importance of shocks to human capital (relative to initial conditions) increases somewhat in the economy where investment in risky assets is not allowed (in the BPB model, the only risky investment in the economy is human capital). The high-return risky asset allows some individuals to get rich (in particular, the rich to get richer), but its absence shuts down the trade-off between investments in stocks and human capital, which is more pronounced for individuals with relatively high ability and low initial human capital.

[14] Apart from the role played by formal education, which is embodied in initial conditions as of age 23, in particular in the form of the human capital stock at this age, we are exploring other mechanisms that contribute to the large role of initial conditions.

Indeed, Figure 8 shows that in the BPS model, individuals who are more likely to invest in human capital (and in particular early on in life) are less likely to participate in the stock market. As shown, when households are young, non-participants have substantially higher ability levels than stock-market participants. It is only in middle age and beyond (as seen in the figure) that ability is similarly distributed across stock-market participants and non-participants. By middle age, as households have accumulated levels of human capital consistent with their innate ability, marginal returns to human capital are no longer substantially higher than the returns on stocks for even those with high innate ability.

At the same time, initial human capital levels are lower for individuals who do not participate in the stock market relative to those who participate, because individuals with low initial human capital levels have higher incentives to invest in human capital early on given the low opportunity cost of this investment. But once individuals accumulate enough human capital by middle age this relationship is reversed. When the opportunity to invest in risky assets is shut down, individuals with high enough ability and low enough initial human capital levels invest even more in human capital, and as a result allocate less time to market work in addition to not investing in stocks. These channels work towards lowering the importance of initial conditions for lifetime inequality (recall that the correlation between ability and initial human capital is a positive and large 0.8).

We next decompose the contribution of initial conditions to lifetime inequality into the relative contributions of h_1, a, and x_1. To do this, we first simulate our model economy without shocks and compute the cross-sectional variance of lifetime earnings and wealth. Next, we repeat the simulation, but each time setting one of h_1, a, and x_1 (one at a time), to their median value, and we recompute the cross-sectional variance of lifetime earnings and wealth. For each of the latter simulations, we report the variance of lifetime earnings as a fraction of the variance in the former simulation. The fraction of the variance in lifetime earnings and wealth due to each of h_1, a, and x_1 is reported in Table 3.

Table 3: Sources of lifetime inequality

Fraction due to h_1	BPS Model	BPB Model
of variance of lifetime earnings	0.73	0.70
of variance of lifetime wealth	0.60	0.59
Fraction due to a		
of variance of lifetime earnings	0.42	0.40
of variance of lifetime wealth	0.21	0.22
Fraction due to x_1		
of variance of lifetime earnings	0.002	0.002
of variance of lifetime wealth	0.48	0.486

