Understanding the Role of Colleges in Intergenerational Mobility

Vinzenz J. Ziesemer†
European University Institute

Preliminary and Incomplete Draft

April 9, 2017

Abstract

Recently available data from administrative sources show large differences in upward mobility rates by the colleges students attend. While expensive and selective colleges are more successful at propelling low income students into high-earning careers, far fewer students from low-income families enter these colleges. Intergenerational mobility is a stated goal of education policies (such as college subsidies, student grants, and loans). Which types of colleges should policy makers target, and how? This paper sets up a heterogeneous agent life-cycle model with overlapping generations that are linked by monetary transfers and the transmission of ability. Colleges are modeled as a learning technology that takes ability, money, and time as inputs. The intergenerational transmission of ability and current education policies (by both colleges and different levels of government) in the model are directly informed by micro data. The remainder of the model is carefully parameterized to match key features of the life-cycle of earnings and the US economy. First, the paper shows how the model captures the relation between college spending and intergenerational mobility. Second, the paper considers a number of policy counterfactuals that illuminate the effect education policies have on intergenerational mobility.

†Email: vinzenz.ziesemer@eui.eu. Telephone: +31 17 66 80 36. Address: Villa La Fonte, Via delle Fontanelle 18, 50014 San Domenico di Fiesole, Italy.

1 Introduction

Measures of intergenerational mobility are indicators of equality of opportunity: to what extent do children from poor families have the same chances in life as their richer peers? The empirical economic literature on these measures has long focused on country-wide indicators: in which countries is intergenerational mobility large (Corak, 2013)? Has intergenerational mobility fallen over time (Mendes Tavares and Urrutia, 2016)?

The most recent wave of empirical work looks beyond these aggregates and considers their underlying heterogeneity. What determines regional variation in intergenerational mobility (Chetty, Hendren, Kline, and Saez, 2014)? Is intergenerational mobility uniform across the income distribution (Landersø and Heckman, 2017)? And, finally, how does intergenerational mobility look over the distribution of colleges?

Recently available data on that final question motivate this paper. As is discussed in section 2, recent work by Chetty, Friedman, Saez, Turner, and Yagan (2017) has shown that intergenerational mobility is almost flat once colleges are controlled for. At the same time, only two characteristics seem to really relate colleges to intergenerational mobility. On the one hand, higher ability students go to colleges whose students are more successful on the labor market. On the other hand, these colleges spend more resources on students and are therefore more expensive. A question remains: to what extent is mobility determined by money versus ability?

This paper captures the relation between ability, money, education, and intergenerational mobility in a model. The model is a a heterogeneous agent life-cycle model with overlapping generations that are linked by monetary transfers and the transmission of ability. Colleges are modeled as a learning technology that takes ability, money, and time as inputs. The intergenerational transmission of ability and current education policies (by both colleges and different levels of government) in the model are directly informed by micro data. The remainder of the model is carefully parameterized to match key features of the life-cycle of earnings and the US economy.

Using this model, the paper than asks a number of policy questions. To what extent is the relationship between instructional expenditures and intergenerational mobility driven by ability, and to what extent by money? A host of current policies, both by educational institutions and various levels of government seek to alleviate financial need, sometimes discriminating based on need or ability. To what extent do each of these policies increase or decrease intergenerational mobility? And finally, where on the spectrum of colleges would further policy intervention be most effective?

Section 3 containts a full description of the model economy. Section 4 describes how the parameters of the model are estimated or calibrated. Section 5 explores the ability of the model to match aspects of the data that were not targeted in the calibration procedure. Section 6 describes policy counterfactuals and their results. Section 7 concludes.

1.1 Literature

The empirical work underlying this paper is recent and referenced throughout this paper. The main reference is Chetty, Friedman, Saez, Turner, and Yagan (2017), who do extensive work in combining a number of existing (in part administrative) datasets to link information on the earnings of two generations, their educational choices, and college characteristics.

There is a macroeconomic literature that tries to understand how education and education policies are related to intergenerational mobility. Lee and Seshadri (2014) argue that a rich life-cycle model with intergenerational links explains a number of intergenerational relationships well, in particular the intergenerational elasticity of earnings. They focus more on development of human capital during childhood, and less on college heterogeneity. In particular, they do not allow for heterogeneous spending on higher education.

Holter (2015) similarly builds a quantitative model of intergenerational mobility. He then investigates the extent to which differences in tax and education policies can explain cross-country differences in intergenerational mobility. This paper features a richer model of intergenerational links, and takes a more granular look at education policies.

Positive implications of education policies on college enrollment are studied by a number of authors. Important papers are by Lochner and Monge-Naranjo (2011b), who consider the effect of student loan policies on the college entry decision of youth that is heterogeneous in ability and family income. Abbott et al. (2013) study the decision to go to college or not in a quantitative model with intergenerational transfers, and find that these transfers are an important adjustment margin that dampen the effects of education policies in equilibrium. Again, they only consider the extensive margin of college entry. Neither of these papers focus on intergenerational mobility. Empirical work on the incidence of financial constrainedness is sumarized in Lochner and Monge-Naranjo (2011a), who find increased evidence for such incidence in recent years.

There is a large normative literature on education policies, often in combination with taxation. Of these, the work of Krueger and Ludwig (2016) is closest to this paper. Their focus is on optimal taxation with (almost) linear instruments, whereas this paper considers the positive impact of actual policies through counterfactuals. Their paper also includes intergenerational transfers, but only has an extensive margin for education (a college-entry choice) and instead considers general equilibrium effects and the importance of the transition between different policy regimes. Bovenberg and Jacobs (2005) find that while education subsidies themselves distribute resources to the well to do, their optimal level is positively related to tax rates. This is because they undo the disincentive effects of taxation on human capital formation. That issue has been been studied in a dynamic theoretical framework by Stantcheva (2015), and in a quantitative framework by Hanushek, Leung, and Yilmaz (2003). Kapicka and Neira (2015) study optimal tax rates when human capital investment itself is risky. In an incomplete market where students cannot borrow against future income, there is a role for government-provided student loans. These are studied in a dynamic framework by Findeisen and Sachs (2016).

2 The Empirics of Colleges and Intergenerational Mobility

Chetty et al. (2017) combine data from federal income tax returns and from the Department of Education to link information on the earnings of two generations to their educational choices, and the characteristics of colleges they attend. These data are available for individuals from the 1980-1982 birth cohorts. In 2014, the time of the last earnings measurement, children of those cohorts were in their early 30s, at which point measures of intergenerational earnings persistence typically stabilize. Parental income is defined as average parental earnings when the children are aged 15 to 19, and child income is measured over the year 2014. The college a student was enrolled in longest counts as the college that the student attended. Their preferred measure of intergenerational mobility induced by a college is the number of students from the bottom 20% by parental earnings that end up in the top 20% of earnings, which they dub the *mobility rate*.

Chetty et al. (2017) first find that, controlling for the college students go to, intergenerational mobility is actually quite large: while there is a strong correlation between parental earnings and child earnings for the United States as a whole, this correlation is only one-third within a given college, and this is true rather irrespective of the type of college. This is illustrated in figure 1. This adds to an established literature in economics that suggests human capital is the main culprit in intergenerational mobility, and suggests that once we understand educational outcomes, we will largely understand intergenerational mobility.

