

Who Values Access to College?*

Kartik Athreya[†] Felicia Ionescu[‡] Urvi Neelakantan[§] Ivan Vidangos[¶]

February 11, 2019

Abstract

A first glance at US data suggests that college—given its mean returns and sharply subsidized cost for all enrollees—could be of great value to most. Using an empirically-disciplined human capital model that allows for variation in college-readiness, we show otherwise. While the top-decile of valuations is indeed large (40 percent of consumption), nearly half of high school completers place *zero* value on access to college. Subsidies to college currently flow to those already best positioned to succeed and least sensitive to them. Even modestly targeted alternatives may therefore improve welfare. As proof of principle, we show that redirecting subsidies away from those who would nonetheless enroll—towards a stock index retirement fund for those who do not even when college is subsidized—increases ex-ante welfare by 1 percent of mean consumption, while preserving aggregate enrollment and being budget neutral.

JEL Codes: E21; G11; I24;

Keywords: Human Capital; Higher Education; Financial Investment

*This paper arises from feedback we received at presentations of our companion paper, “Investment Opportunities and the Sources of Lifetime Inequality,” urging us to separate from the latter the detailed measurement of the value of access to college across US households and the attendant implications for public support. We therefore thank, without implicating, numerous seminar and conference participants and discussants of both papers. We are especially grateful to Gene Amromin, Joao Cocco, Andrew Figura, Francisco Gomes, Oksana Leukhina, Monica Paiella, Luigi Pistaferri, Sangeeta Pratap, and Paolo Sodini for detailed discussions and input. The views expressed in this paper are those of the authors and do not necessarily reflect the views of the Federal Reserve Board, the Federal Reserve Bank of Richmond, or the Federal Reserve System. All errors are ours.

[†]FRB Richmond. P.O. Box 27622, Richmond, VA 23261, kartik.athreya@rich.frb.org, Ph:804-697-8225

[‡]Federal Reserve Board. felicia.ionescu@frb.gov, Ph:202-452-2504,

[§]CAFRAL and FRB Richmond. urvi.neelakantan@cafral.org.in, Ph:+91-22-22694587

[¶]Federal Reserve Board. ivan.vidangos@frb.gov, Ph:202-452-2381

1 Introduction

A first glance at US data suggests that college could be of great value to most, if not all, individuals. The average returns to college completion are high: in recent decades, the lifetime income of college graduates exceeded that of high school graduates by a factor of roughly two (Goldin and Katz, 2007).¹ College graduation can thus substantially increase individuals' lifetime consumption and, presumably, utility.

However, two important features of the data work against this interpretation. First, and perhaps most importantly, college graduation is far from guaranteed. College noncompletion is in fact widespread (Restuccia and Urrutia, 2004; Bound et al., 2010; Johnson, 2013); in recent data, roughly half of all enrollees into public four-year colleges (which absorb roughly two-thirds of all Bachelor's degree enrollees) failed to complete college even eight years after initial enrollment (Bound et al., 2010). In addition, partial completion offers relatively little reward. Second, even conditional on successful college completion, subsequent returns—as embodied in labor earnings—are not only subject to uninsurable aggregate risk as with many diversified financial investments, but also to substantial uninsurable idiosyncratic risk.

Empirically, the ability to successfully complete college depends strongly on *individual-level* characteristics, chiefly those summarizing the preparation with which individuals enter college. This preparation governs both the ability to effectively accumulate human capital (e.g., innate capabilities and being in environs conducive to learning) and the amount of human capital (e.g., math and language literacy) accumulated through the end of high school.² The variability in these characteristics makes the returns to college—whether high or low—fairly predictable for some and uncertain for others. Importantly, for a significant proportion of high school graduates, the ex-ante attractiveness of college may be far lower than the ex-post return suggests, and *predictably* so.

College is also heavily subsidized. In particular, college in the US receives direct subsidization that lowers the cost of attendance for all enrollees, irrespective of their backgrounds. Empirical estimates of the size of this subsidy are large, reducing the direct cost of college by 40 to 50 percent (Athreya and Eberly, 2013). The dependence of college completion on initial preparation immediately suggests that direct subsidies to college may be of widely varying value to high school completers. In particular, the lion's share of the benefits from subsidies may flow to those who already place high value on college enrollment even in their absence.

Motivated by these observations, the first contribution of this paper is to provide a measure

¹This is also consistent with the large empirical literature that attempts to estimate the causal effect of education on earnings and finds very substantial returns. See Card (2001) for a review.

²Hendricks and Leukhina (2017) show that college preparation, as measured by transcript data, is a strong predictor of graduation prospects.

of the value of access to college across various types of individuals. Our goal here is to answer the question “*Who* benefits from access to college?” We do this by computing the value of access to college across agents who differ in characteristics, namely ex-ante preparedness and the ability to finance college. Specifically, we construct a model in which agents who differ ex-ante in their initial human capital, their ability to further accumulate it, and their initial financial wealth decide whether to enroll in college.³ We then use data on life-cycle earnings and college investment behavior to back out—following the approach pioneered in Huggett et al. (2006) and extended in Ionescu (2009)—the joint distribution of ability, initial human capital, and initial wealth. With the joint distribution of initial attributes in hand, we use the model to measure the value of access to college across the spectrum of individual types.

Two findings emerge from this exercise. First, nearly half the population assigns little value (less than 5 percent of consumption) to college. This is striking but emerges naturally from our model and reflects the large noncompletion rate even among those who currently enroll. Second, this low value to many individuals contrasts with the extremely high value that a few enjoy: the top decile of valuations exceeds 40 percent of consumption, pulling up the mean value of college access, in consumption-equivalent terms, to 14.7 percent. The valuations differ starkly across measures of college readiness; for example, in our model, those in the bottom quartile of the ability distribution place little value on college access (only 1 percent of consumption) while the valuations of the top quartile near 30 percent of consumption.

Next—and this is the second contribution of the paper—we assess the extent to which, all else equal, the large subsidies currently in place for those who enroll in college matter. Public support for college attendance takes many forms, including need- and merit-based grants, subsidized student loans, and, perhaps most importantly, subsidies that *directly* and substantially lower the tuition charged by public schools, especially in-state. It is important to note that the latter increases college affordability for *all* enrollees, with no reference to individual attributes. Is this direct subsidy important? We find that the answer is yes: removing the direct subsidy reduces the value of access to college at the mean by more than one-third.

The analysis described above lets us classify the population of high school completers into three groups: those who enroll in college with or without the subsidy (whom we label “always enroll”

³The ability to accumulate human capital—henceforth, “ability”—must be interpreted with great care. It is emphatically not to be taken as any sort of innate measure of intelligence. Rather, in our model it simply represents the totality of forces that govern the effectiveness with which an individual—at the end of high school—can turn *time* into human capital. Needless to say, such effectiveness will be affected by (1) the multitude of experiences prior to completing high school, especially the investments made in their human capital by parents, the school system, and the community at large, as well as (2) any unmodeled ongoing constraints facing the individual (e.g., a complex home environment, etc.). Indeed, it is exactly analogous to TFP in a typical production function, which of course is a “catch-all” and not independently well-defined.

and who comprise 24 percent of the population), those who enroll in the presence of the subsidy and not in its absence (“switchers,” who account for 30 percent of the population), and those who choose not to enroll in either case (“nonenrollees”, who make up the remaining 46 percent). We find that the withdrawal of the subsidy matters most for those who always enroll. This finding has an important implication: for individuals with the poorest earnings prospects, changes in college subsidization are essentially irrelevant, while for those well-prepared to go, subsidies are meaningful. After all, for the latter, the removal increases the costs one-for-one of going to college, something they were fully planning on doing even in the absence of the subsidy. In other words, direct subsidies to college may most benefit those already best-suited to take advantage of college.

As noted above, for those individuals whose initial characteristics leave them ill-suited to reap benefits from college enrollment, current subsidies to attend college are not useful. This suggests that an alternative investment, whose returns are comparable to those of college but are not dependent on individual characteristics or subject to idiosyncratic risk, has the potential to be more valuable than college to a sizable fraction of US high school graduates. Perhaps surprisingly, the stock market meets both criteria.⁴ An indexed equity fund delivers ex-post returns with stochastic properties similar to those from college (Judd, 2000; Psacharopoulos and Patrinos, 2018). And unlike college (more generally, human capital), these returns are “blind” to the characteristics of their owners and, by construction, immune to idiosyncratic risk.⁵

To evaluate this alternative, we calculate the value of access to the stock market exactly as we did for college. We find that on average, the value of stocks (4.4 percent of consumption) is lower than the value of college (14.7 percent). However, it is precisely those who are ill-positioned to benefit from college who place the highest value on stocks. In particular, those in the lowest quartile of the ability distribution value access to stocks at 9.7 percent of consumption and college only at 1 percent of consumption. This, coupled with the finding that nearly a quarter of the population would continue to enroll in college *even in the absence of any subsidy*, while a full 46 percent receives no benefit from it, motivates us to consider whether welfare improvements are possible from an alternative regime. Specifically, we consider a case in which those whom we identify as “always enroll” no longer receive the direct subsidy to college. Instead, the funds are used to subsidize the purchase of a managed stock index fund, available only at retirement, for

⁴Another initially plausible alternative is housing; however, unlike those on a stock market index, the returns to housing do depend on owner behavior and idiosyncratic risk.

⁵Our restriction to an index fund is important because there is some evidence of heterogeneity in returns to individual stock portfolios (Fagereng et al., 2016; Bach et al., 2016). An index, by definition, offers all those who hold it the same returns. Even if we were to allow a correlation to exist between individual types and stock returns, it is likely that it would favor those with the highest levels of education. If those with lower levels of education are also less likely to “do well” in the stock market, this would only bolster the case for a stock-index fund—which gives them the market rate of return—making them better off.

the “nonenrollees”.⁶ The switchers, who clearly depend on the subsidy to enroll in college, are unaffected. Of course, college might be subsidized to deliver socially optimal human capital in the presence of positive externalities, or to generate a specific target level of college enrollment for other reasons. Notice that the alternative we propose ensures that any social benefits from college are preserved: it induces no reduction in enrollment (and hence in the skill premium) and is also budget-neutral.

The mean gain from moving to this regime from the baseline is 1.07 percent, thus making all better off in ex-ante terms (i.e., before knowing one’s initial type). Of course, as we will show, the gains are heterogeneous and in this case flow to those who derived little or nothing from college access.

