Higher Education Policies and Intergenerational Mobility

Vinzenz Ziesemer*

Erasmus University Rotterdam

May 2021

Abstract

Low levels of intergenerational mobility are often taken as a sign of inefficient educational outcomes. This paper finds that efficient higher education policies may actually decrease intergenerational mobility. Policies increase access to college. But they also increase the importance of human capital in earnings versus other factors such as luck, making earnings more persistent over generations. To establish which effect dominates, the paper develops a model of linked generations of heterogeneous agents. The model includes a tractable market for heterogeneous colleges. Parameterizing the model, the paper finds that common and welfare-improving higher education policies may actually decrease intergenerational mobility.

*Email: vinzenz.ziesemer@eui.eu. For comments and discussions, I am grateful to Árpád Ábrahám, Axelle Ferriere, Piero Gottardi, Bas Jacobs, Dirk Krueger, Victor Rios-Rull, Dominik Sachs, an anonymous referee, as well as numerous participants of seminars and conferences.

1 Introduction

1.1 Overview

Intergenerational mobility (IM) is typically measured as the inverse of earnings persistence. The less related children's earnings are to those of their parents, the more mobile society is. Cross-country studies show significant differences in IM, amongst others showing more mobility in Scandinavian countries than in the United States (Corak, 2013).

Education is the component of earnings most obviously related to family background. Therefore, authors and commentators interested in IM tend to focus on the role of education. Indeed, recent work by Chetty, Friedman, Saez, Turner, and Yagan (2017) has delved into the role of higher education and suggests that once we understand college outcomes, we will largely understand IM. These authors also find that low-income students are unlikely to enter selective colleges. However, when they do, they stand a good chance of moving up the income distribution. This result suggests that higher education policies might be effective in stimulating IM. Figure 1 and Appendix A.1 contain the relevant empirical facts. For these reasons, lacking IM is usually taken as a sign of inefficient educational outcomes. Education policies are, in part, designed so that children can overcome their circumstances at birth. If they do not achieve that goal, the thinking goes, something must be off.

But do education policies really increase IM? There is essentially no causal reduced-form evidence regarding this question. The long horizons over which such policies take effect, a lack of historical data, as well as many confounding trends and unobservables stand in the way. Instead, this paper explores the nexus between higher education and IM using a parameterized model. In the model this paper considers, efficient and welfare-improving education policies can reduce standard measures of IM. Thus, the paper concludes that measures of IM should be viewed critically and interpreted with care.

The first contribution of this paper is to establish a theoretical ambiguity. Taking wages as a combination of human capital and luck, as is standard in macroeconomics, yields an insightful decomposition of the intergenerational elasticity of earnings (or IGE, a standard measure of IM). The decomposition makes apparent that there are two key ingredients to intergenerational persistence. The first one is the correlation between the human capital of parents and that of their children. One would expect education policies to reduce this correlation. The second one is the relative importance of education in earnings (versus luck). If education policies increase the variance of educational outcomes, then this will increase the intergenerational persistence of earnings. Which effect dominates is an empirical matter.

Second, the paper contains a modeling contribution. It sets up a model of linked generations