A logical next step to ask, then, is which colleges are particularly good at inducing intergenerational mobility. Interestingly, Chetty et al. (2017) find little relation between their preferred measure of intergenerational mobility and any measures of the average student. In particular, focusing on the average student's SAT score (which is generally considered the best measure of college readiness), mobility is flat across the board. Looking closer, they

Figure 1: Controlling for Colleges

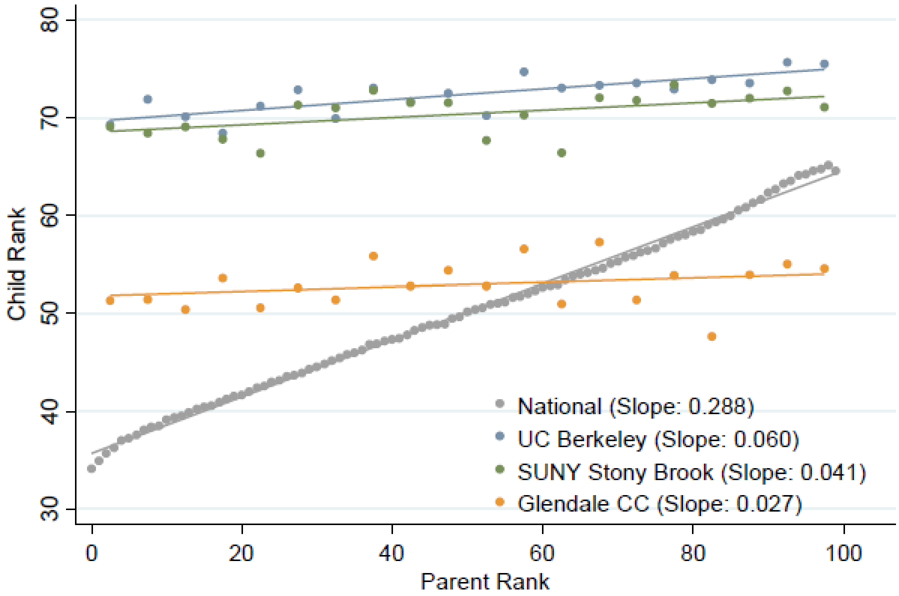

Figure 3 from Chetty et al. (2017): parent versus child earnings ranks.

find that this is easily explained: while colleges with low average SAT scores take many students from low income families (as defined by the share of students from the bottom 20% of parental earnings, a measure of *access*), relatively few of these make it to the top of the earnings distribution (as defined by the top 20% of earnings, a conditional measure of *success*). At more selective colleges, the situation is reversed. Unconditional intergenerational mobility turns out flat as a result. This is displayed in the three left graphs of figure 2.

The same patterns hold true when comparing colleges by their annual cost of attendance (the right three graphs): the expensive colleges are more successful when in it comes to propelling students from low income families to high earning careers, yet they also take less of these students. Thus, it is hard to tell to what extent it is money that creates success, and to what extent it is due to the ability of entering students. At the same time, the relation between college preparedness and spending remains imperfect. This, combined with the existing empirical evidence on financial constrainedness, suggests that money likely plays some role in making the more successful colleges accessible to low income students. So how would intergenerational mobility look without financial constraints? Or without the policies that are in place to alleviate such constraints? To answer such questions, we have a need for model-based counterfactuals.

After accounting for these headline variables, much heterogeneity remains in intergenerational mobility by college. Early results from the literature indicates that this is reflective

Figure 2: Intergenerational Mobility by Colleges

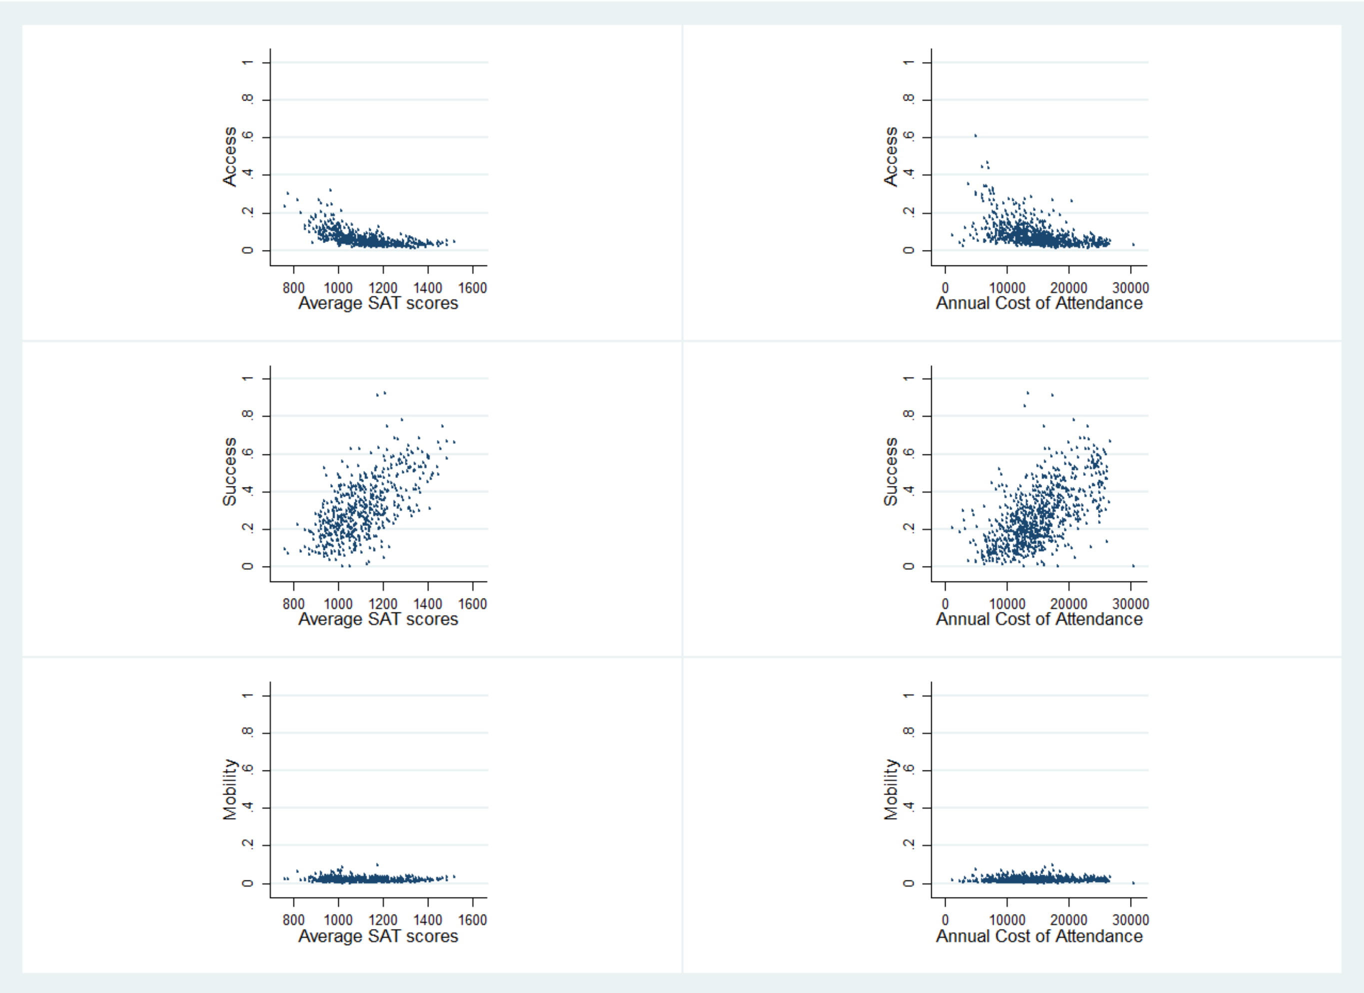

Author calculations based on data from Chetty et al. (2017). Includes all private 4-year colleges for which data are available. Data on colleges from the year 2000 or 2001. Annual cost of attendance is an average. Mobility measures are described in the main text.

of factors that fall outside the scope of this paper. For example, children of immigrants are more upwardly mobile, and often concentrate in certain geographies. Other geographical specifics may also be important. A final often-cited source of heterogeneity is the choice of college major, but interestingly Chetty et al. (2017) already indicate that this is not a strong explanatory factor.

The above graphs focus on average annual costs, but there are some important steps to take when connecting student choices to costs. First, besides loan programs there are many need-based grants from government sources, which we will directly measure from other data when calibrating the model. Second, there are institutional grants, which we can again take from the data.These grants are often merit- and need-based. Third, there are some public colleges that receive substantial government subsidies.

That last point ties into another issue: once we understand the cost of attending college, we need to understand how these costs translate into the educational investment that creates human capital and, eventually, earnings. We tackle this issue by using the data of Chetty et al. (2017), taking a college's instructional expenditures per student as a measure of educational investment. Figure 3 displays average cost of attendance and instructional expenditures per student, for both public and private colleges. A clear pattern emerges. First, higher costs lead to higher instructional expenditures, but the relationship is far from one-to-one. Second, public colleges fill the low-cost space and provide instructional expenditures at about the level of the cheaper (but still more expensive) private colleges. In doing so they significantly flatten the relationship between the cost of education and educational investment at the lower end. We later use these data to calibrate our model.