Figure 8: Ability Distribution of Participants and Non-Participants

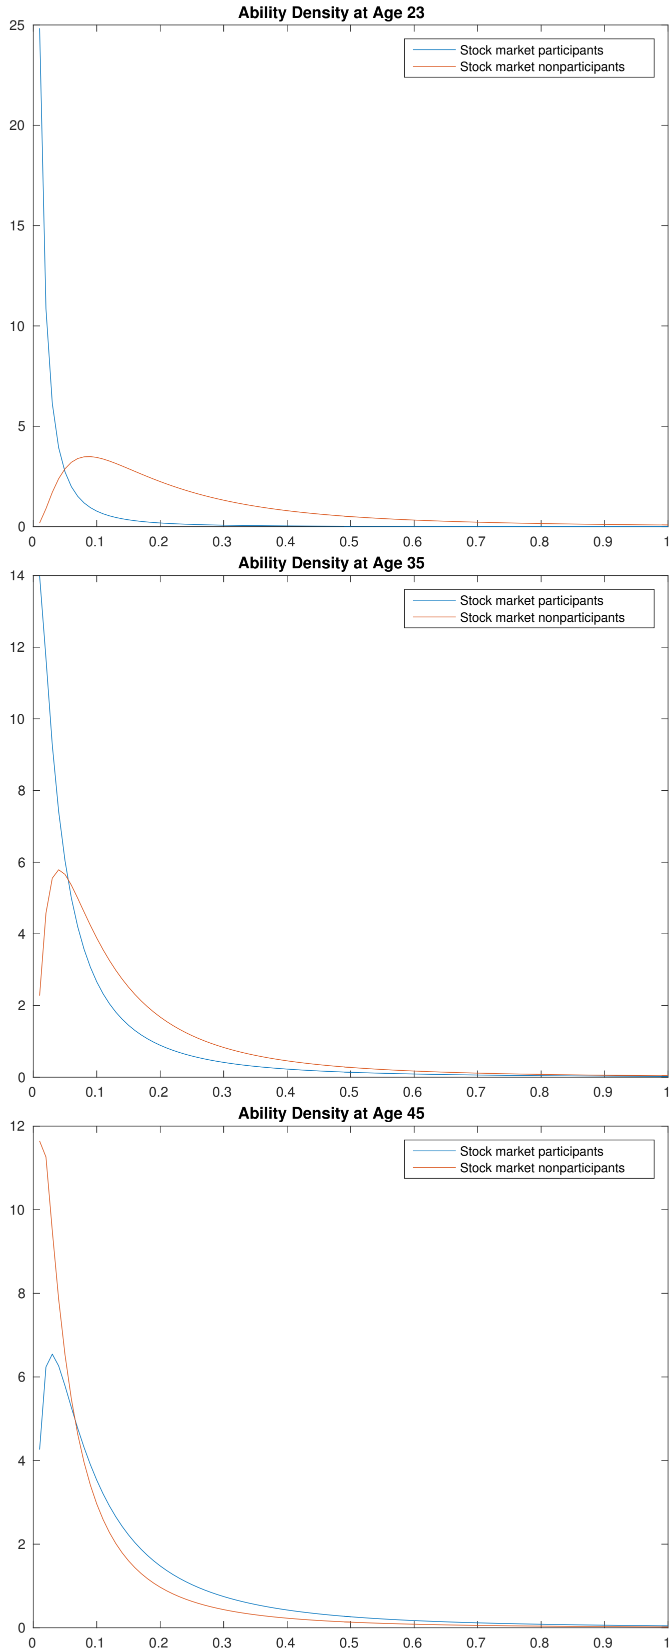

Our results indicate that both initial human capital and ability are quantitatively important for lifetime inequality as measured by both the fraction of the variance in lifetime earnings and in lifetime wealth. This is because both characteristics shape the incentives to invest in human capital; in particular individuals with low initial human capital levels and high ability levels have the highest incentive to invest in human capital early on. This behavior in turn has implications for lifetime earnings but also for the financial investment behavior of these individuals. Indeed, as shown in Figures 9 and 10, investment in stocks over the lifecycle varies greatly by initial human capital and ability.

In particular, individuals with the highest levels of initial human capital (quartile 4) participate in the stock market at the highest rates, while those with the lowest levels participate at by far the lowest rates. Specifically, participation within the top quartile is about 70% at age 23 and reaches 100% participation by age 55 (Figure 9). Quartiles 2 and 3 participate at around a 20%-30% rate early in life, and reach 100% participation after age 57. For the lowest quartile, participation starts at around 10% and remains below 60% throughout the working life. This behavior is a direct consequence of the trade-off between incentives to invest in human capital and stocks that these individuals face. Specifically, the model delivers that time allocation as a function of initial human capital is inversely proportional to its initial level: Those in quartile 1 (the lowest level of initial human capital) invest the most time, while those in the highest quartile invest the least. The intuition is natural. Those with high initial human capital face not only a high opportunity cost of additional accumulation, but also stand to reap only low marginal returns. The reverse holds for those with low initial human capital. The absence of the opportunity to invest in stocks reduces this variation in wealth accumulation behavior across individuals with different initial human capital levels, and as a result its importance for lifetime inequality is relatively lower in the BPB model versus the BPS model.