The fact that there are gains to be had from moving to this regime suggests that current support for college, which is invariant to individual characteristics (i.e., the direct tuition subsidy currently in place) and meant to equalize opportunity, might instead be flowing to those already well-positioned to benefit from college. Our analysis is in no way an indictment of college subsidies as a whole nor is it an exhortation to subsidize stocks per se. Rather, our message—as embodied in these measurements—is that more widespread benefits may be available from a targeted approach that conditions investment subsidies on individual circumstances.

1.1 Related work

Our paper builds on work that is aimed at understanding the role of human capital when the particulars of college education, in terms of its costs and returns as a function of observable enrollee and household characteristics, are modeled explicitly. Important references in this literature include Arcidiacono (2005); Garriga and Keightley (2007); Johnson (2013); and Altonji et al. (2016). Recent work of Abbott et al. (2018) is clearly relevant as well. They develop a rich representation of higher education, allowing for a variety of salient features—gender, labor supply during college, government grants and loans (including private loans), and heterogeneity in familial resources—that have bearing on the measurement in which we are interested.

Our goal is to highlight the uneven flow of current college benefits across individuals and to suggest that a targeted alternative can more efficiently deploy the same public monies without interfering with college enrollment. By contrast, in the work cited above, a primary focus is

⁶This proposal is similar in spirit to policies that encourage asset building among households—such as child development accounts (CDAs)—which many different countries have promoted (see Loke and Sherraden, 2009, for an overview). Though the specifics differ, a common feature of such policies is that money is deposited in an individual’s account early in life and restrictions are placed both on when it can be withdrawn and what the funds can be used for. A full assessment of such policies is well beyond the scope of this paper.

on providing rich policy counterfactuals in which college enrollment is itself a policy target. We emphasize instead heterogeneity among individuals in order to answer the question of who benefits from access to college.⁷ We therefore adopt a partial equilibrium perspective and emphasize the derivation of the (not-directly-observable) joint distribution of learning ability, initial human capital stock, and initial financial wealth. These features, as argued above, are critical to accurately assessing individual-level variation in the valuation of college and of the alternative we study.⁸

We are also informed by the work that emphasizes the bias imparted to measured returns to college by the possibility of noncompletion. Hendricks and Leukhina (2014) allow for selection effects and argue that two layers of selection are important: weakly prepared students disproportionately fail to enroll in college, and those who do enroll fail to complete at high rates.⁹ Our model allows for both effects to operate and thereby avoids overstating the payoff to college. With respect to failure risk, our work builds on earlier work of Restuccia and Urrutia (2004), Akyol and Athreya (2005), and Chatterjee and Ionescu (2012).¹⁰ More recently, Athreya and Eberly (2013) demonstrate that college failure risk hinders low-wealth individuals, even relatively well-prepared ones, from enrolling in college.

Our work is also related to an empirical literature that has examined various aspects of heterogeneity in the returns to college, including differences in ability, college preparedness, and family background among individuals (see Card, 2001), and more recently also in field of study and school quality (e.g., Altonji et al., 2012). Our framework allows for heterogeneity in ability and college preparedness, though it does not explicitly consider heterogeneity in field of study or school quality. However, the two sets of characteristics appear to be positively correlated (Arcidiacono et al., 2012; Hendricks and Schoellman, 2014). In our model, heterogeneity in ability and college preparedness will act as a summary measure of all dimensions of heterogeneity in the returns to college investment.

With respect to stocks, our work follows the literature on portfolio choice in life-cycle models (see, for example, Cocco et al., 2005). In spirit, our work is also closely related to Kim et al. (2013), which also features both education and stock market investment.¹¹

⁷See also Epple et al. (2013) and Cestau et al. (2017) for analysis of higher education policies in the presence of substantial enrollee heterogeneity.

⁸Relatedly, our analysis abstracts from aggregate externalities from education. The empirical evidence on the magnitude of education externalities is mixed. While some authors find evidence of such externalities (e.g. Moretti, 2004a,b), others find that any such externalities are negligible (Acemoglu and Angrist, 2000; Ciccone and Peri, 2006). Nonetheless, our proposed alternative, because it preserves college enrollment, acknowledges that such externalities may be part of the underlying motivation for subsidizing college.

⁹See also, Arcidiacono (2004).

¹⁰The possibility of college failure has also been evaluated in work of Stange (2012) and Ozdagli and Trachter (2011).

¹¹Indeed, in Athreya et al. (2015), we incorporate the elements of Kim et al. (2013) in a model with human

The remainder of the paper is organized as follows. Section 2 describes the model and Section 3 the data we use to calibrate it. Section 4 summarizes the calibration, and the results are reported in Section 5. Section 6 concludes.

2 Model

Our aim is to first measure the value, across different types of individuals, of access to college, and then use these measures to assess the gains available from a more targeted alternative to the current direct subsidy policy. This requires a model rich enough to accommodate the salient types of heterogeneity across individuals, such as college readiness and individual-specific uninsurable shocks, as well as investment opportunities that allow for human capital accumulation via college and time, in addition to standard financial assets (a bond and a stock). Our model thus combines a standard life-cycle consumption-savings problem in which households face uninsurable idiosyncratic risk and have a meaningful portfolio choice, with a model of human capital accumulation that includes both the classical Ben-Porath technology and an investment opportunity that takes the particulars of college seriously. The details are as follows.

2.1 Environment

Time is discrete and indexed by $t = 1, \dots, T$ where $t = 1$ represents the first year after high school graduation. We allow for three potential sources of heterogeneity across agents at the beginning of life: their learning ability, a , which does not change after high school, their initial stock of human capital, h_1 , and their initial assets, x_1 . These characteristics are drawn jointly according to a distribution $F(a, h, x)$ on $A \times H \times X$.

Each period, agents choose how much to consume and how to divide their time between learning and earning, as in Ben-Porath (1967). Agents also decide how much of their wealth to allocate to stocks, s , versus bonds, b . The latter may be used to either borrow or save. Debt is not defaultable and is subject to a borrowing limit, $-\underline{b}$, where $\underline{b} > 0$.

Agents work and accumulate human capital using the Ben-Porath technology until $t = J$. Agents can also accumulate human capital by choosing (in the first period) to attend college. College can be financed using wealth, x , unsecured debt, b , and nondefaultable, unsecured student-loan debt, d . Agents retire in period $t = J + 1$, after which they face a simple consumption-savings and portfolio choice problem.

capital investment (though without four-year college) and show that it can match important life-cycle observations on household stock market participation.

To capture an empirically important dimension of human capital accumulation, we assume that agents may fail to complete college.¹² At the end of four years in college, the probability of completion—which depends on the agent’s learning ability as well as human capital accumulated to that point—is realized.¹³ Those who complete college start their working life with human capital h^{CG} , where CG denotes college graduates. Those who fail to complete start their working life with human capital h^{SC} , where SC denotes “some college,” and those who choose not to go to college start their working life at $t = 1$ with human capital h^{HS} , where HS denotes “high school.”¹⁴

2.2 Preferences

Agents maximize the expected present value of utility over the life cycle:

$$\max E_0 \sum_{t=1}^T \beta^{t-1} u(c_t), \tag{1}$$

where $u(\cdot)$ is strictly concave and increasing. Preferences are represented by a standard time-separable constant relative risk aversion (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 throughout their working lives.

Both within and outside college, agents accumulate human capital using a Ben-Porath technology. However, the incentives to invest in human capital are different in and outside college, for several reasons. First, the rental rate on human capital grows faster for those who complete college, consistent with empirical evidence that shows faster earnings growth for college graduates. Second, college enrollees have access to grant funding, which is not available outside of college, as well as to student loan credit that carries a lower rate of interest than the unsecured credit available to all agents. Access to grants and cheaper credit makes funding consumption while spending

¹²For example, Bound et al. (2010) report, using NLS72 data, that only slightly over half of all college enrollees graduate within eight years of enrollment.

¹³For computational simplicity, we assume that the probability of failure is realized only once, at the end of four years in college. We think this is a reasonable assumption considering that 89 percent of college dropouts are enrolled in college for at least three full years, according to data from the Beginning Postsecondary Students Longitudinal Study (BPS).

¹⁴Note that there is variation in the value of h^{CG} , h^{SC} , and h^{HS} across individuals.

time accumulating human capital relatively easier on the college path than on the no-college path. Third, college graduates enjoy a higher replacement rate on their income at retirement (see Section 4). Finally, the opportunity cost of spending time learning is higher on the no-college path than on the college path. Outside college, we assume that earnings are a function of accumulated human capital, whereas in college they are not: those who work while in college face a relatively low wage rate that does not differ with the level of human capital. This assumption is consistent with evidence that the jobs college students hold do not necessarily value students' human capital stocks. In fact, we assume that working takes time away from human capital accumulation, and that accumulating less human capital decreases the odds of completion. This, too, is consistent with empirical evidence that college jobs do not contribute to human capital accumulation and that students who work while in college are more likely to drop out (see Autor et al., 2003; Peri and Sparber, 2009). Consequently, most college students in the model find it optimal to allocate all of their time to human capital accumulation, which is in line with empirical findings that the majority of full-time students do not work while in college (see Manski and Wise, 1983; Planty et al., 2008). College in the model thus represents a device that can greatly accelerate human capital accumulation but harshly penalizes noncorner solutions for time allocation. Taken together, these factors make human capital accumulation more attractive in college than outside of it during the college years.

2.3.1 Ben-Porath Human Capital Investment

During college and while working, agents accumulate human capital (as in the classic Ben-Porath, 1967, model):¹⁵

$$h_{t+1} = h_t(1 - \delta) + a(h_t l_t)^\alpha \text{ with } \alpha \in (0, 1) \quad (2)$$

Human capital production depends on the agent's learning ability, a , human capital, h_t , the fraction of available time put into it, l_t , and the production function elasticity, α . Human capital depreciates at a rate of δ , which we will allow to differ by education groups.

2.3.2 College and its Financing

From the outset, we have highlighted the fact that ex-ante returns to investment in college vary significantly because of their dependence on individual characteristics. These characteristics drive the likelihood of completion, subsequent earnings, and—via the institutional structure of higher

¹⁵Blandin (2018) shows that the Ben-Porath model does much better at matching various empirical properties of life-cycle earnings than alternative models such as learning-by-doing.

education financing—the cost of college. We now detail the accommodation of these features in our model.