Figure 3: From college costs to educational investment

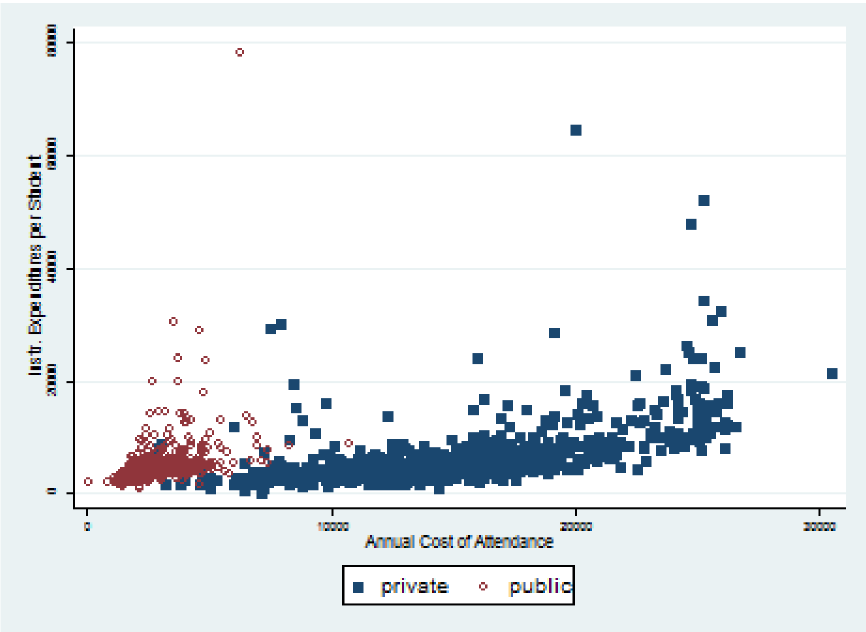

Author calculations based on data from Chetty et al. (2017). Includes all public and private 4-year colleges for which data are available. Data on colleges from the year 2000 or 2001. Annual Cost of Attendance and Instructional Expenditures are averages.

So how should we think of colleges? In this paper, we take straightforward approach: when in college, human capital is formed by ability, time and money. A college is then simply a technology that, given some level of cost, allows a student of given ability to invest money and time into education. How costs are converted into monetary investment follows directly from the data. Throughout the paper, we will take the behavior of colleges as given.

An existing literature explicitly models the behavior of colleges. For example, Epple, Romano, and Sieg (2006) analyze a model with quality maximizing colleges, peer-effects and a preference for low-income students, where price discrimination leads to student sorting over colleges. Epple, Romano, Sarpça, and Sieg (2013) then adapt this framework to include public universities and endow students with idiosyncratic preferences over colleges.

Several of these motives could be important for our paper as well. First, colleges indeed engage in price discrimination, offering discounts to lower income and higher merit students. They can do so even in a competitive market when other colleges have the same preferences as they do, or simply using their endowment income. Thus, these motives do not necessarily contradict our setup. Institutional grants themselves are taken into account in this paper too when we estimate the relationship between student characteristics and the price of educational investment. Second, peer-effects would suggest that sorting is an important issue, but a large quasi-experimental literature has found peer-effects of all kinds to be small (Angrist, 2014), so that it seems safe to ignore them. Third, colleges pricing power would diminish if the number of colleges between which students can choose is large, which it is at all spending levels. This is less so if students have idiosyncratic preferences over colleges (how much less depends on functional form assumptions), but there is no direct empirical evidence for that. Instead, recent work by Hoxby (2016) (see also Hoxby (2015)) suggests that returns to education are rather flat across the distribution of colleges.

3 Model Economy

This paper considers stationary equilibria. There is a continuum of agents with mass one. Each agent spawns a new agent with a mass identical to its own. We refer to the former as parents and the later as children. The timing of the life-cycle is deterministic and equal for all households. The symbol $'$ is used to denote choice variables pertaining to an agent's children.

Each agent goes through a life-cycle from age 0 to age T, representing his working life. There are two special phazes in this life cycle. First, the agent has access to colleges and, if he chooses to enrol in a college, a system of student grants and loans in the first period of his life. Second, at a later point in life (age t^I) he makes an inter-vivos transfer to his children, who begin their life-cycle in the following period. The agent does so because he values the child's expected discounted lifetime utility (at a rate potentially lesser than its own, i.e. the parent is impurely altruistic).

In each period, an agent can use his time to work, enjoy leisure, or to invest in human capital. In any period, he can use his resources to consume, save, and repay student loans. At age 0, he chooses whether to go to college or not, and if so how much to invest in a college education At age t^I he can make an inter-vivos transfer. All resources are expressed in terms of consumption, which is also the model's numeraire. Markets in the model are incomplete in the sense of a Bewley-Aiyagari-Huggett model: agents face idiosyncratic income risk that they cannot insure against. They can borrow using (student) loans and save using a risk-free asset, but face borrowing constraints that potentially constrain their consumption and human capital investment.

Compared to most models of college choice, human capital is continuous in this paper, and formed by a Ben-Porath (1967) accumulation function that takes both time and money as inputs. This allows the model to capture the full effect of education and policies, rather than just the effect on those at the margin of college entry. Such a margin is included too. Those who choose not to go to college or are no longer in college accumulate human capital by a function that is also of the Ben-Porath (1967) type, but one that takes only time as an input and has separately identified parameters. Both functional forms have been shown to be successful in capturing many aspects of human capital formation in the existing macroeconomic literature.

Each agent in the model economy is linked to their parents in three ways. First, agent's ability to accumulate human capital is correlated with that of their parents. Second, parents endogenously decide how much financial resources to make available to their children as they make initial decisions on human capital investment. Third, the availability of grants and loans depends on parental wealth, as US education policies have historically been means-dependent. These mechanisms are important in assessing the impact of education policies on human capital investment decisions: when family income-dependent policies change and make more or less resources available, parental transfers are a major compensating margin. And the more persistent ability is across generations, the more correlated wealth and ability will be, reducing the influence of education policies.

The economy contains detailed features of the policy environment in the United States, in particular: taxes, educational subsidies and grants, and student loans: average labor tax rates are non-linear and based on the US tax code, as are other taxes. Section **??** provides a detailed overview of student aid in the United States in 2003, the year to which the model will be calibrated. The Stafford loan system is explicitly modeled in this paper. To capture subsidies and grants from institutions and all levels of government, the model employs a flexible specification that allows estimation of these items directly from the data. All this is important for two reasons. First, as Landersø and Heckman (2017) argue, policies are key to understanding students' incentives to go to college, and therefore have a first-order impact on intergenerational mobility. Second, it is precisely these policies that we will later evaluate.

3.1 Individual's problem

Let s_t denote the stochastic state of the agent's life-cycle at age t, and s^t a history of stochastic states up to age t: $s^t = [s_t, s_{t-1}, \ldots, s_1, s_0]$. These histories are suppressed in most of the below, but made explicit where the arguments of the maximization problem are listed.