Figure 1: Access, Success, and Mobility in the Data

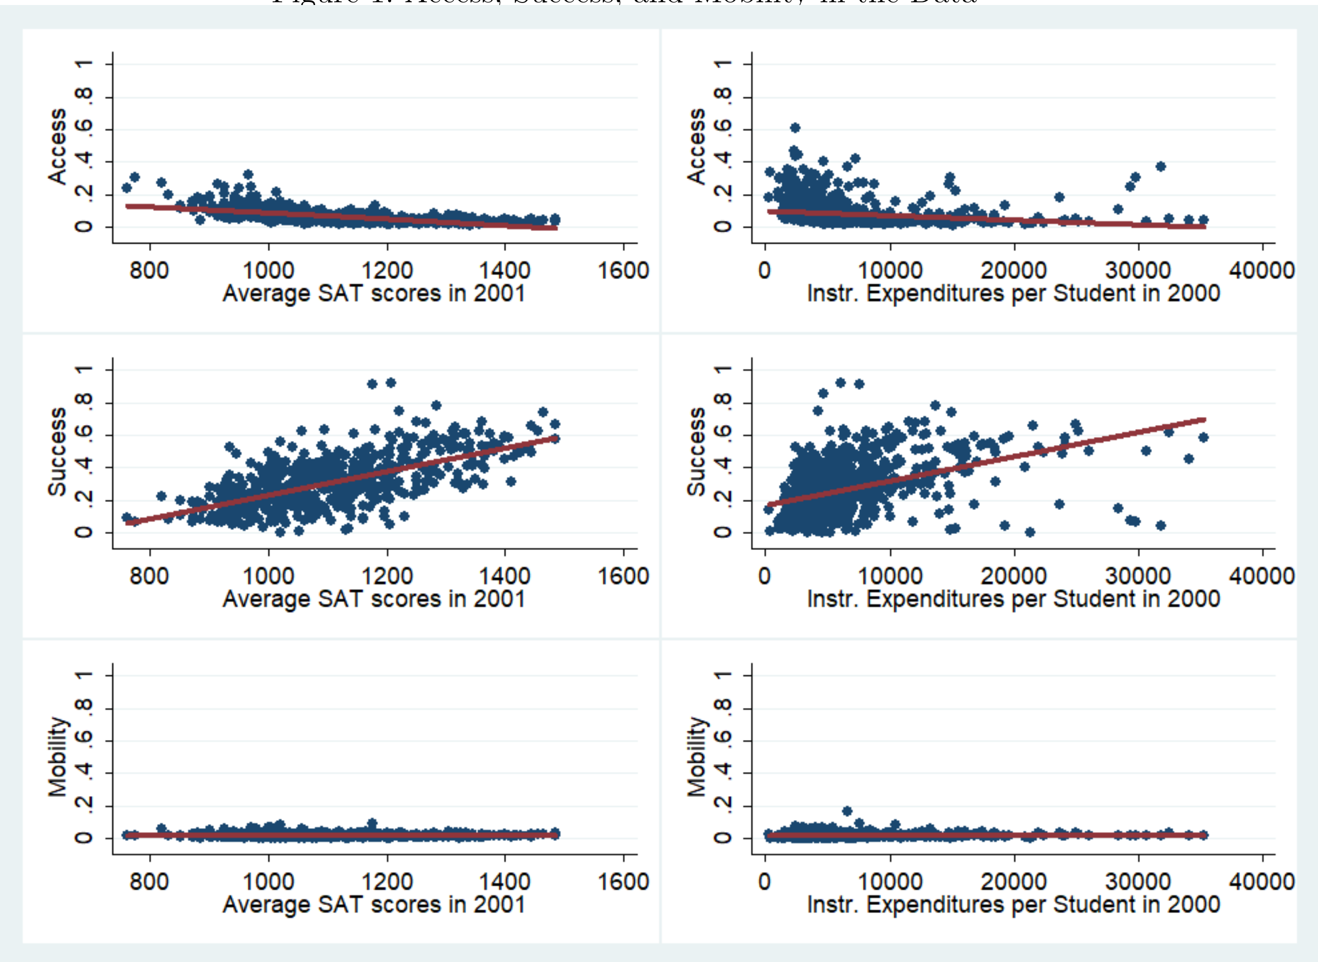

Each point represents a college. 'Access' is defined as the share of low-income students (defined as a family income in the bottom 20%) going to the college. 'Success' is defined as the likelihood that a student from a low-income family reaches the top 20% of the income distribution after attending the college. 'Mobility' is the product of the two, i.e., the share of the college's students that go from the bottom to the top of the distribution. The figure contrasts the three measures with two indicators of college selectiveness. For the graphs on the left, this is the average SAT score of an entering student. For the graphs on the right, it is the college's yearly average instructional expenditures per student (in US Dollars). See Appendix A.1 for further details.

of heterogeneous agents featuring heterogeneous colleges. The innovation is to model colleges as operating in a perfectly competitive market, which keeps the supply side of college education tractable while allowing for sufficient heterogeneity to match relevant facts. Moreover, in order to study how higher education policies affect IM, the model includes parental links, human capital formation, and earnings over the life cycle.