Those who invest in college face the possibility of noncompletion, which decreases with the level of human capital accumulated during college. Specifically, the probability of completion, $\pi(h_5(h_1, a, l_{1,\dots,4}^*))$, is an increasing function of the amount of human capital accumulated after completing four years in college, h_5 , which in turn increases with the initial human capital stock, h_1 , the agent’s learning ability, a , and the amount of time, $l_{1,\dots,4}^*$, that she chooses to allocate to human capital accumulation (versus working) while in college.

Those who work in college can earn a maximum of w_{col} if they allocate all of their time to working. Working during college diverts time away from human capital accumulation and therefore increases the probability of noncompletion.

There are several possible sources of college financing: savings, x , borrowing, b , earnings from working while in college, $w_{col}(1-l)$, merit- and need-based aid, $\kappa(a, x_1)$, and student loans. Notice that because financial aid varies with individual characteristics, so too will the net cost of college. Agents are allowed to take out student loans up to $\min[d_{max}, \max(\bar{d} - x_1, 0)]$. Here, \bar{d} represents a measure of the “full” cost of attending college; it includes tuition, books and supplies, room and board, etc. Students can borrow up to this amount less x_1 as long as the student loan limit, d_{max} , is not exceeded. We think of x_1 not literally as the amount the student has saved (since savings are negligible for most at this point in the life cycle), but rather as the so-called “expected family contribution (EFC),” i.e., the savings to which the student has access. The EFC is a quantity that federal student loan programs calculate as a part of determining the amount of subsidized lending to which an enrollee is entitled.

Students choose how much to take out in student loans, d , in total (i.e., the stock of debt with which they will emerge from college) at the beginning of college and receive equal fractions of that loan ($\frac{d}{4}$) in each period (year) of college. While attending college, students are billed annually, in the amount \hat{d} , which represents the direct cost of college (e.g, tuition and fee payments levied by the college). After college, students repay their loan in equal payments, p , which are determined by the loan size, d , 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.

The return to human capital is in the form of earnings during working life, which are subject to shocks, as described below.

2.4 Earnings

During an agent's working life, their earnings are given by:

$$y_{it} = w_t(1 - l_{it})h_{it}z_{it}$$

where w is the rental rate of human capital, $(1 - l_{it})$ is the time spent working, and z_{it} is the stochastic component. Both the rental rate on human capital and the stochastic component vary between college graduates, CG , and those with no college, NC (which includes college dropouts and high school graduates). The latter consists of a persistent component $u_{it} = \rho u_{i,t-1} + \nu_{it}$, with $\nu_{it} \sim N(0, \sigma_\nu^2)$, and a transitory independent and identically distributed (iid) component $\epsilon_{it} \sim N(0, \sigma_\epsilon^2)$. The variables u_{it} and ϵ_{it} are realized in each period over the life cycle and are not correlated.

The rental rate of human capital evolves over time according to

$$w_t = (1 + g)^{t-1}$$

where g is the growth rate. This rate is higher for college graduates than for those with no college.

2.5 Means-Tested Transfer and Retirement Income

We allow agents to receive means-tested transfers, τ_t , which depend on age, income, and assets. Following Hubbard et al. (1994), we specify these transfers as

$$\tau_t(t, y_t, x_t) = \max\{0, \underline{\tau} - (\max(0, x_t) + y_t)\} \quad (3)$$

These transfers capture the net effect of the various US social insurance programs that are aimed at providing a floor on income (and thereby on consumption).

After period $t = J$, in which agents start retirement, they receive a constant fraction of their earnings in the last working period, $\varphi(y_J + \tau_J)$, which they allocate between risky and risk-free investments. We allow the income replacement rate for college graduates to differ from the rate for all other agents.

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 . Savings will earn the risk-free interest rate, R_f . We assume that the borrowing rate, R_b , is higher than the savings rate: $R_b = R_f + \omega$. Debt is nondefaultable and comes with a borrowing limit $\underline{b} > 0$.

Risky assets

Risky assets, or stocks, earn stochastic return $R_{s,t+1}$ in period $t + 1$, given by:

$$R_{s,t+1} - R_f = \mu + \eta_{t+1}, \quad (4)$$

where η_{t+1} , the period $t + 1$ innovation to excess returns, is assumed to be iid 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

$$x_{t+1} = R_j b_{t+1} + R_{s,t+1} s_{t+1}$$

with $R_j = R_f$ if $b \geq 0$ and $R_j = R_b$ if $b < 0$.

2.7 Agent's Problem

The agent chooses whether or not to invest in college (and, if investing in college, how much student debt to take on), how much to consume, how much time to allocate to learning, and asset positions in stocks and bonds (or borrowing) to maximize expected lifetime utility.

We solve the problem backward, starting with the last period of life when agents consume all their available resources. The value function in the last period of life is set to $V_T^R(a, h, x) = u(x)$.

Retired agents do not accumulate human capital. They face a simple consumption-savings problem and may choose to invest in both risk-free and risky assets. The value function is given by

$$V^R(t, a, b, s, y_J + \tau_J) = \max_{b', s'} \left\{ \frac{c^{1-\sigma}}{1-\sigma} + \beta V^R(t+1, a, b', s', y_J + \tau_J) \right\} \quad (5)$$

where

$$\begin{aligned} c + b' + s' &\leq \varphi(y_J + \tau_J) + R_j b + R_s s \\ b' &\geq \underline{b} \\ s' &\geq 0 \end{aligned}$$

In the above, $R_j = R_f$ if $b \geq 0$, and $R_j = R_b$ if $b < 0$. The only uncertainty faced by retired individuals pertains to the rate of return on the risky asset.

2.7.1 Problem in Working Phase for those with No College

We use $V_J^R(t, a, b, s, y_J + \tau_J)$ from Equation 5 as a terminal node for the adult's problem on the no-college path. We solve

$$V^{HS}(t, a, h, b, s, z) = \max_{l, b', s'} \left\{ \frac{c_t^{1-\sigma}}{1-\sigma} + \beta EV^{HS}(t+1, a, h', b', s', z') \right\} \quad (6)$$

where

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

2.7.2 Problem in Working Phase for those who Attended College

As before, we use $V_J^R(t, a, b, s, y_J + \tau_J)$ from the retirement phase as a terminal node and solve for the set of choices in the working phase $j = 5, \dots, J$ of the life cycle. We further break down the working phase into a student loan post-repayment period and a repayment period. In the post-repayment period, $t = P + 1, \dots, J$, the problem is identical to the one for working adults on the no-college path.

During the repayment period, $t = 5, \dots, P$, agents have to repay their student loans with a per-period payment

$$p = \frac{d}{\sum_{t=1}^{P-5} \frac{1}{R_g^t}}.$$

The value function is given by

$$V^j(t, a, h, b, s, z) = \max_{l, b', s'} \left\{ \frac{c_t^{1-\sigma}}{1-\sigma} + \beta EV^j(t+1, a, h', b', s', z') \right\}, j = CG, SC \quad (7)$$

where

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

$R_j = R_f$ if $b \geq 0$ and $R_j = R_b$ if $b < 0$.

2.7.3 Problem in College

For the college phase $t = 1, \dots, 4$ of the life cycle, we first take into account the likelihood 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. The value function is given by

$$V^C(t, a, h, b, s, z) = \max_{l, b', s'} \left\{ \frac{c^{1-\sigma}}{1-\sigma} + \beta EV^C(t+1, a, h', b', s', z') \right\} \quad (8)$$

where

$$\begin{aligned} c + b' + s' &= w_{col}(1-l) + R_b b + R_s s + \frac{d}{4} - \hat{d} + \kappa(a, x_1) \\ l &\in [0, 1] \\ h' &= h(1-\delta) + a(hl)^\alpha \\ d &\leq \min[d_{max}, \max[\bar{d} - x_1, 0]] \\ b' &\geq \underline{b} \\ s' &\geq 0. \end{aligned}$$

While in college, the rental rate of human capital is set to a relatively low value (see Section 4), which means that human capital is not productive until graduation. As noted earlier, this reflects evidence that the jobs in which college students work do not necessarily utilize or augment students'

human capital stocks. The set of skills involved in these jobs is different from the one students acquire in college and use after graduation. An implication of this assumption is that in the model college students find it optimal to allocate all of their time in college to human capital accumulation, a result that is consistent with the empirical findings that the majority of full-time college students do not work while in school. Finally, people who choose to work while in school most likely drop out of college, as numerous studies attest.

Our model captures the large variety of resources that are available to students to finance their college education in addition to those obtained from working during college. Every year in college, students use a combination of personal and family savings, captured by x , unsecured borrowing, b , student loans, $\frac{d}{4}$, and merit- and need-based grants, $\kappa(a, x_1)$, to pay for direct college expenses, \hat{d} . We assume that the direct cost paid each period while in college incorporates the existing large subsidy that students receive to finance college, as we discuss in the calibration section. As described before, our model captures the key features of the student loan program with agents being allowed to borrow up to the full college cost minus the expected family contribution, $\bar{d} - x_1$, as long as they do not hit the constraint d_{max} . Importantly, agents use the loan amount to pay for college expenses while in college.¹⁶

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), on financial assets from the Survey of Consumer Finances (SCF), and on college enrollment and completion rates from the Beginning Postsecondary Student Longitudinal Survey (BPS) 2004/2009 and the National Education Longitudinal Study (NELS:1988).

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 life cycle. To this end, we first estimate life-cycle 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

¹⁶Without this feature, college enrollees would have an advantage over the no-college group in being able to borrow at a subsidized rate to finance consumption over the life-cycle.

Figure 1: Life-cycle earnings statistics



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 US 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 life-cycle 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.¹⁷ We then take the regression coefficients on the age indicators (we use the latest year 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 life-cycle profiles that we use in the calibration. Figure 1 shows the life-cycle 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 composed 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, Individual Retirement Accounts (IRAs)/Keoghs, thrift-type retirement accounts, and other managed assets.

As in the case of earnings, we construct life-cycle profiles of asset holdings, controlling for time effects. The results (in 2014 dollars) are reported in Figure 2.¹⁸

3.3 College enrollment and completion


We use data from the Beginning Postsecondary Student Longitudinal Survey (BPS) 2004/2009 and the National Education Longitudinal Study (NELS:1988) to match enrollment and completion rates. Specifically, we estimate correlations of ability and initial wealth, and of initial human capital and initial wealth, to match college enrollment rates for three groups of initial wealth (expected family contributions) based on NELS:1988 data and to match college completion rates based on the BPS 2004/2009 dataset for students who enrolled in college in the year 2003-04.