When it comes to the change in the importance of the ability level for lifetime inequality in the absence of the opportunity to invest in stocks, results are a bit more mixed as shown in Table 3. In particular, the importance of ability slightly increases in the BPB model when measured as a fraction of the variance of lifetime wealth but it increases more when it is measured as the fraction of the variance of lifetime earnings. This is because the model delivers that while learning ability is directly related to earnings over the life cycle, learning ability, *all else equal*, is inversely related to equity investment. As shown in Figure 10, households in the lowest quartile of ability participate at very high rates in the stock market relative to those in other quartiles. The intuition is simply that for low-ability households, the effective rate of return from human capital is much lower than from equity investment. Further, their earnings profile is relatively flat, which means that their participation rate also remains flat over the life cycle. In contrast, the high initial investment in human capital, particularly for quartiles 2 and 3, and the steeper earnings profile, particularly for quartile 4, is associated with these groups exhibiting a steeply increasing stock market participation rate over the life cycle. For these households, learning, especially when young, is a better investment than earning and investing in equities.

At the same time, the contribution of initial wealth to lifetime inequality as measured by the fraction of the variance in lifetime earnings due to x_1 is insignificant. This is because unlike ability and initial human capital, initial wealth does not shape the incentives to invest in human capital and consequently does not affect the time allocated to work, which is what generates labor earnings. Initial wealth, on the other hand, does contribute quite a bit to lifetime inequality as measured by the fraction of the variance in lifetime wealth since initial asset positions are relevant for wealth accumulation over life.