In the below, c denotes consumption, l leisure, e time investment in human capital, d goods or monetary investment in human capital, a assets, b student loans, and v inter-vivos transfers. k denotes the choice between going to college ($k = 1$) or not ($k = 0$). $I_{[x]}$ is an indicator function that equals one when x is true and zero otherwise. q_0 is permanent parental income, which is described in further detail below. The same goes for student loan repayment functions B and borrowing constraints $\underline{a}$ and $\underline{b}$. $\mathbb{E}$ is the usual expectations operator. Denote the vector of control variables $\mathbf{z}(s^t)$:

$$\mathbf{z}(s^t) = [c_t(s^t), l_t(s^t), e_t(s^t), a_{t+1}(s^t), b_{t+1}(s^t)].$$

The problem is then as follows:

$$\begin{aligned}
V(\alpha, q, a_0) = \max_{\substack{\left\{\left\{\mathbf{z}(s^t)\right\}_{s^t}\right\}_{t=0}^{T-1},\\ \left\{k(s^0), d(s^0)\right\}_{s^{t0}},\\ \left\{v(s^{t^I})\right\}_{s^{t^I}}}} & \mathbb{E}_{s^t,\alpha'} \left\{ \sum_{t=0}^{T-1} \beta^t \frac{(c_t^{\nu} l_t^{1-\nu})^{1-\sigma}}{1-\sigma} + \omega \beta^{t^I} V(\alpha', q', v) \right\} \\
& \text{subject to } \forall t \in \{0, \ldots, T-1\}: \\
& c_t(1+\tau_c) \leq (1 - l_t - e_t) w h_t(d_t, e_t, h_{t-1}, \alpha) x(s_t)(1 - \tau_n(\cdot)) - d_t - v I_{[t=t^I]} \\
& \qquad + a_t(1 + r(1-\tau_a)) - a_{t+1} + B_t(b_t) - b_{t+1} \\
& c_t, d_t, v \geq 0, \quad 0 \leq l_t, e_t \leq 1, \quad l_t + e_t \leq 1, \quad k \in \{0, 1\} \\
& a_{t+1} \geq -\underline{a_t}(q), \quad 0 \geq b_{t+1} \geq -\underline{b_t}, \quad b_0 = 0, \quad a_{t+1} b_{t+1} = 0.
\end{aligned}$$

Human capital is formed using time e_t, and also money $\tilde{d}(d)$ when and if the agent is in college ($k = 1$). The functional form is as follows:

$$h_{t+1} = \begin{cases} h_t(1-\delta_h) + \alpha(e_t h_t)^{\beta_1} (\tilde{d}(d, q, h_t))^{\beta_2} & \text{if} \quad k = 1 \text{ and } t = 0 \\ h_t(1-\delta_h) + \alpha(e_t h_t)^{\beta_0} & \text{otherwise.} \end{cases}$$

These are much-used variations of the well-known Ben-Porath (1967) function, which has been shown to capture a number of stylized facts regarding the life-cycle of education investments. We would expect that $\beta_0, \beta_1, \beta_2 > 0$. We further assume that h_0 and α_0 are perfectly correlated, so that only one of them is a state variable at the start of life. The monetary input $\tilde{d}$, importantly depends on the choice variable d but is not the same: here is where we account for institutional aid as well as student grants at the local, state, and federal level. That is why $\tilde{d}$ also depends on human capital h and on q, permanent parental income, which is determined by the previous generation: $q' = wh$.[1] Grants can only be used towards the cost of education, a feature of current policies that Lochner and Monge-Naranjo (2011b) argue is important in understanding how ability and parental wealth are connected to college choice. (Particularly, they argue that this feature explains the positive relationship between ability and college attendance given family income for constrained individuals.) The discussion of the calibration of the model further elaborates on these points.

Leisure and consumption enter periodic utility multiplicatively. The utility function is tied down by parameters σ and ν. With this functional form, the elasticity of inter-temporal substitution is given by $\frac{1}{1-\nu(1-\sigma)}$, and the Frisch elasticity by $\frac{1-\nu(1-\sigma)}{\sigma(l/(1-l))}$.

The agent is uncertain about the next realization of his idiosyncratic earnings state s_t, which follows a first-order discrete Markov process with transition matrix $\Gamma_s(s_t, s_{t+1})$. All agents start out from the same state: $s_0 = \bar{s}$. The earnings shock $x(s_t)$ combines with his human capital h_t and the wage w of human capital to determine his individual wage rate.

The agent is also uncertain about his child's ability, and remains so throughout. Ability is discrete and drawn from the joint distribution of parents' and children's ability $\Gamma_\alpha(\alpha, \alpha')$.

A government charges taxes on consumption τ_c, labor income $\tau_n((1 - l_t - e_t)wh_t x(s_t))$, and capital income τ_a. Labor income taxes are non-linear in labor income. The government's budget, after consideration of education policies, is balanced by neutral (or wasteful) spending G that does not influence any of the economy's variables.

In it's most general version, the student loan system mimics the 2003 Stafford loan system as follows. At age 0, college-going students fall into one of three eligibility categories q on the basis of their parents' earnings at the time they become independent decision makers. If parental permanent earnings (q) are not higher than y^* (the 'poorest' category) and the household qualifies for subsidized loans up to $\underline{b}^s$ as well as unsubsidized loans up to $\underline{b}^u$. If parental permanent earnings are above y^* but not higher than y^{**} the household can only borrow at the unsubsidized rate up to $\underline{b}^s + \underline{b}^u$. For the remaining households their parents' high earnings enable them to borrow up to $\underline{a}^p$ on the private market at rates generally below the subsidized rate (and not subject to the same repayment schedules as student loans from

[1] In reality, policies are heterogeneous across colleges and states, but typically depend on a number of indicators of families' ability to pay for college. Here, the permanent component of parental income is used as a parsimonious proxy. Transitory components would have the potential to make the problem non-convex because parents could adjust their choices to make their children qualify for student aid (which is something policy makers indeed attempt to rule out).

the government). Access to these loans do not depend on whether they enroll in college or not. If they need to borrow more, they also have access to the same credit from the government as those in the category below. Interest rates r^s and r^u are set exogenously. Interest on subsidized loans is forgiven during the period in which they are paid out. Otherwise, agents cannot borrow at age 0. The model also imposes that those who take out student loans do not save assets at the same time, which is captured by the complementarity constraint $a_{t+1}b_{t+1} = 0$. This structure essentially follows Abbott, Gallipoli, Meghir, and Violante (2013). Thereafter, the natural borrowing constraint applies: all loans must be repaid by the end of working life.

After the colege-going period, students pay down their debt by a constant amount π every period for m periods. Since pay-down is linear, we can provide an analytical solution for π_t. When $1 \leq t < 1+m$ and $b_1 < 0$:[2]

$$\pi_t = \begin{cases} -\frac{r^s}{1-(1+r^s)^{-m}} b_1 & \text{if} \quad q \leq y^* \text{ and } -\underline{b}^s \leq b_1 \\ \frac{r^s}{1-(1+r^s)^{-m}} \underline{b}^s - \frac{r^u}{1-(1+r^u)^{-m}} (b_1 + \underline{b}^s) & \text{if} \quad q \leq y^* \text{ and } b_1 < -\underline{b}^s \\ -\frac{r^u}{1-(1+r^u)^{-m}} b_1 & \text{if} \quad y^* < q \text{ and } b_1 < 0. \end{cases}$$

Otherwise, $\pi_t = 0$.

Matching all this to the notation above:

$$\underline{a}_t(q) = \begin{cases} 0 & \text{if} \quad q \leq y^{**} \text{ and } t = 0 \\ \underline{a}^p & \text{if} \quad y^{**} < q \text{ and } t = 0 \\ 0 & \text{if} \quad t = T-1 \\ \infty & \text{otherwise,} \end{cases}$$

and

$$\underline{b}_t = \begin{cases} \underline{b}^s + \underline{b}^u & \text{if} \quad t = 0 \\ 0 & \text{otherwise.} \end{cases}$$

The student loan repayment function looks as follows:

$$B_t(b_t) = \begin{cases} b_t & \text{if} \quad t = 0 \\ -\pi_t & \text{otherwise.} \end{cases}$$

All of this only holds for college-going students. Non-college going students cannot borrow.

[2]Note that b_1 is a negative number, while $\underline{b}^s$ and $\underline{b}^u$ are positive.

3.2 Stationary Equilibrium

The production function takes the following functional form:

$$F(K,H) = K^{\theta} H^{1-\theta}.$$

Here, H denotes the aggregate effective supply of human capital hours. θ is the capital share of total factor income.

Labor and goods markets are perfectly competitive. We model the economy as closed to labor and goods, and open to capital. This reduces the number of general equilibrium conditions that must be cleared numerically, and is arguably as realistic as assuming an economy that is entirely closed to capital. Additionally, general equilibrium effects through capital formation are by no means a focus of this paper.

Firms borrow capital from households, who receive an international real interest rate $\tilde{r}$ and loose a share δ of their capital to depreciation. This share is exempt from capital taxation. The openness assumption yields an equilibrium condition relating the capital-labor ratio to the exogenous interest rate, which, together with the income share of labor ties down the marginal product of labor.