Third, the paper uses the model to assess the impact of higher education policies on IM. In order to do so, it must capture several empirical ingredients to human capital formation. Careful parameterization proceeds as follows. (i) Quasi-experimental results establish the effectiveness of colleges in delivering future earnings. College effectiveness, in turn, is a critical determinant of the effectiveness of policies. (ii) The availability of income- and need-dependent student grants is estimated from survey data, while parental transfers and the US student loan system are modeled explicitly. Jointly, these tie down the sources of college financing. (iii) Human capital formation after college is parameterized by the path of earnings over the adult life cycle, which allows for the measurement of earnings at specific ages in the model. (iv) The intergenerational persistence of learning ability is quantified using test scores of matched pairs of parents and children. That completes the links between parents and children. In this way, the model consciously abstracts from childhood formation, meaning that policy changes at the college level do not change parenting choices with regards to pre-college education. While this limitation comes at some cost, it greatly reduces complexity and allows for greater focus on the role of the key channels.

The parameterization must also make sense of the non-human capital components in earnings (e.g., unemployment, bad health, disability, and any other form of good or bad luck), which are the other key ingredient to the analysis. In order to capture these components systematically, the paper follows the literature on incomplete markets by using estimated earnings processes. The paper shows that its human capital-based approach explains the facts on IM well. The model also matches several other moments related to college enrollment and labor markets, none of which are targeted in its parameterization.

Counterfactuals establish the main result: compared to a situation with no policies at all, merely removing borrowing constraints reduces IM. The ambiguity is at work here: policies that remove borrowing constraints increase mobility in human capital, but also increase inequality in human capital. The latter effect increases the role of persistent human capital in earnings, causing a reduction in IM. A simple transition analysis shows that the reaction of IM is monotone in time. Both channels work at the same time, so that IM begins to decrease as college-going children are relieved of borrowing constraints. The reaction is also relatively quick: almost all of the change in IM has occurred by the time one generation of parents and children have lived under the same set of policies.

Current education policies, such as student grants and loans, are found to reduce IM even further. The same channels are at work again, but now mobility in human capital itself is reduced as well. This is explained by the role of parental transfers: As before, education policies cause increased inequality in human capital, which translates into higher earnings inequality. But under current policies, many students remain borrowing constrained. Since high-income parents are now even better off, they spend even more on their children's education. As a result, the mobility of human capital falls.

These results have several practical implications. They show a trade-off between individual welfare and IM. Policymakers should, therefore, interpret measures of IM with care and refrain from targeting them, at least under standard welfarist considerations. One should also be careful when comparing countries, regions, or periods by IM. Competing channels are at work, so that one cannot infer the efficiency of policies. Landersø and Heckman (2017) provide direct evidence for the importance of this point. They find that IM is larger in Denmark than in the United States when measured by earnings, but the same is not true when measuring mobility by educational attainment. The results in this paper show that this is not just a matter of the measures used: as the counterfactual results show, education policies may well reduce the mobility of human capital even if the latter is precisely measured.

This paper includes two further sets of results. The first is on the role of college heterogeneity and financial constraints. The empirical literature in Appendix A.1 suggests that college heterogeneity may play an important role in IM. Indeed in the model, even when students are essentially unconstrained in their extensive margin of college choice (going or not), they may still be constrained in their intensive margin of college choice (which college to go to). Thus, accounting for college heterogeneity is crucial in understanding the welfare implications of higher education policies.

Second, the model yields a decomposition of the current persistence of earnings across generations. Roughly half of earnings persistence is determined before the start of adult life. Thus, while it is true that childhood is perhaps most important in determining IM, higher education, and adult life are well worth studying. Of the remaining persistence, about a third is due to money from parents, and two thirds to government policies. Parental resources are an essential source of college financing, and their responsiveness reduces the impact of policies. Nevertheless, higher education policies play a significant role in determining IM.