Figure 9: Investment in stocks by quartiles of initial human capital

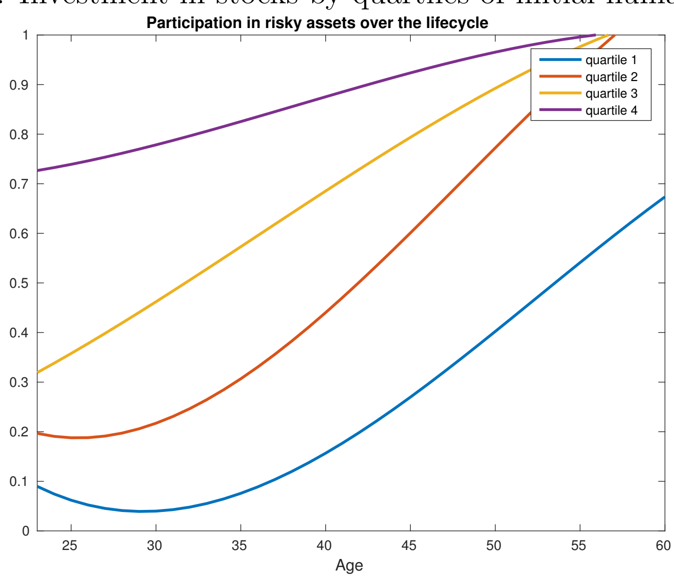

Figure 10: Investment in stocks by quartiles of ability

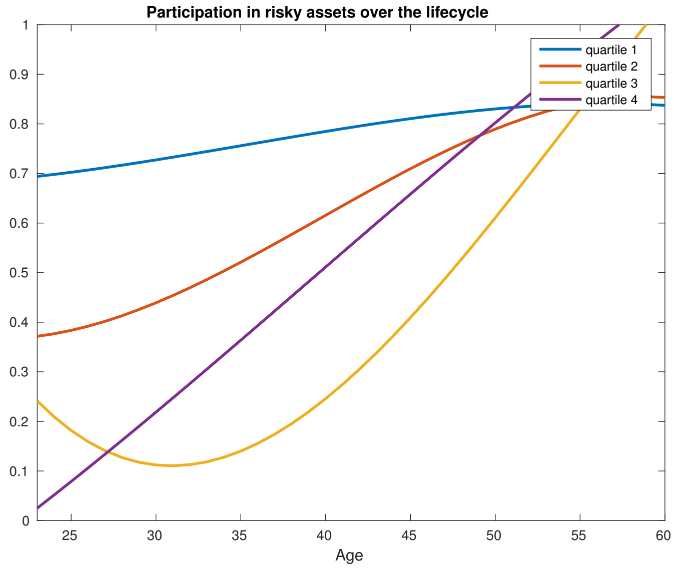

5.1.4 The role of credit and lifetime inequality

When it comes to wealth and income (as opposed to earnings) inequality, a natural question is the extent to which borrowing constraints matter. Early in life, borrowing constraints may significantly affect human capital investment. Later in life, borrowing constraints may affect the ability of households to achieve desired financial investment positions. Moreover, because households who vary in their human capital attainment will also vary in their earnings, households' willingness to bear risk will differ in the population. Thus borrowing constraints may systematically change the distribution of risk-bearing capacity in the economy. Therefore, if access to credit is unequal, so too will be lifetime outcomes, not just for earnings but also for wealth.

We study the role of credit in a model where households may invest in both stocks and bonds. Specifically, we run a version of our BPS model where we shut down the possibility to borrow and a version of it where credit is cheap (the interest rate for credit is set to the riskfree interest rate). As shown in Table 5, we find that access to credit lowers the contribution of (age-23) initial conditions by 5 percentage points. However, less expensive credit does not affect the role played by initial conditions for lifetime inequality.

We find that in the model people borrow primarily to smooth consumption while they invest in human capital rather than to invest in stocks. Therefore credit unavailability hinders households' efforts to smooth consumption intertemporally. This in turn diminishes the benefits to human capital accumulation as households' living standards can only rise once human capital payoffs are realized. This effect is particularly strong for households with low initial human capital levels who are the ones borrowing the most in the face of steep investment profiles in human capital, and therefore earnings. Taking away the option to borrow will therefore increase the importance of initial conditions, in particular that of initial human capital.

At the same time, cheap credit could induce borrowers to invest more in both human capital accumulation and stocks. We find that households only do the former: households who borrow allocate more time to human capital investment and use credit to smooth consumption in the face of lower current (and potentially higher future) earnings. This in turn, will lower the importance of initial conditions, but the quantitative effect is small given that households do not significantly change their stock market investment behavior despite having access to cheaper credit.

Table 4: Role of credit in the BPS Model

Fraction due to all initial conditions	No borrowing	Cheap credit
of variance in lifetime earnings	99.3	95.4
of variance in lifetime wealth	99.4	97

Our findings suggest that credit availability has important implications for the temporal resolution of lifetime inequality. Credit plays a significant role for the investment in human capital, in particular for those individuals for whom this investment gives the higher returns and/or supposes a relatively low opportunity cost. Consequently, an immediate question that arises is how does the role of credit change when we account for a richer, more realistic credit program to invest in human capital. In particular, we plan to study the implications of the current student loan program and college-debt financing for lifetime inequality and for the role played by initial conditions. Our approach and findings are discussed in the following section.

5.2 College investment and sources of lifetime inequality

To understand the role played by investment in a college education, we proceed in two steps. First, we derive results for outcomes in a setting where human capital is accumulated in college, i.e. we compute the CS model. Second, we compare outcomes, and especially, the decomposition of the resolution of lifetime inequality in this instance, to the BPS model, where human capital is more smoothly and risklessly accumulated. Note that our CS model includes a rich characterization of formal education that captures several key features of the college investment. Specifically, recall that we model college as an investment opportunity that includes the following features: college is a time-consuming, lumpy and risky project, which individuals may use personal assets or costly borrowing to finance. This richness allows us to discipline the model to match an array of facts on both college enrollment and completion.

More importantly, the exercise we propose informs us as to the role played by the particulars of real-world human capital investment opportunities: riskiness, lumpiness, and non-defaultable debt instruments used to finance it. First, with respect to the role played by the risky, lumpy human capital investment in college, we aim to disentangle the effects of two specific forces. On the one hand, college allows people to catch up by greatly enhancing their human capital in one step via college, so initial conditions may matter less (in particular that of h_1). On the other hand, the risk of an investment in college may discourage some people from making this investment and this would increase the importance of initial conditions, in particular that of initial human capital and financial wealth, which may interact particularly strongly. For instance, wealth-poor households will have a relatively lower willingness to take the risk of college than their better-heeled counterparts–even if both have the same initial levels of ability and human capital. The model allows us to quantify the net effect and compare it to that of financial assets for lifetime inequality.