For simplicity, all student grants are assumed to be under the control and paid for by the model's government, including institutional aid.

Let $x^*_\tau(\iota_t)$ denote a decision rule given states $\iota_t \in \mathcal{I}_t$ for choice variable x_τ. Let $\mathbb{I}_t$ denote a generic subset of the Borel sigma algebra of age-specific state-space $\mathcal{I}_t$.

Definition 1. *A* ***stationary equilibrium*** *of the model economy is defined as:*
wages w;
allocations K, H;
government spending G;
decision rules, each $\forall \iota_t \in \mathcal{I}_t$ whenever they are defined for t, for consumption $\{c_t(\iota_t)\}_{t=0}^{T-1}$, leisure $\{l_t(\iota_t)\}_{t=0}^{T-1}$, assets $\{a_{t+1}(\iota_t)\}_{t=0}^{T-1}$, goods $d_0(\iota_0)$ and time $\{e_t(\iota_t)\}_{t=0}^{T-1}$ investment in human capital, college choice $k(\iota_0)$ (equal to 1 if the agent enrolls in a college, 0 if not), and the inter-vivos transfer $v'(\iota_{t^I})$;
age-specific measures $\lambda_t(\mathbb{I}_t)$, and the resulting overall measure $\lambda(\mathbb{I})$ on $\mathbb{I} \in \times\mathcal{I}_t$;
such that given *international interest rates $\tilde{r}$, tax functions τ_c, τ_n, τ_a, τ_i, sets A and S, transition matrices $\Gamma_\alpha(\hat{\alpha}, \alpha)$ and $\Gamma_s(s_t, s_{t+1})$, as well as the parameters of the model,* ***the following holds****:*

- *the decision rules solve the households' problem as described in subsection 3.1;*

- *the firms make profit maximizing decisions; as a result, their profits are zero and prices of the inputs to production equal their marginal products:*

$$\tilde{r} = F_1(K, H) - \delta_a,$$

$$w = F_2(K, H);$$

- $\lambda(\mathbb{I})$ *is a fixed point of its law of motion that is generated by the following:*
  - *the decision rules of the households,*
  - *the laws of motion for assets and human capital,*
  - *the transition matrices of productivity shocks* $\Gamma_s(s_t, s_{t+1})$*,*
  - *the distribution over the initial states at independence which is consistent with* $\Gamma_\alpha(\hat{\alpha}, \alpha)$*, parental wealth, and the decisions made by parents on schooling and inter-vivos transfers;*
- *the market for labor clears:*

$$H = \sum_{t=0}^{T-1} \int_{\mathcal{I}_t} (1 - l_t - e_t) x_t h_t \, \mathrm{d}\lambda_t;$$

- *the market for goods clears (net investment in assets is zero since the equilibrium is stationary):*

$$F(K, H) = \sum_{t=0}^{T-1} \int_{\mathcal{I}_t} c_t \, \mathrm{d}\lambda_t + G + \int_{\mathcal{I}_0} d_0 \, \mathrm{d}\lambda_0;$$

- *and the government balances its budget:*

$$G + \sum_{t=0}^{t^C-1} \int_{\mathcal{I}_t} k(B_t - b_{t+1}) \, \mathrm{d}\lambda_t + \int_{\mathcal{I}_0} k(\tilde{d} - d) \, \mathrm{d}\lambda_0$$
$$= \sum_{t=0}^{T-1} \int_{\mathcal{I}_t} (c_t \tau_c + a_t \tau_a) \, \mathrm{d}\lambda_t + \sum_{t=0}^{T-1} \int_{\mathcal{I}_t} (n_t w(h_t) x(s_t) \tau_n(\cdot)) \, \mathrm{d}\lambda_t + \sum_{t=t^C}^{t^C+m-1} \int_{\mathcal{I}_t} \pi_t \, \mathrm{d}\lambda_t.$$

4 Parameterization

I now proceed to discuss the parameterization of the model. The parameter space consists of three parts: Some parameters are estimated outside of the model. These are described in subsection 4.1. Further, some parameters are set directly (either because they have obvious counterparts in reality or because they are readily available in existing literature), and some are set to match moments of our model to their counterparts in the data. These are all described in subsection 4.2.

Table 1: Intergenerational transmission of ability ($\Gamma_\alpha(\alpha, \breve{\alpha})$)

	Quintile	Value	Child				
			Bottom	2	3	4	Top
	1	$\rho + \gamma(-1.40)$	0.37	0.25	0.17	0.14	0.07
	2	$\rho + \gamma(-0.53)$	0.26	0.24	0.21	0.16	0.14
Parent	3	ρ	0.16	0.21	0.22	0.23	0.18
	4	$\rho + \gamma(0.53)$	0.12	0.18	0.22	0.23	0.25
	5	$\rho + \gamma(1.40)$	0.07	0.12	0.18	0.25	0.38

4.1 Estimation

A number of important drivers are estimated outside of the model using micro data. These are, in particular, the transmission of ability, the idiosyncratic earnings uncertainty, and the dependency of grants on ability and permanent parental income.

Ability transmission The intergenerational transmission of ability is determined by $\Gamma_\alpha(\hat{\alpha}, \alpha)$. To calibrate this part of the model, we do not choose a functional form. Instead, we directly employ data from the NLSY79 (National Longitudinal Study of Youth '79) and the Children of the NLSY79 datasets, which contain scores on tests taken by mothers and their children. As part of the former study, women aged about 16 to 23 were asked to take an AFQT (Armed Forces Qualification Test) in 1981. They have been tracked since, and their children were also tested using a variety of metrics. This allows us to establish a connection between the ability of mothers and their children. The test we use to assess the ability of children is the PIAT Math test, who were 7 years old on average when taking the test. We then sort both mothers and children into quintiles on their respective scores, and determine a transition matrix. That matrix is shown in table 1. Ability is quite persistent at the tails, but not so much in the middle of the distribution. Our procedure up to this point essentially follows Abbott et al. (2013).

Because the AFQT score is constructed to generate percentiles, we assume a normal distribution of ability. Each state is assigned the expected value of an observation in the corresponding quintile of a normal distribution. These values, for a standard normal distribution, are reported in table 1. Denoting the discretized standard normal distribution (whose values are found in table 1) by $\tilde{\alpha}$, and its lowest entry by $\underline{\tilde{\alpha}}$, the distribution of α is formed as follows:

$$\alpha = \tilde{\alpha}\gamma + \rho.$$

We still need to calibrate the average (ρ) and standard deviation (γ) of our (discretized) normal distribution, which is described below.

Earnings uncertainty As is typical in calibrated model of the macro-economy, the model setup of this paper restricts idiosyncratic earnings uncertainty to be of a first-order Markov form, so that only one state is requires to track the idiosyncratic component in earnings. This process is ideally calibrated based on an empirical study of hourly wages that allows for significant heterogeneity in the systemic component of wage profiles. As Guvenen, Kuruscu, and Ozkan (2014) note, the closest such study is by Haider (2001). Two complications now arise: that paper uses an ARMA model for log wages, which would take an additional state variable to track the moving average of wages, and its estimates are based on yearly data while the calibration period in this paper is four years. These issues are resolved as follows: the ARMA process estimated by Haider (2001) is simulated, after which every four simulated periods are summed to one, and an AR(1) process is estimated on the resulting series using maximum likelihood. Taking this approach, we use both the best possible measurement of the idiosyncratic component of wages, and the best possible approximation of that process in the context of our model. The estimates of the autoregressive coefficient and error term variance are then used to create a discrete and symetric first-order Markov process with two states, which has the same persistence and unconditional variance as the estimated model. The final $\Gamma_s(s_t, s_{t+1})$ and $x(s_t)$ are shown in table 2. $\bar{s}$ is set to be the lower of the two states.

Table 2: Idiosyncratic earnings process

		To	
		1	2
From	1	0.72	0.28
	2	0.28	0.72
	Value	0.72	1.28

Student grants and college subsidies Appendix A.2 provides an overview of the landscape of US education policy around 2003. In part, student aid consists of student loans, which are modeled explicitly in this paper and parameterized below. For the remained, a plethora of student grants from federal, state, and local governments, as well as tuition discounts based on family income and merit create a wedge between individual investment in college (d), and actual instructional expenditure ($\tilde{d}$) that one would expect to be effective in the creation of human capital. I now lay out the mapping between these two variables and then parameterize it.