The remainder of this section (1.2) discusses related literature. Section 2 uses theory to demonstrate why higher education policies have an ambiguous effect on IM. Section 3 contains the model. Section 4 describes how the parameters of the model are either estimated or set to match moments. Section 5 explores the ability of the model to match aspects of the data that were not targeted in the parameterization of the model. It also discusses what we can learn from the model's positive implications. Section 6 contains results from counterfactual policies. Section 7 concludes and makes suggestions for future research.

1.2 Literature

There is a quantitative macroeconomic literature that connects education policies to IM. Lee and Seshadri (2014) argue that a rich life-cycle model with intergenerational links can explain well several intergenerational relationships, in particular the intergenerational elasticity of earnings. They focus more on the development of human capital during childhood, and less on college heterogeneity. Holter (2015) similarly builds a quantitative model of IM. He then investigates the extent to which differences in tax and education policies can explain cross-country differences in IM. Herrington (2015) also uses a structural model to look at the effect of taxes and education policies on inequality and IM, comparing policies in the United States to those in Norway. Restuccia and Urrutia (2004) build a model of intergenerational mobility and decompose sources of persistence, taking into account both early and college education. They find that parental investment in education accounts for one-half of the intergenerational correlation in earnings, underlining the motivation for this paper.

The current paper features a richer model of inter-generational links, and takes a more granular look at higher education policies. It is also the first to discuss the theoretically ambiguous effect of these policies. Some earlier theoretical work on IM, starting with Becker and Tomes (1979), is reviewed in Grawe and Mulligan (2002). Earlier work also discusses heterogeneous human capital formation in overlapping generations models, but not in relation to IM. Key references are Heckman, Lochner, and Taber (1998a) and Heckman, Lochner, and Taber (1998b).

A number of authors studies Positive implications of education policies on college enrollment. A relevant contribution by Lochner and Monge-Naranjo (2011b) considers the effect of student loan policies on the college entry decision of youth that is heterogeneous in ability and family income. Abbott, Gallipoli, Meghir, and Violante (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. They only consider one type of college. Empirical work on the incidence of financial constrainedness is summarized in Lochner and Monge-Naranjo (2011a), who find increased evidence for such incidence in recent years. In particular, and much in line with this paper, they emphasize the importance of college heterogeneity in understanding financial constraints. None of these papers focus on IM.

There is a large normative literature on education policies, often in combination with taxation. While these papers answer a different question, their insights guide the discussions of policy optimality in this work. Krueger and Ludwig (2016) focus is on optimal taxation with (almost) linear instruments. Their paper also includes intergenerational transfers, but only has one type of college, and considers general equilibrium effects as well as the importance of the transition between different policy regimes. Bovenberg and Jacobs (2005, 2011) find that while education subsidies themselves distribute resources to the well to do, their optimal level may still be positively related to tax rates. This is because they undo the disincentive effects of taxation on human capital formation. The same issue has been been studied in a dynamic theoretical framework by Stantcheva (2017), and in a quantitative framework by Hanushek, Leung, and Yilmaz (2003). Further, 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 (2016a). Finally, Findeisen and Sachs (2016b) study an economy with the same motivation for education subsidies as in Bovenberg and Jacobs, but with an extensive margin for college choice and under financial constraints. In that case, the government wants to efficiently target those who would optimally be students from a social standpoint, but who would not enter college in the absence of policy intervention. It can do so by need-dependent grants, essentially using parental income as a tag of financial constraints. Insights from this normative literature guide some of the discussion of policy in this paper. Readers more interested in the political economy of education reform may be referred to a series of papers by Fernandez and Rogerson (1995; 1996; 1998; 2003).