The BPS 2004/2009 is one of several National Center for Education Statistics (NCES)-sponsored studies that is a nationally representative dataset with a focus on postsecondary education indicators. BPS cohorts include beginners in postsecondary schools who are surveyed at three points in time: in their first year in the National Postsecondary Student Aid Study (NPSAS) and then three


¹⁷By 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 of the life-cycle profiles is qualitatively similar.

¹⁸Averages for risky and risk-free assets are taken conditional on ownership.


Figure 2: Average Life-Cycle Assets (SCF)



(a) Total



(b) Risky



(c) Risk-free

and six years after first starting their postsecondary education in follow-up surveys. BPS collects data on a variety of topics, including student demographics, school experiences, persistence, borrowing/repayment of student loans, and degree attainment six years after enrollment. Our sample consists of students aged 20-30 who enroll in a four-year college following high school graduation. For demographic characteristics, we use SAT (and converted ACT) scores as a proxy for ability and expected family contribution (EFC) as a proxy for wealth.

The NELS:1988 is a nationally representative sample of eighth graders who were first surveyed in the spring of 1988. A sample of these respondents were then resurveyed through four follow-up surveys in 1990, 1992, 1994, and 2000. We use the third follow-up survey when most respondents completed high school and report their postsecondary access and choice. As in the BPS, demographic information, including SAT scores and EFC, is available. We use this dataset to compute college enrollment rates by EFC. Our sample consists of recent high school graduates aged 20-30

who have taken the SAT (or ACT).

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 financial 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. These parameters are listed in Table 1. We present the details of the calibration in the next section, followed by the model fit relative to data.

Table 1: Parameter Values: Benchmark Model

Parameter	Name	Value
T	Model periods (years)	58
J	Working periods (after college)	34
β	Discount factor	0.96
σ	Coeff. of risk aversion	3
R_f	Risk-free rate	1.02
R_b	Borrowing rate	1.11
\underline{b}	Borrowing limit	\$17,000
μ	Mean equity premium	0.06
σ_η	Stdev. of innovations to stock returns	0.157
α	Human capital production function elasticity	0.7
g_{NC}, g_{CG}	Growth rate of rental rate of human capital	0.01, 0.02
δ_{NC}, δ_{CG}	Human capital depreciation rate	0.021, 0.038
ψ_{NC}, ψ_{CG}	Fraction of income in retirement	0.682, 0.93
$\underline{\tau}$	Minimal income level	\$17,936
$(\rho_{NC}, \sigma_{\nu_{NC}}^2, \sigma_{\epsilon_{NC}}^2)$	Earnings shocks no college	(0.951, 0.055, 0.017)
$(\rho_{CG}, \sigma_{\nu_{CG}}^2, \sigma_{\epsilon_{CG}}^2)$	Earnings shocks college	(0.945, 0.052, 0.02)
$(\mu_a, \sigma_a, \mu_h, \sigma_h, \varrho_{ah}, \varrho_{ax}, \varrho_{hx})$	Parameters for joint distribution of ability initial human capital, initial wealth	(0.29, 0.50, 79.71, 45.26, 0.66, 0.36, 0.42)
\hat{d}	Annual direct cost of college	\$7,100
\bar{d}	Full cost of college	\$53,454
d_{max}	Limit on student loans	\$23,000
R_g	Student loan rate	1.068
w_{col}	Wage during college	\$17,700

4.1 Calibration

4.1.1 Preference parameters

The per period utility function is CRRA as described in the model section. We set the coefficient of risk aversion, σ , to 3, which is consistent with values found in the literature.¹⁹ The discount factor used ($\beta = 0.96$) 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.

4.1.2 Human capital parameters and earnings shocks

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. To parameterize the stochastic component of earnings, z_{it} , we follow Abbott et al. (2018), who use the National Longitudinal Survey of Youth (NLSY) data using CPS-type wage measures to estimate parameters for the idiosyncratic persistent and transitory wage shocks. For the persistent shock, $u_{it} = \rho u_{i,t-1} + \nu_{it}$, with $\nu_{it} \sim N(0, \sigma_\nu^2)$ and the transitory shock, $\epsilon_{it} \sim N(0, \sigma_\epsilon^2)$, they report the following values: for high school graduates, $\rho = 0.951$, $\sigma_\omega^2 = 0.055$, and $\sigma_\nu^2 = 0.017$; for college graduates, $\rho = 0.945$, $\sigma_\omega^2 = 0.052$, and $\sigma_\nu^2 = 0.02$. We use the first set of values for individuals with no college as well as for those with some college education and the second set of values for those who complete four years of college.

As previously noted, the rental rate of human capital in the model evolves according to $w_t = (1 + g)^{t-1}$. The growth rate g is calibrated to match the average growth rate in mean earnings observed in the data. We obtain 0.01 for individuals with no college degree and 0.02 for college graduates.

Given the growth rate in the rental rates, the depreciation rates are set so that the model produces the rate of decrease of average real earnings at the end of the working life. The model implies that at the end of the life cycle negligible time is allocated to producing new human capital and, thus, the gross earnings growth rate approximately equals $(1 + g)(1 - \delta)$. We obtain 0.021 for individuals with no college degree and 0.038 for college graduates.

¹⁹We conduct robustness checks on this parameter by looking at alternative values such as the upper bound of $\sigma = 10$ considered reasonable by Mehra and Prescott (1985) as well as lower values such as $\sigma = 2$. The results are available upon request.

4.1.3 Distribution of initial characteristics: financial assets, ability and human capital

The distribution of initial characteristics (ability, human capital, and financial 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) and college enrollment and college completion rates across three wealth groups (proxied by expected family contributions). The estimation proceeds as follows. First, for the distribution of initial financial assets, x_1 , we use data from the BPS 2004/2009.²⁰ Second, we calibrate the joint distribution of ability and initial human capital to match the key properties of the earnings distribution over the life cycle reported earlier using March CPS data. Third, we estimate the correlations of ability and initial wealth, and of initial human capital and initial wealth, to match college enrollment rates based on NELS:1988 data, and college completion rates based on BPS 2004/2009 data.

The dynamics of the earnings distribution implied by the model are determined in several steps: i) we compute the optimal decision rules in the model using the parameters described above for an initial grid of the state variable; ii) we simultaneously compute college, human capital, and financial investment decisions and compute the life-cycle 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 earnings from 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.²¹ We find the vector of parameters $\gamma = (\mu_a, \sigma_a, \mu_h, \sigma_h, \rho_{ah})$ 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

²⁰The proper notion of wealth at age 18 is not unambiguous. In particular, while 18-year-olds typically do not have substantial wealth of their own, they may have access to alternative sources of wealth that are not directly measured, most notably, inter vivos transfers from their parents. In this context, therefore, the EFC seemed to us to be the most appropriate measure of wealth available to high school graduates.

²¹In practice, the grid is defined by 20 points in human capital and in ability.

data on earnings, and $m_j(\gamma)$, $d_j(\gamma)$, and $s_j(\gamma)$ are the corresponding model statistics.²²

We choose the correlations of ability and initial wealth, and of initial human capital and initial wealth, that best replicate college enrollment and college completion rates by wealth levels (see further details in the next subsection). Our estimation delivers a correlation between ability and initial human capital stock of 0.67 and a correlation between initial wealth and ability and initial human capital of 0.36 and 0.42, respectively.

4.1.4 College parameters

We first specify parameters related to students' ability to borrow public funds to finance their college education. We set the full cost of college to $\bar{d} = \$53,454$. This figure is the enrollment-weighted average of the full annual cost at a public four-year institution (\$33,849) and a private four-year institution (\$78,570) between academic years 2003-04 and 2007-08, weighted by the fraction of students attending each in the data (60 and 40 percent, respectively). Recall from the agent's problem that this is not what students actually pay but rather a parameter that influences the borrowing limit on student loans, which is set at $\min[d_{max}, \max[\bar{d} - x_1, 0]]$. We set the other parameters that governs this constraint, i.e., the direct limit on student loans, to $d_{max} = \$23,000$, and the interest rate to $R_g = 1.068$, respectively, as specified in the Department of Education guidelines for the students who enter college in 2003-04.

We turn next to parameters that determine the actual privately borne cost of college (both direct and in terms of forgone earnings). These costs are of course the relevant ones for anyone deciding whether to enroll. First, the annual direct cost of college, \hat{d} , (which, as described earlier, incorporates direct subsidies and is what students are actually billed), is \$7,100. This is a proportion (53 percent) of an enrollment-weighted average of the direct cost at public and private four-year institutions.²³ The key point is that the subsidy reduces the direct cost of college by 47 percent. Next, we set the wage during college, w_{col} , to \$17,700 based on NCES data. Lastly, we must spell out nonloan financial aid available to enrollees. Here, we use the BPS data to set both merit- and need-based aid. Merit-based aid increases with an enrollee's high school GPA: for example, those in the top GPA quartile receive merit-based aid equivalent to about 63 percent of the total college cost, while those in the bottom quartile receive about 10 percent. In the model, we use ability as a proxy for high school GPA. To be consistent with the data, we assume that the fraction of the total college cost covered by merit-based aid increases with ability. We calibrate it to ensure that the average aid amount within each ability quartile in the model matches the

²²For details on the calibration algorithm, see Huggett et al. (2006) and Ionescu (2009).

²³For details on how these costs are calculated, see Ionescu and Simpson (2016).

average aid amount within each GPA quartile in the data.

Turning to need-based aid, we observe in the BPS data that students in the bottom quartile of EFC receive about 7 percent of the college cost in the form of need-based aid, on average, while those in the top quartile of EFC receive, on average, no need-based aid. Since EFC is a function of wealth, we assume that the fraction of the total college cost received in the form of need-based aid decreases with wealth. We calibrate this so that the average need-based aid within each quartile of initial wealth in the model matches average need-based aid within the corresponding quartile of EFC in the data.

We have specified the parameters governing access to public student loans as well as the costs that any individual seeking to attend college faces. We now specify the risks associated with college investment. In the model, we assume that the probability of college completion is a function of the amount of human capital accumulated at the end of college, h_5 (which in turn is a function of ability, initial human capital, and the time spent learning during college). We use cumulative GPA scores in the BPS data as a proxy for h_5 . In the data, we observe the fraction of the student population that obtained each of the grades listed in Table 2. In the model, we divide the distribution of h_5 into groups according to these percentages, and we assign each group the completion probability listed in the first column of the table.²⁴ For example, an agent in the group with the highest level of h_5 will face a 70 percent probability of completion.