With respect to capturing the effect of college-financing instruments, we note first that households frequently finance college (and their consumption while attending) by reducing assets or taking on nondefaultable debt. These factors are likely to put downward pressure on household investment in stocks, and hence, on future wealth. They are, however, partially offset by positive effects on lifetime labor income, as those who obtain a college degree tend to earn higher earnings and face lower unemployment risk. As a whole, these features of college can each be expected to affect the impact of any pre-college distribution of "innate" ability and "initial" human capital. We plan to explore these issues in versions of our CS model where borrowing for college takes different forms relative to the current student loan program and assess the relevance of these alternative specifications of student loan lending for the importance of initial conditions for lifetime inequality. Specific arrangements that we have in mind are different eligibility criteria for student loans, college subsidies, and a richer menu of repayment options that allow for different levels of contingency, such as income-contingent repayment plans as the ones recently proposed by the Administration. We aim to derive policy-relevant implications regarding student loan subsidies, need-based aid, and the choice of regime governing the consequences of student loan default. When combined with what will be a more accurate assignment to pre-college heterogeneity in ability and human capital, our approach has promise for delivering more accurate implications not only for the temporal resolution of uncertainty and inequality, but also for the implications of the policies noted above.

[TBC]

6 Conclusion

Previous research has found that a significant fraction of inequality is determined early in life, with estimates ranging from half to as much as 90 percent. Our goal in this paper is to quantify the contribution of three investments—human capital, stocks, and credit—to the early-in-life resolution of income, wealth, and consumption inequality. Consistent with previous work, we find that initial conditions contribute significantly to variations in lifetime earnings and wealth.

When we unpack the roles played by the three investments that we consider we find, first, that around 70 percent of income inequality and 60 percent of wealth inequality is attributable to variations in human capital as of age 23. Innate ability plays a relatively smaller, but still significant role, contributing to about 40 percent of earnings inequality and 20 percent of wealth inequality. These results are not surprising in a setting such as ours where households have the ability to influence their lifetime earnings through early investments in human capital, and those with high ability and low initial human capital, in particular, have a high incentive to do so.

The contribution of initial ability and human capital to income and wealth inequality is slightly higher in a model where we also allow agents to invest in stocks, which adds a second risky, high-return asset (the first being human capital) to their investment choices. Individuals with the highest opportunity cost of investing in human capital, namely those with relatively high levels of initial human capital, choose to invest in stocks instead starting early in life. Because they have high initial human capital, these individuals also have relatively high earnings and wealth. Stocks offer them an additional investment option by which to accumulate more wealth. When the option to invest in stocks is absent, this channel is shut down, which lowers the importance of initial human capital for lifetime inequality.

Given the role played by human capital as of age 23, it becomes especially important to measure the extent to which the investment in college contributes to the accumulation of human capital by this age. We find that [results to be added.]

References

Abbott, B., G. Gallipoli, C. Meghir, and G. L. Violante (2013, February). Education Policy and Intergenerational Transfers in Equilibrium. Working Paper 18782, National Bureau of Economic Research.

Akyol, A. and K. Athreya (2005). Risky Higher Education and Subsidies. *Journal of Economic Dynamics and Control 29*(6), 979–1023.

Altonji, J. G., A. A. Smith Jr., and I. Vindangos (2013). Modeling Earnings Dynamics. *Econometrica 81*(4), 1395–1454.

Autor, D., F. Levy, and R. Murnane (2003). The Skill Content of Recent Technological Change: An Empirical Exploration. *Quarterly Journal of Economics 118*(4), 1279–1333.

Ben-Porath, Y. (1967). The Production of Human Capital and the Life-cycle of Earnings. *Journal of Political Economy 75*(4), 352–365.

Board of Governors of the Federal Reserve System (2014). Federal Reserve Statistical Release, G19, Consumer Credit. Technical report, Board of Governors of the Federal Reserve System.