This paper considers the adult part of the life cycle in isolation. Incorporating the effects of policies on the development of children in the earlier stages of the life cycle would likely increase the effects sizes reported in this paper. This is discussed further in subsection 4.1. The literature yields several insights into human capital formation during the earlier phases of the life cycle. For example, Cunha, Heckman, Lochner, and Masterov (2006) and Cunha and Heckman (2007), and Caucutt and Lochner (2012) study the complementarity between investments during different phases of the life cycle. Generally, this literature finds that complementarity to be strong, so that investment in earlier human capital formation is more effective. Indeed, this paper finds that IM is largely determined by age 18. Following this line of investigation, recent work has tried to disentangle the determinants of the income gradient in childhood performance (e.g. Caucutt, Lochner, and Park, 2017), much as this paper does for the adult part of the life cycle. Holmlund, Lindahl, and Plug (2011) attempt to synthesize a growing literature on the effect of parents' schooling on children's schooling. Overall, the causal effect of changes to parents' schooling on children's schooling appears to be small relative to the total correlation between parents' schooling and that of their children. That is in line with this paper's assumptions, since it takes ability at the start of adult life (and its transition across generations) as given.

2 Decomposing Intergenerational Mobility

The most commonly used measure of intergenerational persistence is the intergenerational elasticity of earnings (IGE), measured as β^{IGE} in the regression equation below:

$$\log(y') = \beta_0 + \beta^{IGE} \log(y) + \epsilon \tag{1}$$

Here, y is a measure of parental earnings, and y' measures the earnings of their children. As we will see later, measurements of β^{IGE} in the literature have a wide range between 0.3 and 0.6, suggesting that a 1% increase in parental earnings may lead to 0.3% up to 0.6% higher earnings for children. In other words, earnings are persistent over generations but not perfectly: they regress to the mean.

As is well known, common estimators of the equation above (such as OLS) are unbiased estimators of

$$\beta^{IGE} = \frac{Cov(\log y', \log y)}{Var(\log y)}.$$

A large macroeconomic literature measures labor wage rates as the multiple of two components: human capital (h), and some idiosyncratic shock to income (x). The latter represents any form of luck not related to ability or education. Abstracting from labor supply for this exposition, this yields the following measure of log earnings:

$$\log y = \log h + \log x.$$

If we now assume that luck is entirely independent from human capital as well as from parental characteristics, we can write:

$$\begin{aligned}\beta^{IGE} &= \frac{Cov(\log h, \log h') + Cov(\log x, \log h')}{Var(\log h) + Var(\log x)} \\ &= \underbrace{\frac{Var(\log h)}{Var(\log h) + Var(\log x)}}_{\text{Weight of } h} \underbrace{\frac{Cov(\log h, \log h')}{Var(\log h)}}_{\text{LPC of } h' \text{ on } h} + \underbrace{\frac{Var(\log x)}{Var(\log h) + Var(\log x)}}_{\text{Weight of } x} \underbrace{\frac{Cov(\log x, \log h')}{Var(\log x)}}_{\text{LPC of } h' \text{ on } x}\end{aligned} \tag{2}$$

Now the IGE is a weighed mean of two (log-)linear projection coefficients, which I will refer to as LPCs: the projection coefficient of (the log of) children's human capital on that of their parents ($\frac{Cov(\log h, \log h')}{Var(\log h)}$), and the projection coefficient of (the log of) children's human capital on (the log of) parental luck ($\frac{Cov(\log x, \log h')}{Var(\log x)}$). The respective weights are the variance of log human capital relative to the total variance of earnings ($\frac{Var(\log h)}{Var(\log h)+Var(\log x)}$), and the variance of log income shocks relative to the total variance of earnings relative to the total variance of earnings ($\frac{Var(\log x)}{Var(\log h)+Var(\log x)}$).[1]

[1]When an economy is in steady state, we have that $Var(\log y) = \sqrt{Var(\log y)}\sqrt{Var(\log y')}$, so that the IGE measures the correlation of earnings across generations $Cor(\log y', \log y)$. Under that same assumption, the decomposition can be further rewritten as:

$$\frac{Var(\log h)}{Var(\log h) + Var(\log x)} Cor(\log h, \log h') + \frac{Var(\log x)}{Var(\log h) + Var(\log x)} \Big[Cor(\log x, \log h') \frac{\sqrt{Var(\log h)}}{\sqrt{Var(\log x)}} \Big].$$

Here, the IGE is a weighed mean of the correlation between two generations' human capital on the one hand

In other words, the IGE does measure how children's human capital relates to the two components of their parents' income: human capital and luck. But the two components may not be equally important: as we will find, parental human capital is more strongly related to children's human capital than is parental luck. As a consequence, the weights of the two projection coefficients are important: if luck plays a larger role in the overall variation of income, then the IGE shrinks.