Table 2: Completion Rates by GPA in College

Completion rate	Grades
0.07	mostly Cs and Ds
0.30	mostly Cs
0.45	mostly Bs and Cs
0.56	mostly Bs
0.67	mostly Bs and As
0.70	mostly As

4.1.5 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 in the literature (e.g., Mehra and Prescott, 1985). The

²⁴We define the completion rate in the data as the fraction of students who had earned a bachelor's degree by June 2009.

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 0.09 to match the average borrowing rate of $R_b = 1.11$ (Board of Governors of the Federal Reserve System, 2014).

We assume a uniform credit limit (\underline{b}) across households. We obtain the value for this limit from the SCF. The SCF reports, for all individuals who hold one or more credit card, the sum total of their credit limits. We take the average of this over all individuals in our sample and obtain a value of approximately \$17,000 in 2014 dollars. Note that, when we take the average, we include those who do not have any credit cards. This ensures that we are not setting the overall limit to be too loose. Lastly, in our baseline model, we assume for the time being that the returns to both risky assets (human capital and financial wealth) are uncorrelated.

4.2 Model vs. Data

We start by presenting the model predictions for targeted data moments for the baseline economy, and we then describe model predictions for key nontargeted data moments.

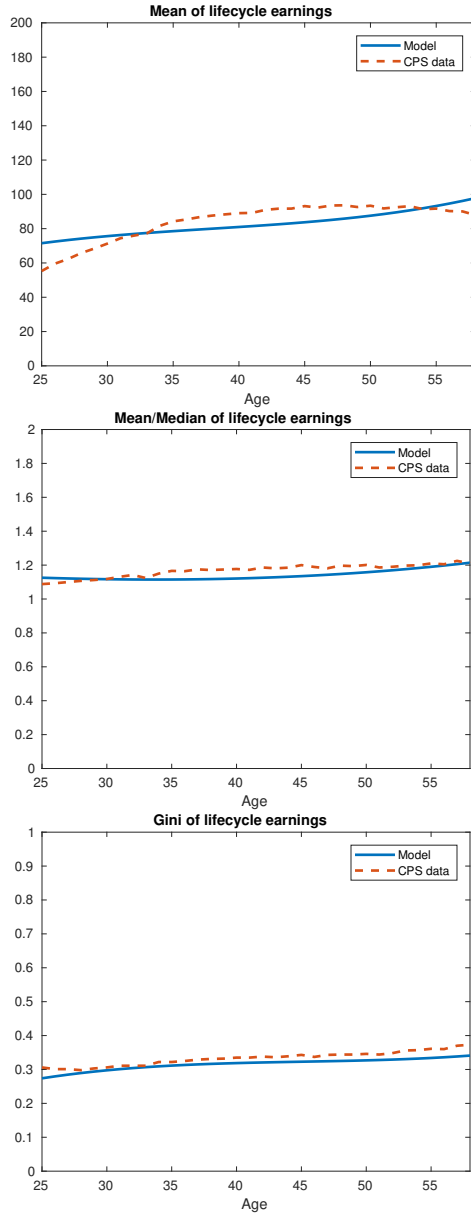
4.2.1 Targeted Moments

This section presents measures of goodness of fit for the baseline model. Figure 3 shows the earnings moments for a simulated sample of individuals in the model versus the CPS data.²⁵ As the figure shows, the model does a reasonably good job of fitting the evolution of mean earnings over the life cycle, though the model’s profile is less hump-shaped than 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 first and last few years of the life cycle.

We next look at the model’s predictions for college investment behavior by initial wealth. Table 3 shows college enrollment and completion rates by level of initial financial wealth; where “low” refers to the bottom quartile, “medium” to the two middle quartiles, and “high” to the top quartile of the distribution of initial wealth. As can be seen, the baseline calibration captures well the fact that both enrollment and completion rates are strongly increasing in the level of initial wealth.

²⁵As a measure of goodness of fit, we use $\frac{1}{3(J-4)} \sum_{j=5}^J |\log(m_j/m_j(\gamma))| + |\log(d_j/d_j(\gamma))| + |\log(s_j/s_j(\gamma))|$. This represents the average (percentage) deviation, in absolute terms, between the model-implied statistics and the data. We obtain a fit of 8 percent (where 0 percent represents a perfect fit).

Figure 3: Life-cycle earnings statistics



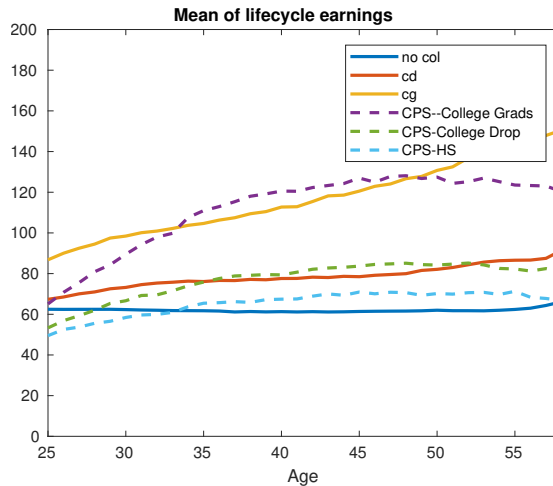
4.2.2 Nontargeted Moments

We show next that our model performs well along relevant nontargeted dimensions. Given our focus on the payoff to investment opportunities, key among these is life-cycle earnings across groups defined by educational attainment. Note that our calibration only targeted overall earnings and, not earnings by education group. Figure 4 shows that our model nonetheless delivers the pattern seen in the data of mean earnings by education group over the life cycle.

Table 3: Targeted Moments Data vs Model: Enrollment and Completion

Initial wealth	Benchmark	Data
<i>College Enrollment</i>	54	47
Low	35	34
Medium	55	47
High	74	62
<i>College Completion</i>	49	45
Low	43	37
Medium	49	45
High	57	60

Figure 4: Life-Cycle Earnings by Education Group



We next examine college enrollment and completion behavior by individual characteristics. As seen in Table 4, the model predicts that both college enrollment and completion rates are increasing in ability and in initial human capital. While there is no direct data counterpart to the notions of ability and initial human capital as represented in the Ben-Porath setting, we see that when college investment behavior is ordered by SAT score—arguably the most widely used measure of college readiness—the model’s implications are qualitatively borne out in the data.

We now look at the model’s predictions for financial wealth. Figure 5 shows the mean wealth accumulation over the life cycle for total assets as well as for risky and risk-free assets. Overall, the model is consistent with the overall trajectory of wealth accumulation, but it underpredicts mean wealth by age. We note that mean wealth in the US data is strongly influenced by the extreme right tail of the distribution. Indeed, this has led models aimed at capturing the skewness of wealth to employ earnings processes in which agents receive extremely large but transitory shocks to earnings with extremely low probability (Castaneda et al., 2003). As a result, the presence or

Figure 5: Life-Cycle Wealth Accumulation

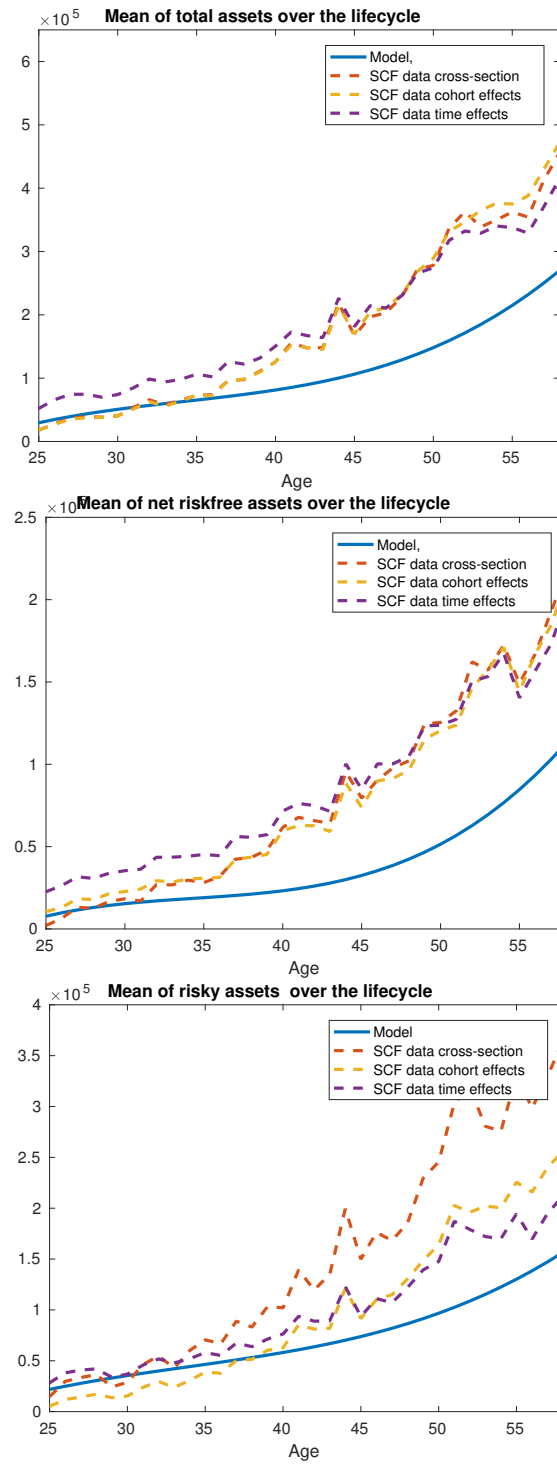
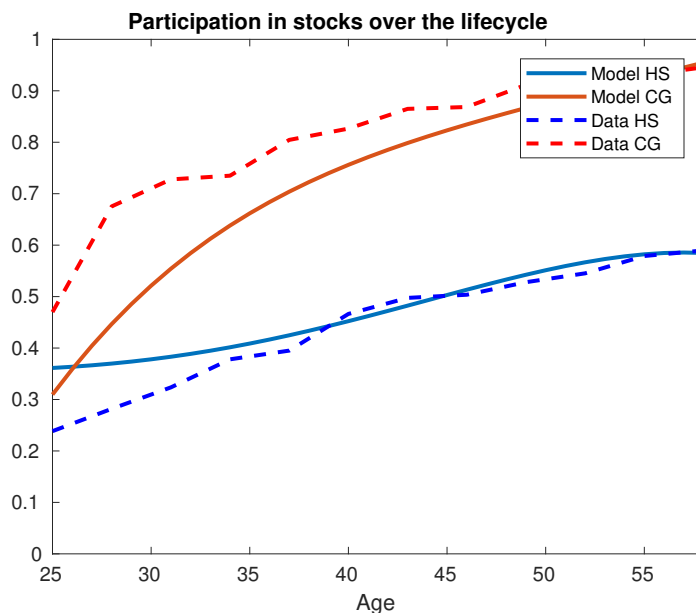


Table 4: Nontargeted Moments: Enrollment and Completion by Characteristics

Characteristic	Ability	Initial Human Capital	Data: SAT scores
<i>College Enrollment</i>			
Low	9	26	53
Medium	63	65	65
High	85	64	85
<i>College Completion</i>			
Low	20	27	30
Medium	42	48	50
High	64	68	69

absence of such improbable shocks is unlikely to be quantitatively important for wealth at the individual level.