Bound, J., M. F. Lovenheim, and S. Turner (2010). Why Have College Completion Rates Declined? An Analysis of Changing Student Preparation and Collegiate Resources. *American Economic Journal: Applied Economics 2*(3), 129–57.

Bowen, W. G., M. M. Chingos, and M. S. McPherson (2009). *Crossing the Finish Line: Completing College at America's Public Universities*. Princeton University Press.

Braxton, J., A. Hirschy, and S. McClendon (2003). *Understanding and Reducing College Student Departure*. Jossey-Bass, San Francisco.

Browning, M., L. Hansen, and J. Heckman (1999). *Handbook of Macroeconomics*. Elsevier Science B.V Amsterdam.

Cagetti, M. (2003). Wealth Accumulation over the Life Cycle and Precautionary Savings. *Journal of Business & Economic Statistics 21*(3), 339–353.

Card, D. (1999). The Causal Effect of Education on Earnings . Volume 3, Part A of *Handbook of Labor Economics*, pp. 1801 – 1863. Elsevier.

Caucutt, E. M. and K. B. Kumar (2003). Higher Education Subsidies and Heterogeneity: A Dynamic Analysis. *Journal of Economic Dynamics and Control 27*(8), 1459–1502.

Caucutt, E. M., L. Lochner, and Y. Park (2015). Correlation, Consumption, Confusion, or Constraints: Why do Poor Children Perform so Poorly? Working paper, National Bureau of Economic Research.

Cocco, J. (2005). Portfolio choice in the Presence of Housing. *Review of Financial studies 18*(2), 535–567.

Garriga, C. and M. Keightley (2007). A General Equilibrium Theory of College with Education Subsidies, In-school Labor Supply, and Borrowing Constraints. Working paper series 2007-051a, Federal Reserve Bank of St. Louis.

Goldin, C. and L. F. Katz (2007). Long-Run Changes in the Wage Structure: Narrowing, Widening, Polarizing. *Brookings Papers on Economic Activity* (2), 135.

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

Heckman, J. J. and A. B. Krueger (2005). *Inequality in America: What Role for Human Capital Policies?* MIT press.

Hendricks, L. and O. Leukhina (2014). The Return to College: Selection Bias and Dropout Risk.

Hubbard, R., J. Skinner, and S. Zeldes (1994). Expanding the Life-cycle Model: Precautionary Saving and Public Policy. *The American Economic Review 84*(2), 174–179.

Huggett, M., G. Ventura, and A. Yaron (2006). Human capital and earnings distribution dynamics. *Journal of Monetary Economics 53*(2).

Huggett, M., G. Ventura, and A. Yaron (2011). Sources of lifetime inequality. *The American Economic Review 101*(7), 2923–2954.

Hungerford, T. and G. Solon (1987). Sheepskin Effects in the Returns to Education. *The Review of Economics and Statistics 69*(1), pp. 175–177.

Ionescu, F. (2009). The Federal Student Loan Program: Quantitative Implications for College Enrollment and Default Rates. *Review of Economic Dynamics 12*(1), 205–231.

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

Lemieux, T. (2010). What Do We Really Know about Changes in Wage Inequality? In *Labor in the New Economy*, pp. 17–59. University of Chicago Press.

McGrattan, E. and E. Prescott (2000). Is the Stock Market Overvalued? *Federal Reserve Bank of Minneapolis Manuscript*.

Mehra, R. and E. C. Prescott (1985). The Equity Premium: A Puzzle. *Journal of Monetary Economics 15*(2), 145–161.

Mincer, J. (1991, September). Education and Unemployment. Working Paper 3838, National Bureau of Economic Research.

Ozdagli, A. K. and N. Trachter (2011). On the Distribution of College Dropouts: Household Wealth and Uninsurable Idiosyncratic Risk. Working paper series, Federal Reserve Bank of Boston.

Stinebrickner, R. and T. R. Stinebrickner (2003). Working During School and Academic Performance. *Journal of Labor Economics 21*(2), pp. 473–491.

Storesletten, K., C. I. Telmer, and A. Yaron (2005). Consumption and Risk Sharing over the Life Cycle. *Journal of Monetary Economics 52*(7), 1331–1358.