This paper studies the effect of higher education policies on measures like the IGE. Amongst other things, higher education policies relieve financial constraints. So how does the IGE change when financial constraints are removed? Table 1 contains output from a model with and without financial constraints from Section 6. It shows the components of the expression above.

The first component (row 1 of Table 1) equals the correlation of human capital across generations in steady state. Financial constraints in education restrain the human capital of children from disadvantaged families. Human capital in these families tends to be low in the model. Taking away these financial constraints then makes these children's human capital less like that of their parents, so that the correlation of human capital across generations is lower. The second component (row 2 of Table 1) of our weighted sum goes the same way. Parental luck increases children's human capital when they are financially constrained, but less so in the unconstrained case: without constraints, children's potential outcomes do not depend on the financial situation of their parents. Thus, the second component is also smaller in the absence of financial constraints. So far, education policies would be expected to reduce the IGE.

The weights of the two components (rows 3 and 4 of Table 1), make the effect of education policies ambiguous. As I will argue later, the variance of the log income shock (row 4) is unlikely to vary much due to education policies. However, the variance of log human capital (row 3) may vary. Removing financial constraints has two effects here: first it removes variation between students of the same ability to acquire human capital but with different family circumstances. Second, it increases variation between students with different abilities, as high ability students see a larger increase in human capital.[2] As we will see, the second effect dominates, so that removing financial constraints increases the variance of log human capital.[3]

Since the first component (the LPC of h' on h) is larger than the second (the LPC of h' on

$(Cor(\log h, \log h'))$, and a measure of the influence of parental luck on children's human capital on the other $\left(\left[Cor(\log x, \log h')\frac{\sqrt{Var(\log h)}}{\sqrt{Var(\log x)}}\right]\right)$.

[2]This may be either because they are more constrained, or because the additional investment yields higher

Table 1: The effect of financial constraints on the IGE

	Constrained	Unconstrained
LPC of h' on h	0.48	0.46
LPC of h' on x	0.03	0.00
$Var(\log h)$	0.07	0.16
$Var(\log x)$	0.08	0.08
β^{IGE} (wage rates)	0.24	0.29

x), an increased variance of log human capital actually increases the IGE, since it increases the weight of the larger component. This moves the IGE in a way opposite of the first two components, which both reduce the IGE when financial constraints are removed. The overall effect of education policy on the IGE is not clear *ex ante*. Instead, it becomes a matter of measurement. As Table 1 already shows, the remainder of this paper finds that the variance effect (row 3 of Table 1) dominates for the model's equivalent of current US higher education policies. As a result, education policies actually decrease IM. The paper also demonstrates that the result does not depend on the specific measure of mobility discussed here (the IGE), but also applies to ranked correlations (a popular alternative).

3 A Model of Intergenerational Earnings Persistence

The following model of intergenerational persistence focuses on higher education. The first subsection focuses on the modeling of heterogeneous colleges, given the importance the empirical literature on IM and on financial constraints attributes to that heterogeneity. It sketches a competitive market for colleges.[4] As a result of that competitive market, students can choose how much to invest in their own education, with colleges just translating spending into investment. The setup is simple and tractable, yet allows for a rich heterogeneity of colleges.

In the second subsection, college choice is embedded into a model of the labor market. Higher education takes place in the first period after the start of adulthood. After that, agents go through a life cycle of earnings and related choices, and have children of their own. The resulting model describes earnings persistence and the role of higher education

returns for these students. The latter will turn out to be more important.

[3]Han and Mulligan (2001) already pointed out that increased heterogeneity in ability decreases IM, but did not discuss the ambiguous role of higher education policies.