Figure 6: Stock Market Participation over the Life Cycle by Education Groups



Finally, as shown in Figure 6, our model’s prediction for the stock-market participation rate is consistent with the data, over the entire life cycle and by education groups. This result is driven primarily by the presence of human capital. Human capital is an attractive investment early in life, especially for those with a combination of high learning ability and relatively low initial human capital: the opportunity cost of spending time learning—forgone earnings—is relatively low, the marginal return to learning is high, and the horizon over which to recoup any payoff from learning is long. Further, anticipating rising earnings over the life cycle, households who invest in

human capital early in life will desire, absent risk, to avoid large positive net positions in financial assets when young. As they age and accumulate human capital, these households will find further investment in human capital less attractive as the marginal return decreases and opportunity cost increases. These high earners will then accumulate wealth and participate in the stock market at high rates. This mechanism, illustrated in detail in Athreya et al. (2015), delivers a profile of aggregate stock market participation that is consistent with the data, as Figure 6 shows.

5 Results

We now use the model described and parameterized above to address the questions posed at the outset: What is the value of access to college across the population? How do these valuations depend on the direct college subsidies currently in place? Are there simple alternatives that illustrate the potential for gains in welfare? We will demonstrate that these valuations vary widely, and that they depend strongly on existing subsidies. These findings then motivate us to study the implications of reallocating subsidies away from inframarginal college attendees and towards inframarginal nonenrollees.

5.1 The Value of Access to College

We assess the value of access to college across individuals by comparing our baseline economy (where individuals have access to both college and stocks) to an economy with no college. The value of access to college for each agent type is calculated as the consumption-equivalent gain obtained from moving from the no-college environment to the baseline.

Whenever we reduce the set of options available to agents (compared to the baseline economy), we allow individuals to fully reoptimize given the new constraints they face. We do this because we do not want to overstate the value of any given investment opportunity by forcing the agent away from fully exploiting whatever substitution possibilities they may have available. In this instance, absent college, agents can make different decisions about stock market investment (as well as about how much time to spend on human capital accumulation) than in the baseline.

We measure the value of access to college via the following steps. First, for each agent type—where type is defined by a profile of *initial* characteristics $s_1 = (a, x_1, h_1)$ —we compute the ex-ante expected utility associated with starting life in each of the two economies: the baseline (B) and the no-college economy (NC). The expectation arises from uncertainty in income and in the rate of return on stocks to which all agents are exposed. Thus, the relevant object is $V^k(s_1)$ where s_1 represents the agent type’s pre-college-decision state vector (i.e., the initial state) and k denotes

the economy in question: $k \in \{B, NC\}$. The value function $V^k(s_1)$ is, by definition, the maximized expected utility for an agent with initial characteristics s_1 .

Second, given measures of V^k for all types s_1 (and for all k), the value of access is calculated as follows. Let the constant-consumption equivalent to any value V^k be defined as the constant value of consumption $\bar{c}^k(s_1)$ that satisfies $\sum_{t=1}^T \beta^{t-1} \frac{\bar{c}^{1-\sigma}}{1-\sigma} = V^k(s_1)$, or:

$$\bar{c}^k(s_1) = \left[\frac{1-\beta}{1-\beta^T} (1-\sigma) V^k(s_1) \right]^{\frac{1}{1-\sigma}}$$

Given \bar{c}^k for all k , we can compute the value of access to the baseline economy for an agent type s_1 currently in economy k —in terms of the net percentage difference in the relevant constant-consumption equivalents—as:

$$\gamma^k = \frac{\bar{c}^B(s_1)}{\bar{c}^k(s_1)} - 1$$

Note that because we hold prices (interest rates and stock returns) and income processes (conditional on educational attainment) fixed, absent any compensation available to agents in the no-college economy (or, in the next section, in the no-stocks economy), agents cannot be made worse off by moving from the economy with more restricted investment opportunities to the baseline. Accordingly, the value of access is nonnegative for all agent types.

We find that the mean value (across all agent types) of access to college is 14.7 percent of consumption. This is a substantial amount, corresponding to about \$6,000 per year (under the assumption that mean annual consumption is about \$40,000 annually). The relatively large valuation reflects that college provides an efficient way to accumulate human capital and that college completion yields a mean earnings premium that is quite large.

Our framework allows us to further ask who in the population derives the most value from access to college. There is, of course, an entire distribution of agent types, and therefore an entire distribution of valuations we could report. To keep results tractable while informative of the basic trade-offs present, we begin by reporting valuations averaged over all individuals within, respectively, the bottom quartile (“low”), the two middle quartiles (“middle”), and the top quartile (“high”) of learning ability, initial human capital, and initial wealth.

Table 5 shows that mean valuations vary greatly across ability, initial human capital, and initial wealth groups. Starting with ability, we see that for those in the lowest quartile of the ability distribution, access to college is worth very little—about 1 percent in consumption-equivalent terms. The value rises rapidly with ability, reaching about 28 percent of consumption, on average, for those in the top quartile of ability. For the latter quartile, college completion is nearly guaranteed

Table 5: Average Valuation of College Access by Initial Characteristic

Initial Characteristic	Valuation (% of Consumption)
<i>Ability</i>	
Low	1.0%
Middle	15.1%
High	27.5%
<i>Human Capital</i>	
Low	13.3%
Middle	13.2%
High	19.0%
<i>Wealth</i>	
Low	9.3%
Middle	13.7%
High	22.0%

and, given the large college earnings premium in the economy, the opportunity to invest in college is worth a lot.

Looking next across groups defined by initial wealth (the bottom panel of the table), we see that a similar story emerges, though the gains do not rise quite as sharply as in the case of ability. This likely reflects the absence of significant credit constraints in these economies (which in turn reflects our reading of the literature on the empirical strength of borrowing constraints for education). The presence of risk of course does make a difference, whereby greater initial wealth allows a greater number of those with low ability and low initial human capital (i.e., those facing a higher risk of noncompletion) to invest in college anyway. As for the role of initial human capital (the middle panel of the table), the results are less clearly ordered, reflecting a trade-off: those with low initial human capital face both a higher marginal return to investing in college and a higher risk of noncompletion.

5.2 Financial Assets and Stocks as Alternatives to College


The heterogeneity in the value of access to college shown above arises from heterogeneity in initial characteristics that affect the likelihood of college completion as well as from the idiosyncratic risk that affects returns conditional on completion. And here, our assessment that a large group would attend college even absent the subsidy is relevant. It suggests that shifting at least some public funding towards assets whose returns do not depend on either individual characteristics or idiosyncratic risk might be valuable to many, or even potentially *all*, from an ex-ante perspective.

A stock market index fund has both these properties while delivering broadly similar mean returns to college.

Given this, using the same methodology described above, we first compute the valuation of access to the stock market across all different agent types. The mean valuation for stocks across all types is 4.4 percent of consumption—substantially lower than that of college. This is not entirely surprising, both because stock market participation is low and because an alternative savings instrument is available in the form of the risk-free asset.

Looking beyond simple averages, however, Figure 7 shows the cumulative distribution function (CDF) of valuations for both college and stocks, and it reveals in each case a great deal of heterogeneity. The horizontal axis in the figure shows the value of college and stocks expressed as a percent of consumption (as described above). Note that since the valuations cannot be negative, the valuations range from zero to one. Starting with college, the figure shows that nearly 40 percent of agents place no value on access to college (the CDF of college valuations at zero is about 0.4). At the opposite end, nearly a quarter of the population value access to college at 30 percent or more of consumption (the CDF at 0.3 of consumption is about 0.75). As for stocks, the figure shows that two-thirds of the population assign a near-zero value to the opportunity to invest in the stock market, while less than 10 percent of the population values stocks access at more than 20 percent of consumption.

Figure 7: CDF of Valuations for College and Stocks



We turn next to the question of *who* values access to each investment opportunity. To facilitate comparison, we include the previously-reported valuations for access to college. As shown in the right-most column of Table 6, those in the lowest quartile of the ability distribution find access to stocks much more valuable (9.7 percent of consumption) than access to college (1.0 percent), and

Table 6: Average Valuations by Initial Characteristic

Initial Characteristic	Valuation	
	College	Stocks
<i>Ability</i>		
Low	1.0%	9.7%
Middle	15.1%	2.3%
High	27.5%	3.4%
<i>Human Capital</i>		
Low	13.3%	5.9%
Middle	13.2%	4.6%
High	19.0%	2.6%
<i>Wealth</i>		
Low	9.3%	4.9%
Middle	13.7%	4.6%
High	22.0%	3.6%

it is only when ability rises (i.e., for those with “medium” or “high” ability) that college asserts itself as the more valuable of the two investment opportunities.²⁶

Notice as well that the sensitivity of the value of access to stocks to initial conditions is uniformly (far) lower than it is for access to college. For example, the valuation for stocks access ranges from 9.7 percent for the bottom quartile of ability down to 3.4 percent for the highest quartile—a much smaller range than the corresponding range of valuations for college access. The diminished importance of ability for the value of stocks access can be understood as before—by noting that those who have a high capability of acquiring human capital (high-ability individuals) can augment their earnings and not rely on stocks to generate income or wealth. As for the importance of initial wealth, we see that it plays only a secondary role when compared with the roles played by ability and initial human capital.


²⁶At first glance, our finding that low-ability types value access to stocks more than high-ability types may seem puzzling. However, there are two reasons for this. First, because average consumption is lower for low-ability types than for high-ability types, a smaller absolute valuation for stocks among low-ability types will translate into a higher percentage change in terms of consumption-equivalent welfare. Second, recall that our counterfactual exercise compares the no-stocks economy *where college is still available* to the baseline. For high-ability types, the fact that college remains available means that they still have access to a high-value investment: their value function in the no-stocks economy is therefore not significantly lower than in the baseline. In contrast, with the loss of stocks, low-ability types—most of whom derive no value from college—cannot easily reoptimize by investing in college at high rates when stocks are unavailable. The result is a larger gap in value between the no-stocks economy and the baseline for them.