First, let us call the sticker price (a college's headline figure for tuition and fees) s. A quick look at the data of Chetty et al. (2017) shows that the relation between instructional expenditure and sticker prices over colleges is close to linear, so that we write (recycling some notation, and with the superscript D referring to data variables):

$$\tilde{d}^D = \alpha_0 + \alpha_1 s^D. \tag{1}$$

Next, let us relate total aid g^D (from colleges and all levels of government) to sticker prices to capture general subsidies, and also to family income and human capital. These latter two capture need- and merit-based aid. Indeed we will see that a linear relationship does a good job at capturing this relationship:

$$g^D = \beta_0 + \beta_1 q^D + \beta_2 s^D + \beta_3 h^D. \tag{2}$$

Because sticker prices are paid either through private expenditures or from aid (which we have defined broadly), we have $s^D = g^D + d^D$, and instructional expenditures are determined as follows:

$$\tilde{d}^D(d^D, q^D, h^D) = \alpha_0 + \frac{\alpha_1}{(1-\beta_2)}[\beta_0 + \beta_1 q^D + d^D + \beta_3 h^D]. \tag{3}$$

Assuming for a moment that we can indeed measure these variables in the data, we still need to connect the data variables to those in the model. Here, there are two issues at play. First, the numeraire in the model is different from the numeraire in the data, which influences all terms that are not linear in the numeraire. Second, the unit of measurement for human capital will be different. We resolve this by rewriting equation 3 as follows:

$$\frac{\tilde{d}^D}{\bar{y}^D} = \frac{\alpha_0}{\bar{y}^D} + \frac{\alpha_1\beta_0}{(1-\beta_2)\bar{y}^D} + \frac{\alpha_1\beta_1}{(1-\beta_2)}\frac{q^D}{\bar{y}^D} + \frac{\alpha_1}{(1-\beta_2)}\frac{d^D}{\bar{y}^D} + \frac{\alpha_1\beta_3}{(1-\beta_2)\bar{y}^D}\bar{h}^D + \frac{\alpha_1\beta_3\sigma_h^D}{(1-\beta_2)\bar{y}^D}\frac{h^D - \bar{h}^D}{\sigma_h^D}. \tag{4}$$

Here, $\bar{y}^D$ are average earnings as measured in the data. $\bar{h}^D$ and σ_h^D are also assumed measurable in the data, and represent the mean and standard deviation of h^D. Now, note that this is an equation relating normalized instructional expenditure $\frac{\tilde{d}^D}{\bar{y}^D}$ to normalized parental income $\frac{q^D}{\bar{y}^D}$, normalized personal education expenditure $\frac{d^D}{\bar{y}^D}$, and normalized human capital $\tilde{h}^D = \frac{h^D - \bar{h}^D}{\sigma_h^D}$. All of these terms have clear model counterparts, while the coefficients are measurable in the data.

Rewriting for the model counterpart of equation 4, we get (recycling some more notation, and with the superscript M referring to data variables):

$$\tilde{d}^M = a_0 + a_1 q^M + a_2 d^M + a_3 \tilde{h}^M. \tag{5}$$

Here, $a_0 = \alpha_0 \frac{\bar{y}^M}{\bar{y}^D} + \frac{\alpha_1\beta_0}{(1-\beta_2)}\frac{\bar{y}^M}{\bar{y}^D} + \frac{\alpha_1\beta_3}{(1-\beta_2)}\frac{\bar{y}^M}{\bar{y}^D}\bar{h}^D$, $a_1 = \frac{\alpha_1\beta_1}{(1-\beta_2)}$, $a_2 = \frac{\alpha_1}{(1-\beta_2)}$, and $a_3 = \frac{\alpha_1\beta_3\sigma_h^D}{(1-\beta_2)}\frac{\bar{y}^M}{\bar{y}^D}$.

Next, we turn to measurement. Data from Chetty et al. (2017) on 4-year colleges for the year 2000 (who are in turn based on the IPEDS dataset from the NCES) allow the estimation of equation 1. The National Postsecondary Student Aid Study (NPSAS) by the NCES for the year 1995-1996 links surveys of student finances to characteristics of the colleges they are enrolled in. In this dataset we find total aid received from all sources (except Stafford and PLUS loans), tuition and fees (before any aid), parental income, as well as SAT scores (combined scores) which function as a proxy for human capital. I use these data to estimate equation 2. Zeros were treated as missing values. The regression is weighted by the NCES's full sample weights.

Table 3 contains Ordinary Least Squares estimates for the two equations, as well as a measure of the explanatory power of the linear model and the number of observations used. Finally, the parameters of equation 5, that will directly be fed into the model, are displayed as well. From the NPSAS we have that $\sigma_h^D = 226.1$ and $\bar{h}^D = 930.0$ when assuming a normal distribution on the SAT score data (calculated from percentile data), which is also the assumption in the model. $\bar{y}^D$ is 31,141 in 1995 USD according to the OECD.

Table 3: Regression results

	(1) $\tilde{d}^D$	(2) g^D			(4) $\tilde{d}^M$
Constant	2984.29	131.48	Constant:	a_0/y^M	0.13
	(257.40)	(370.00)			
q^D		-0.01	q^M :	a_1	-0.01
		(0.02)			
s^D	0.27	0.52	d^M :	a_2	0.57
	(0.03)	(0.03)			
h^D		1.95	$\tilde{h}^M$:	a_3/y^M	0.01
		(0.56)			
R^2	0.14	0.42			
Observations	1255	~6900			

4.2 Moment Matching

The calibration targets the year 2003. The below describes the moments used, together with the parameters that they are informative of. This subsection ends with an overview.

Life-cycle The model period is set to four years. Model ages are set as close as possible to their counterparts in reality: working life starts at age 18 ($t = 0$), retirement at 66 ($T = 12$). Child birth occurs at age 28, which is the average age of mothers at child birth[3], so that children start their working life when the parent is aged 46. Inter-vivos transfers are made

[3]Calculated from 2010 data provided by the Center for Disease Control and Prevention (CDC).

during the period before that ($t^I = 6$).[4]

Production We use values for discounting (β, yearly value 0.987) and depreciation (δ_a, yearly value 0.012) that are standard in the literature. We adjust these values for our model period. The international interest rate ($\tilde{r}$) is set such that the post-depreciation yearly rate r is 1%. θ is set equal to the capital share of total factor income in the data (0.33).[5]

Preferences ν and σ are set to match average hours worked and the elasticity of intertemporal substitution. The former is taken to be 39%, based on a daily time endowment of 16 hours and a reported weekly 43.35 hours of total market work in 1985 Aguiar and Hurst (2007). For the latter we rely on a meta-study by Havranek (2013), who finds that the literature's best estimate for this elasticity is 0.3-0.4 after correcting for publication bias. We use the midpoint of that range.

Inter-vivos transfers Abbott et al. (2013) do extensive empirical work on inter-vivos transfers using data from the NLSY97.[6] They estimate average total inter-vivos transfers between age 16 and 22 to be \$30,566 in 2000 dollars (0.84% of GDP), and we set ω to match this figure with our one-off inter-vivos transfer.

Human capital The ability distribution parameters ρ and γ are strongly related to the initial distribution of human capital, h_0, which is assumed to be normally distributed as well, and perfectly correlated with ability. The assumption on the correlation with ability follows Guvenen, Kuruscu, and Ozkan (2014), who point out that Huggett, Ventura, and Yaron (2011) find ability and initial human capital to be strongly correlated at age 20 when estimating them jointly (a correlation of 0.792). This reduces the initial state space by one variable and requires less parameters to be calibrated, thereby reducing computational complexity. The distribution of initial human capital is then formed as follows:

$$h_0 = (\tilde{\alpha} - \underline{\tilde{\alpha}}) + \psi.$$

Thus, the lowest level of initial human capital in the economy is normalized to ψ, and the level of ψ controls the standard deviation (and thereby the mean) of this normally distributed variable by shifting the distribution up and down. Note that, to implement equation 5, we can use that in this setup $\bar{h} = \psi - \underline{\tilde{\alpha}}$ and $\sigma_h = 1$.