5.3 The Role of Direct College Subsidies

Our results thus far have revealed high valuations for college in an environment that features large direct subsidies to the direct cost of college. To what extent do these valuations depend on current college subsidies? To answer this question, we consider an economy in which the direct cost of college is not subsidized, i.e., the nearly 50 percent reduction in college costs coming from the subsidy is made unavailable.

Figure 8 reports the CDFs for the valuations of college access in both the baseline economy and the no-subsidy counterfactual. The mean valuation across all types for access to unsubsidized college (9.2 percent of consumption, not shown) is over 5 percentage points lower than in the baseline. It is therefore clear that for many, subsidies are an important part of the high value they derive from college.

Figure 8: CDFs of Valuations for College With and Without Subsidy




However, the valuation of the college subsidy is not uniform across the population. To see this, note for example in Figure 8 that for about 40 percent of the population, the value of access to college is zero *with or without* the presence of the college subsidy. That is, the subsidy makes no difference for these individuals' valuation of college.

Figure 9 further characterizes *who* benefits from the provision of the college subsidy. The figure shows the CDFs of college valuations (with and without the subsidy) separately for the groups corresponding to low, medium, and high levels of ability. This figure confirms that there are sizable differences in the valuation of the college subsidy across agent types. Furthermore, it is only individuals with *high* levels of ability who assign a substantial value to the provision of the subsidy. Perhaps more importantly, the figure also suggests that those individuals who are best

positioned to benefit from the opportunity to invest in college (i.e., individuals with high ability) continue to assign a relatively high value to college access even in the absence of the subsidy. (For example, for the high-ability group, the fraction of individuals who place no value on college access barely moves, from about 9 percent in the presence of the subsidy to only about 12 percent in the absence of it.)

Figure 9: By Ability



In the Appendix, we show that the same result holds when we look at different levels of initial human capital and initial wealth (Figures 11 and 12). In those cases, too, it is only individuals with *high* levels of initial human capital or initial wealth who assign a substantial value to the provision of the subsidy; and those individuals who are best positioned to benefit from the opportunity to invest in college (i.e., those with high initial human capital or high initial wealth) continue to assign a relatively high value to college access even in the absence of the subsidy.

5.4 Targeted Reallocation of College Subsidies

The above counterfactual enables us to classify agents into three groups: those who enroll in college with or without the subsidy (whom we label “always enroll” and who comprise 24 percent of the population), those who enroll in the presence of the subsidy and not in its absence (“switchers,” who account for 30 percent of the population), and those who choose not to enroll in either case (“nonenrollees”, who make up the remaining 46 percent).

The fact that nearly a quarter of the population would continue to enroll in college even in the absence of any subsidy while a full 46 percent receives no benefit from it motivates us to consider whether gains can be had from reorienting funds away from those who do not require subsidies to enroll in college (the “always enroll” group) and towards those who are not positioned to benefit from the college subsidy (the “nonenrollees”). Specifically, we consider an experiment in which college is no longer subsidised for the “always enroll” group. The proceeds (i.e., the total amount of the college subsidy that was accruing to this group) are divided among the “nonenrollee” group in the form of a stock index fund that is available in the first year of retirement.²⁷ The “switchers” continue to receive college subsidies as in the baseline.

Before we turn to the results, we make two important observations about this experiment. First, it is revenue-neutral, as it requires no additional resources beyond what is currently being spent on the provision of college subsidies. Second, and more substantively, it leaves college investment decisions unchanged. Existing college subsidies presumably reflect a public desire to internalize potential positive externalities, or otherwise modify college-attendance decisions. Our approach ensures that these goals continue to be met, as monies are moved without interfering with college investment decisions. In particular, we do not evaluate the myriad policies that would be more “invasive,” such as taking funds away from those whose enrollment decisions depend on the subsidy. As a result, our welfare implications are also easy to interpret: the college premium will not change precisely because human capital allocations themselves will not change.


5.4.1 Findings

Starting with those who always enroll, we find that these agents strongly prefer the baseline to the alternative subsidy regime: the consumption-equivalent value of the status quo for these individuals (relative to the alternative) is 7.7 percent. Conversely, members of the set of “nonenrollees” strongly prefer the alternative subsidy regime to the status quo: they would experience a -5.5 percent loss in moving from the alternative regime to the baseline. Because the size of the nonenrollee

²⁷Note that the per-capita size of the subsidy will differ across the two groups simply because the groups differ in size.

group is nearly double that of those who always enroll, the mean welfare implication of the status quo relative to the alternative regime is slightly negative (-1.07 percent). That is, from behind the veil of ignorance, welfare gains appear available from a more targeted allocation of subsidies that preserves existing college enrollment and is budget neutral.

Figure 10: CDFs of Value of Moving From Alternative Subsidy Regime to Baseline



This finding can also be gleaned from Figure 10, which shows the CDF of the value of moving *from the alternative subsidy regime to the baseline*. The figure highlights the presence of large differences in the valuation of the current regime across the population. The “jump” in the CDF at zero shows that the measure of switchers is substantial (30 percent). It is also easy to see in the figure the group with negative valuations (corresponding to “nonenrollees”) and the group with positive valuations (corresponding to “always enroll”).

Our second main finding from this exercise is that *initial* conditions (a, h_1, x_1) matter systematically and substantially for the impact of policy. This can be seen clearly in Table 7, which reports mean valuations for groups defined, in turn, by a single dimension of initial conditions. The entries report valuations across the bottom, middle two, and top quartile of ability, initial human capital, and initial wealth.

Two points are immediate from the table. First, and naturally, there is a clear positive relationship between attributes that ease college investment and the relative gains from remaining in the status quo. Conversely, we see that it is those with initial characteristics least helpful for college success (e.g., the groups $(a_{Low}, h_{1,Low}, x_{1,Low})$) who suffer from the status quo. These types of individuals would prefer to receive stock-market funds as opposed to the promise of subsidized

Table 7: The Gains and Losses from Reallocating Direct College Subsidies

Initial Characteristic	Welfare Change
<i>Ability</i>	
Low	7.7%
Middle	0.1%
High	-5.3%
<i>Human Capital</i>	
Low	6.7%
Middle	0.0%
High	-2.5%
<i>Wealth</i>	
Low	4.1%
Middle	1.0%
High	-1.9%

college should they choose to enroll. Second, the quantitative strength of both gains and losses is nontrivial. That is, just as with college overall, our model suggests the presence of a wide variation in the value of the status quo policy across the initial characteristics of US high school completers.

Our demonstration of a welfare-improving alternative reallocation of public support towards a simple financial asset is intended as an illustration of the potential misdirection of funds that we may currently be experiencing. While we have argued that indexed financial assets are the sensible alternative, even here there are any number of policies one might consider. We intend, therefore, for this part of the analysis to make two points. First, our finding that the subsidy has the most value to the well-prepared suggests that whatever motivations one places on subsidizing college (e.g., externalities, or the promotion of equality of opportunity) to increase total investment, the current structure of support—universal subsidies conditional on enrollment—does not seem fully consistent with them. Second, while college does not differ from stocks as an investment in terms of mean returns, it clearly does as a target for public support. Our paper suggests that such divergence may warrant further consideration.

6 Concluding Remarks

Using a rich model of human capital investment, we show that the value of access to college varies greatly in the population. While a small group of well-prepared individuals value college access very highly, we estimate that 40 percent of US high school completers place no value on access to college at all. The latter thus derive no benefits from the direct subsidies that currently reduce college costs

sharply. Because receiving these benefits requires enrollment and hinges on college completion, subsidies flow instead to the best-prepared high school graduates, most of whom would continue to enroll in the absence of those subsidies. Even modestly targeted alternatives may therefore improve welfare. We provide one: redirecting subsidies away from those who would nonetheless enroll—into a stock fund for those who do not enroll even with the subsidy—increases ex-ante welfare by 1 percent of mean consumption. Unlike college subsidies, this alternative arrangement benefits the large group that is poorly poised for collegiate success, while yielding mean ex-ante returns comparable to the mean returns accruing to college completers.

To prevent any misinterpretation of our findings, we stress that they must not be read as any kind of sweeping statement about, or indictment of, college education. Instead, they are suggestive of the importance of college *readiness*. Our results show that heterogeneity in college readiness (as summarized in the pair (a, h_1)) drives heterogeneity in college returns, so much so that poor preparedness almost fully nullifies the high ex-post (i.e., conditional on college completion) payoffs of college. Put another way, the high current payoffs to college completion contain a clear signal about the importance of pre-college preparation. Our findings provide additional perspective—in line with the large body of work of James Heckman and coauthors²⁸—on why early-childhood environments are critical in determining the effectiveness with which individuals can acquire human capital. They may well hold the key to helping individuals unlock the benefits of college.

References

- Abbott, B., G. Gallipoli, C. Meghir, and G. L. Violante (2018). Education policy and intergenerational transfers in equilibrium. *Forthcoming Journal of Political Economy*.
- Acemoglu, D. and J. Angrist (2000). How large are human-capital externalities? evidence from compulsory schooling laws. *NBER macroeconomics annual 15*, 9–59.
- Akyol, A. and K. Athreya (2005). Risky Higher Education and Subsidies. *Journal of Economic Dynamics and Control 29*(6), 979–1023.
- Altonji, J., P. Arcidiacono, and A. Maurel (2016). Chapter 7 - the analysis of field choice in college and graduate school: Determinants and wage effects. Volume 5 of *Handbook of the Economics of Education*, pp. 305 – 396. Elsevier.

²⁸See, for example, Heckman (2006) and Heckman et al. (2013).

- Altonji, J. G., E. Blom, and C. Meghir (2012). Heterogeneity in human capital investments: High school curriculum, college major, and careers. *Annu. Rev. Econ.* 4(1), 185–223.
- Arcidiacono, P. (2004). Ability sorting and the returns to college major. *Journal of Econometrics* 121(1), 343–375.
- Arcidiacono, P. (2005). Affirmative action in higher education: how do admission and financial aid rules affect future earnings? *Econometrica* 73(5), 1477–1524.
- Arcidiacono, P., V. J. Hotz, and S. Kang (2012). Modeling college major choices using elicited measures of expectations and counterfactuals. *Journal of Econometrics* 166(1), 3–16.
- Athreya, K., F. Ionescu, and U. Neelakantan (2015). Stock market participation: The role of human capital. Working paper, Federal Reserve Bank of Richmond.
- Athreya, K. B. and J. Eberly (2013). The supply of college-educated workers: the roles of college premia, college costs, and risk. Working paper, Federal Reserve Bank of Richmond.
- 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.
- Bach, L., L. E. Calvet, and P. Sodini (2016). Rich pickings? risk, return, and skill in the portfolios of the wealthy.
- Ben-Porath, Y. (1967). The production of human capital and the life cycle of earnings. *Journal of Political Economy* 75(4), 352–365.
- Blandin, A. (2018). Learning by doing and ben-porath: Life-cycle predictions and policy implications. *Journal of Economic Dynamics and Control* 90, 220–235.
- 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.
- Browning, M., L. P. Hansen, and J. J. Heckman (1999). Micro data and general equilibrium models. *Handbook of macroeconomics* 1, 543–633.

- Card, D. (2001). Estimating the Return to Schooling: Progress on Some Persistent Econometric Problems. *Econometrica* 69(5), 1127–1160.
- Castaneda, A., J. Diaz-Gimenez, and J.-V. Rios-Rull (2003). Accounting for the us earnings and wealth inequality. *Journal of political economy* 111(4), 818–857.
- Cestau, D., D. Epple, and H. Sieg (2017). Admitting students to selective education programs: Merit, profiling, and affirmative action. *Journal of Political Economy* 125(3), 761–797.
- Chatterjee, S. and F. Ionescu (2012). Insuring Student Loans Against the Financial Risk of Failing to Complete College. *Quantitative Economics* 3(3), 393–420.
- Ciccone, A. and G. Peri (2006). Identifying human-capital externalities: Theory with applications. *The Review of Economic Studies* 73(2), 381–412.
- Cocco, J. (2005). Portfolio choice in the presence of housing. *Review of Financial Studies* 18(2), 535–567.
- Cocco, J., F. Gomes, and P. Maenhout (2005). Consumption and portfolio choice over the life cycle. *Review of Financial Studies* 18(2), 491–533.
- Epple, D., R. Romano, S. Sarpça, and H. Sieg (2013, August). The u.s. market for higher education: A general equilibrium analysis of state and private colleges and public funding policies. Working Paper 19298, National Bureau of Economic Research.
- Fagereng, A., L. Guiso, D. Malacrino, and L. Pistaferri (2016, November). Heterogeneity and persistence in returns to wealth. Working Paper 22822, National Bureau of Economic Research.
- Garriga, C. and M. Keightley (2007). A general equilibrium theory of college with education subsidies, in-school labor supply, and borrowing constraints. federal reserve bank of st. Technical report, Louis Working Paper Series 2007-051A.
- 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., R. Pinto, and P. Savelyev (2013). Understanding the mechanisms through which an influential early childhood program boosted adult outcomes. *American Economic Review* 103(6), 2052–86.
- Heckman, J. J. (2006). Skill Formation and the Economics of Investing in Disadvantaged Children. *Science* 312(5782), 1900–1902.

- Hendricks, L. and O. Leukhina (2014). The return to college: Selection and dropout risk.
- Hendricks, L. and O. Leukhina (2017). How risky is college investment? *Review of Economic Dynamics* 26, 140 – 163.
- Hendricks, L. and T. Schoellman (2014). Student abilities during the expansion of us education. *Journal of Monetary Economics* 63, 19–36.
- 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), 265–290.
- Ionescu, F. (2009). The Federal Student Loan Program: Quantitative Implications for College Enrollment and Default Rates. *Review of Economic Dynamics* 12(1), 205–231.
- Ionescu, F. and N. Simpson (2016). Default risk and private student loans: Implications for higher education policies. *Journal of Economic Dynamics and Control* 64, 119–147.
- Johnson, M. T. (2013). Borrowing constraints, college enrollment, and delayed entry. *Journal of Labor Economics* 31(4), 669–725.
- Judd, K. L. (2000). Is education as good as gold? a portfolio analysis of human capital investment. *Unpublished working paper. Hoover Institution, Stanford University.*
- Kim, H. H., R. Maurer, and O. S. Mitchell (2013, December). Time is money: Life cycle rational inertia and delegation of investment management. Working Paper 19732, National Bureau of Economic Research.
- Loke, V. and M. Sherraden (2009). Building assets from birth: A global comparison of child development account policies. *International Journal of Social Welfare* 18(2), 119–129.
- Manski, C. F. and D. A. Wise (1983). *College choice in America*. Harvard University Press.
- McGrattan, E. R. and E. C. Prescott (2000). Is the stock market overvalued? *Federal Reserve Bank of Minneapolis Quarterly Review* 24(4), 20–40.
- Mehra, R. and E. C. Prescott (1985). The equity premium: A puzzle. *Journal of Monetary Economics* 15(2), 145–161.

- Moretti, E. (2004a). Human capital externalities in cities. In *Handbook of regional and urban economics*, Volume 4, pp. 2243–2291. Elsevier.
- Moretti, E. (2004b). Workers’ education, spillovers, and productivity: evidence from plant-level production functions. *American Economic Review* 94(3), 656–690.
- 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.
- Peri, G. and C. Sparber (2009, July). Task specialization, immigration, and wages. *American Economic Journal: Applied Economics* 1(3), 135–69.
- Planty, M., W. Hussar, T. Snyder, S. Provasnik, G. Kena, R. Dinkes, A. KewalRamani, and J. Kemp (2008). The condition of education 2008.
- Psacharopoulos, G. and H. A. Patrinos (2018). Returns to investment in education: a decennial review of the global literature. *Education Economics* 26(5), 445–458.
- Restuccia, D. and C. Urrutia (2004). Intergenerational persistence of earnings: The role of early and college education. *The American Economic Review* 94(5), 1354–1378.
- Stange, K. M. (2012). An empirical investigation of the option value of college enrollment. *American Economic Journal: Applied Economics* 4(1), 49–84.

A Appendix

In this Appendix, we present additional results conveying the importance of individual-level characteristics for the valuations individuals place on access to college relative to the alternative we study. Figures 11 and 12 show the CDFs of college valuations—with and without the direct college subsidy—separately for groups corresponding to low, medium, and high levels of initial human capital and initial wealth. See Section 5.3.

Figure 11: By Initial Human Capital

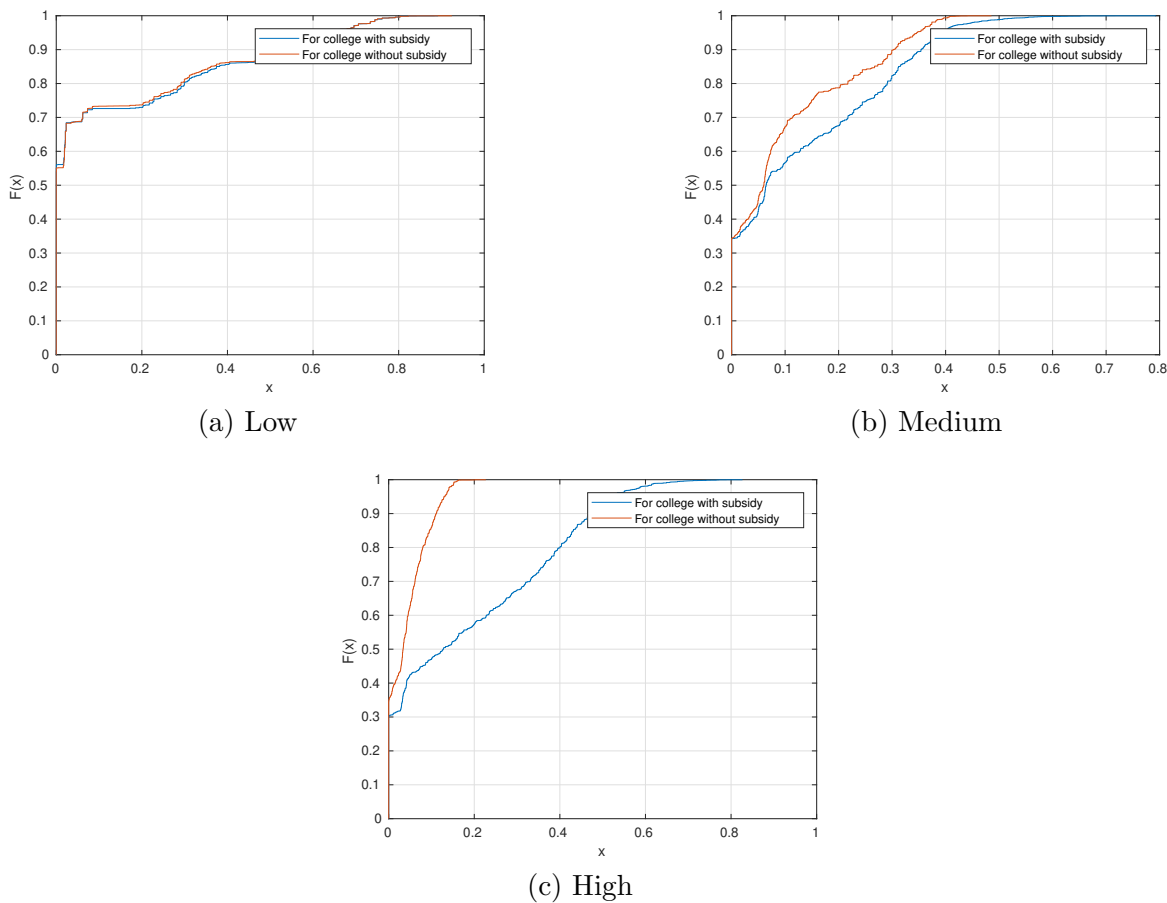
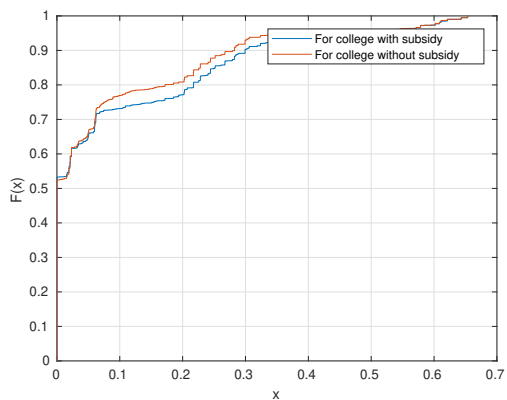




Figure 12: By Initial Wealth



(a) Low



(b) Medium



(c) High