[4] In the problem formulation above, interest is accrued and interest income taxes are paid over the inter-vivos transfer, which we assume to be set aside in the last period and transferred to the child at the start of the next period.

[5] Data are available from the OECD for 2003.

[6] The NLSY97 surveys a nationally representative sample of individuals in much the same manner as the NLSY79, starting in 1997. Participants were aged 12 to 16 when they first participated.

γ, ρ, and the parameters in our human capital production functions (β_0, β_1, β_2, and δ_h), are used to capture features of age-earnings profiles, aggregate inequality, college enrollment, and overall spending on tertiary education.

Using data from Chetty et al. (2017) which is on the relevant cohorts, I find that 75% of the population go to college. For those where data on the type of degree are available, 62% of college-goers enter a four year college. Four year colleges are the relevant empirical counterpart for the model, so that extrapolating to the entire population I target a college-going rate of 46%.

Huggett, Ventura, and Yaron (2011) do empirical work to establish the distribution of patterns of life-cycle earnings, taking into account time fixed-effects. The sample consists of men who are attached to the labor force. We calculate the following moments from these data.

1. Average earnings at age 60 over average earnings at age 24. (1.32)
2. Average earnings at age 44 over average earnings at age 24. (1.57)
3. Average earnings at age 36 over average earnings at age 24. (1.48)
4. The variance of log earnings at age 44 versus at age 24. (1.32)

To tie down spending on education, the calibration targets the share of GDP spent on tertiary education from private sources. Because private spending in our model is very narrowly defined as direct spending by households, we take the NIPA account on private spending on higher education for 2003 as the counterpart in the data, which is 0.83% of GDP.

Finally, the moment used to tie down ψ is overall gross earnings inequality. The OECD provides a yearly series of 90th over 10th percentile earnings, based on full-time employed men. For the United States, that figure was 4.71 in 2003. The underlying data are from the Current Population Survey of the Bureau of Labor Statistics, which counts persons who work 35 hours or more per week as full-time employed. The equivalent time use in our model is 0.3125.

Student loans I follow the structure of Abbott, Gallipoli, Meghir, and Violante (2013) to model and calibrate the student loan system around 2003. The parameters y^*, y^{**}, $\underline{b}^s$, $\underline{b}^u$, r^s, r^u, and $\underline{a}^p$ are informed by the following moments:

1. The fraction of students with subsidized Stafford loans was 37.3% in 2000 (Abbott et al., 2013).
2. The fraction of students with unsubsidized Stafford loans was 21.2% in 2000 (Abbott et al., 2013).

3. The subsidized loan limit which was 43% of GDP per capita.[7]
4. The unsubsidized loan limit which was 45% of GDP per capita.[8]
5. Subsidized Stafford aid was 0.18% of GDP (College Board, 2013).
6. Unsubsidized Stafford aid was 0.15% of GDP (College Board, 2013).
7. Private sector loans were 0.06% of GDP (College Board, 2013).

The repayment period length m is set to 20 years in either case. While the repayment period has typically been 10 years, this is easily extended.

Tax policies Guvenen, Kuruscu, and Ozkan (2014) zoom in on labor income taxation, and find that US policy for 2003 is well matched by the functional form below, where $\bar{y}$ are the average United States earnings (and the same parameter that is used in the implementation of equation 5).[9]

$$\tau_n(\cdot, \bar{y}) = 1.2088 - 0.00942 \left(\frac{n_t w(h) x(s_t)}{\bar{y}} \right) - 0.94261 \left(\frac{n_t w(h) x(s_t)}{\bar{y}} \right)^{-0.10259}.$$

We take the consumption and capital income tax rates from McDaniel (2007): $\tau_c = 0.075$, and $\tau_a = 0.232$ for 2003.

Overview Table 4 provides an overview of parameters set outside of the model and their values. Further, parameter σ can be solved for analytically given our chosen target of the elasticity of intertemporal substitution (EIS): $\sigma = 1 + \frac{1}{(EIS * \nu)} - \frac{1}{\nu}$. Parameter $\bar{y}$ is calibrated internally (for each period to which the model is applied): it must match the model simulated average earnings up to some level of precision. Table 5 lists parameters that were set to match moments: it displays the final parameter values, together with the moments that inform them and their empirical value. Percentages refer to either GDP or GDP per capita.

[7] According to FinAid (2016), the aggregate subsidized loan limit in 2003 was $17,125, which was 43.16% of GDP per capita at the time.

[8] According to FinAid (2016), the aggregate unsubsidized loan limit in 2003 was $18,000, which was 45.37% of GDP per capita at the time.

[9] We simply apply this formula to periodic model incomes, normalized by average wage. This would be equivalent if incomes were indeed constant during the model period. For the purposes of this paper, we consider it a good enough approximation.

Table 4: Parameters set outside of the model

	Value	Moment
β	0.949	Discount rate
δ_a	0.047	Asset depreciation rate
$\tilde{r}$	0.130	Pre-depreciation real interest rate
θ	0.314	Capital share of income

Table 5: Parameters

Parameter	Value	Moment	Model	Data
σ		Elasticity of intertemporal substitution	-	0.35
ν		Average labor supply		0.39
ω		Average inter-vivos transfer		0.84
γ		Variance of log earnings: age 44 vs. age 24		1.32
ρ		Average earnings: age 44 vs. age 24		1.57
β_0		Share of males with some college		0.53
β_1		Average earnings: age 36 vs. age 24		1.48
β_2		Total private spending on tertiary education		0.83%
δ_h		Average earnings: age 60 vs. age 24		1.32
ψ		Gross earnings inequality (P90 vs. P10)		4.71
y^*		Fraction of students with subsidized Stafford loans		7.3%
y^{**}		Fraction of students with unsubsidized Stafford loans		21.2%
$\underline{b}^s$		Subsidized loan limit		43.16%
$\underline{b}^u$		Unsubsidized loan limit		45.37%
r^s		Subsidized Stafford aid		0.18%
r^u		Unsubsidized Stafford aid		0.15%
$\underline{a}^p$		Private sector loans		0.08%
$\bar{y}$	-	Internally calibrated	-	-

5 Model Validation

6 Policy Experiments

7 Conclusion

A Appendix

A.1 Computation

The computational procedure that produces the results in this paper proceeds as follows.

1. Guess parameter values.
2. Guess the initial value function at independence V (because we interpolate between grid points, guesses at grid points are sufficient). (Prices can be solved analytically in our framework.)
3. Solve the individual's problem. I elaborate on this below. This results in an updated function V.
4. Simulate households. This is done by randomly assigning initial states, and then simulating a household for many generations. Using a large number of households and a large number of generations per household, we arrive at a stationary distribution of the economy.
5. Update average income and V^C (i.e. repeat from step 2) until convergence. Howard improvement steps speed up this procedure significantly.
6. Update parameter values (i.e. repeat from step 1) until moments are matched.

The individual problem is in principle standard, although involved, and can be solved using standard methods. In this paper, the optimization problem is solved by backward induction. Value functions are (multi-dimensionally) linearly interpolated for continuous states. Optimization over choice variables is either done using first-order conditions, or by a 'double grid' method. In the 'double grid' optimization method, an optimum over a rough grid is found first, after which a new grid is formed around the maximum value (the interpolating function is kept constant). This method works well for well-behaved objective functions, and achieves reasonable precision on a small number of grid points. The choice of methods is largely driven by their robustness, which is important given the large number of choice and state variables.

The individual problem is first solved under the assumption that in the optimum we always have $l_t + e_t < 1$. In that case, analytical expressions for c_t, d_0, and a_{t+1} can be derived (in all but the last stage of the life-cycle) using first-order conditions, the budget constraint, Inada conditions, as well as the remaining bounds. Under this assumption, optmization can then proceed by performing a double-grid search over l_t and where relevant e_t (and a root-finding routine in combination with the relevant first-order condition to find v'). It is then verified ex-post that choices for which $l_t + e_t < 1$ are indeed optimal for all but a very small mass of agents in the economy.

The code for this procedure was written in Fortran90 and parallelized using OpenMPI.

A.2 Student Aid in the United States

This appendix describes US tertiary education policies around 2003. Fuller (2014) provides a more detailed description of the history of US education policies.

Table 6: Education policies in 2003

Student Aid (as % of total, or in millions of 2012 USD)	2003
Grants (non-institutional)	
Pell	54%
Other Federal (mostly military)	19%
State	27%
Total	$ 27,461
% of GDP	0.19%
Institutional Grants	$ 22,470
% of GDP	0.16%
Public Sector Loans	
Stafford, subsidized	44%
Stafford, unsubsidized	38%
PLUS	11%
Other Federal	4%
State and Institution Sponsored Loans	3%
Total	$ 56,280
% of GDP	0.40%
Private Sector Loans	$ 8,900
% of GDP	0.06%

Sources: author calculations; data from College Board (2013); CPI from the St. Louis FRED database; GDP from the World Bank's WDI.

Table 6 provides an overview of the student aid landscape in 2003. Government intervention generally consists of grants and loans. The largest uniform grant program is the Pell grant program, which provides grants to college students depending on financial need. Other programs are sizable but spread thin, with most of the money coming from states or military (including veteran) related programs. Institutional grants are of a similar order of magnitude as non-institutional grants. This uncovers a serious issue with using headline costs of college to calibrate models with an extensive margin: institutional grants are essentially discounts to attending a college, and given their size the headline costs can hardly be taken to be the price of college. In addition, colleges often discount prices based on both financial need and merit. To account well for that complicated landscape, this paper relies on linked micro survey data on students and the colleges they go to.

Public sector loans, the other major policy instrument, largely consist of Stafford loans. These loans, in their subsidized version, provide student loans to students from lower income families at below market rates. Interest accrued during college is forgiven. Unsubsidized loans have higher interest rates and no accrual forgiveness, but are easily available to students regardless of family income. Subsidized and unsubsidized loans are subject to a joint limit, in addition to a separate limit on subsidized loans. Stafford loans are explicitly modeled and calibrated in this paper.

The other major loan programs are PLUS loans and Perkins loans. The availability of PLUS loans in practice strongly depends on parental credit scores, and are essentially a way for parents to transfer privately borrowed funds to children. This mechanism is separately present in the model through inter-vivos transfers, so that PLUS loans are not modeled explicitly. The Perkins loan program is small in size, and also not modeled.

Private student loan markets were small in 2003. Why this is so, not only in the United States but globally, is a topic of research in its own right. Here I put forward the following narrative: in the face of regular consumer bankruptcy regulation, private student loan markets are unlikely to develop at all (cf. Lochner and Monge-Naranjo (2011b)). Public student loans, presumably for the same reason, have historically been exempted from discharge in regular bankruptcy proceedings. Importantly, this exemption was extended to any non-profit entity in 1985, allowing many financial institutions to structure their loans such that they were immune to discharge (Consumer Financial Protection Bureau, 2012) and the private student loan market to develop.

Despite the discharge exemption, private student loans are not widely available: for example, 90% of undergraduate and 75% of graduate private student loans in 2012 were co-signed (MeasureOne, 2013). Without a co-signer, students typically lack a credit history that would allow them to take out a loan at competitive interest rates, but those that do take out these loans tend to get them at rates that are attractive compared to unsubsidized Stafford loans (Institute for Higher Education Policy, 2003). Consequently, this paper models them as available to a subset of those that do not qualify for subsidized Stafford loans based on family income. The model structure thereby coincides with that in Abbott, Gallipoli, Meghir, and Violante (2013), as do most modeling choices regarding student loans.

Default on student loans or otherwise is not modeled. It is precisely the exemption from discharge that makes this a less relevant issue for the purposes of this paper: students may default, but then still have to repay their student loans. In fact, the College Board (2013) documents that over 90% of federal student loan dollars that enter default are eventually recovered.

References

Abbott, B., G. Gallipoli, C. Meghir, and G. L. Violante (2013). Education policy and intergenerational transfers in equilibrium.

Aguiar, M. and E. Hurst (2007). Measuring trends in leisure: The allocation of time over five decades. *The Quarterly Journal of Economics 122*(3), 969–1006.

Angrist, J. D. (2014). The perils of peer effects. *Labour Economics 30*, 98–108.

Ben-Porath, Y. (1967). The production of human capital and the life cycle of earnings. *The Journal of Political Economy*, 352–365.

Bovenberg, A. L. and B. Jacobs (2005). Redistribution and education subsidies are siamese twins. *Journal of Public Economics 89*(11).

Chetty, R., J. Friedman, E. Saez, N. Turner, and D. Yagan (2017). Mobility report cards: The role of colleges in intergenerational mobility.

Chetty, R., N. Hendren, P. Kline, and E. Saez (2014). Where is the land of opportunity? the geography of intergenerational mobility in the united states. *The Quarterly Journal of Economics 129*(4), 1553–1623.

College Board (2013). Trends in student aid 2013. Technical report.

Consumer Financial Protection Bureau (2012). Private student loans. Technical report.

Corak, M. (2013). Income inequality, equality of opportunity, and intergenerational mobility. *The Journal of Economic Perspectives 27*(3), 79–102.

Epple, D., R. Romano, S. Sarpça, and H. Sieg (2013). The us market for higher education: A general equilibrium analysis of state and private colleges and public funding policies.

Epple, D., R. Romano, and H. Sieg (2006). Admission, tuition, and financial aid policies in the market for higher education. *Econometrica 74*(4), 885–928.

FinAid (2016, November). Historical loan limits.

Findeisen, S. and D. Sachs (2016). Education and optimal dynamic taxation: The role of income-contingent student loans. *Journal of Public Economics 138*, 1–21.

Fuller, M. B. (2014). A history of financial aid to students. *Journal of Student Financial Aid 44*(1), 4.

Guvenen, F., B. Kuruscu, and S. Ozkan (2014). Taxation of human capital and wage inequality: A cross-country analysis. *The Review of Economic Studies 81*(2), 818–850.

Haider, S. J. (2001). Earnings instability and earnings inequality of males in the united states: 1967–1991. *Journal of labor Economics 19*(4), 799–836.

Hanushek, E. A., C. K. Y. Leung, and K. Yilmaz (2003). Redistribution through education and other transfer mechanisms. *Journal of Monetary Economics 50*(8), 1719–1750.

Havranek, T. (2013). Publication Bias in Measuring Intertemporal Substitution. *Working Papers IES 2013*(15).

Holter, H. A. (2015). Accounting for cross-country differences in intergenerational earnings persistence: The impact of taxation and public education expenditure. *Quantitative Economics 6*(2), 385–428.

Hoxby, C. M. (2015). Computing the value-added of american postsecondary institutions.

Hoxby, C. M. (2016). The dramatic economics of the u.s. market for higher education. *NBER Reporter 3*.

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

Institute for Higher Education Policy (2003). Private loans and choice in financing higher education. Technical report.

Kapicka, M. and J. Neira (2015). Optimal taxation with risky human capital. *CERGE-EI Working Paper Series* (553).

Krueger, D. and A. Ludwig (2016). On the optimal provision of social insurance: Progressive taxation versus education subsidies in general equilibrium. *Journal of Monetary Economics 77*, 72–98.

Landersø, R. and J. J. Heckman (2017). The scandinavian fantasy: The sources of intergenerational mobility in denmark and the us. *The Scandinavian Journal of Economics 119*(1), 178–230.

Lee, S. Y. T. and A. Seshadri (2014). On the intergenerational transmission of economic status.

Lochner, L. and A. Monge-Naranjo (2011a). Credit constraints in education.

Lochner, L. J. and A. Monge-Naranjo (2011b). The Nature of Credit Constraints and Human Capital. *The American Economic Review 101*(6), 2487–2529.

McDaniel, C. (2007). Average tax rates on consumption, investment, labor and capital in the OECD 1950-2003.

MeasureOne (2013). Private student loan report 2013. Technical report.

Mendes Tavares, M. and C. Urrutia (2016). An empirical investigation on the great gatsbys curve.

Stantcheva, S. (2015). Optimal taxation and human capital policies over the life cycle.