[4]The terms 'college' and 'institution of higher education' are used interchangeably. Later, the model will be brought to the data in such manner that the higher education phase represents the entire higher education career.

in it. Modeling choices are highlighted as they appear. The individual's decision problem is specified in full. The subsection ends by defining the stationary equilibrium of the model.

3.1 Competitive Colleges

Students are defined by their learning ability α. When going to a college, that learning ability combines with time spent studying e and money invested in education d to form human capital $h(\alpha, d, e)$. To have money invested in education, the student must go to a college, which charges price $\tilde{d}(d, q, \alpha)$ for an investment of d.[5] In principle, the college can condition that price on the student's parental income status q (so that the price of college becomes need-based) and on the ability α of the student (which makes the price merit-based).

The student's decision-making problem is discussed in detail in the following subsection. For now, it suffices to say that the student chooses from available colleges based on the price he must pay for d since that is the only thing the college has to offer. Peer effects, whether through learning or networking, as well as the signaling value of going to a college, are not modeled explicitly. The model will be parameterized to match the actual earnings returns to investing in education so that it does not necessarily matter for the purposes of this paper whether these returns are due to actual learning or other sources such as signaling.

Private colleges are indexed by their level of educational spending per student d, which is the same for each student in the college. Colleges do not face any fixed costs. Instead, they have access to an education technology in which they simply incur the cost of educational investment for each student.[6] New colleges can freely enter any market for a d type college, free of cost. Suppose they either value profits (for-profit colleges) or their existence (not-for-profit colleges). Then we have the following:

$$\forall q \forall \alpha \quad \tilde{d}(d, q, \alpha) \leq d.$$

Any type of student receives an educational investment that is at least as large as their spending on college. If this condition was violated for any type of student, new colleges would enter and offer the same services at a lower price until some college offers $\tilde{d}(d, q, \alpha) = d$.

Can colleges exist for whom $\tilde{d}(d, q, \alpha) < d$ for some type of student? Yes, if they have other income that they choose to invest in their students' education. A clear example of

[5] Price here is meant to refer to the price for college, and not for food, lodging, and the like. In this paper, those are considered consumption items and will be treated as such when connecting theory and data.

[6] This is a simplification that does not come at much of a cost. Because costs now scale linearly in the number of students, the number of students in each college will be indeterminate. This is not important for the purposes of this paper.

such income would be endowment income. In any case, due to the result above, any college pricing schedules can always be written as follows:

$$\tilde{d}(d, q, \alpha) = d - g^I(q, \alpha), \quad \text{where} \quad g^I(q, \alpha) \geq 0.$$

Here, $g^I(q, \alpha)$ is a grant or discount received from the college. This is the same pricing schedule we observe in college pricing data. Colleges typically post a sticker price, from which they offer discounts in the form of explicit institutional student grants. Because of this practice, we can separately observe the discounts in data on institutional student grants. As a result, we do not need to make assumptions on colleges' objective functions for the purposes of this paper.

The result of all of the above is a model of 'translated spending'. Students decide what to spend on higher education, and through competitive colleges, that same amount (plus any grants received) is invested in their human capital. This way of thinking about goods investment in human capital is much in line with the macroeconomic literature. Indeed, the empirical literature assigns a rather limited role to admissions luck (e.g. Dale and Krueger, 2002).

Some papers in other strands of the literature explicitly model the behavior of colleges, which may result in a wedge between investment and spending. For example, Epple, Romano, and Sieg (2006) analyze a model with quality maximizing colleges, peer-effects, and a preference for low-income student enrollment, 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. This work yields interesting insights into the behavior of colleges. For a number of reasons, however, these frameworks are less applicable to the macroeconomic questions this paper asks. First, all these papers take colleges as given, both in numbers and in terms of characteristics, and can therefore not explain why we see the colleges that we see. Second, the strategic interactions these models describe typically become less relevant when the number of colleges grows large. Idiosyncratic preferences by students over colleges can maintain college pricing power even then, but it is not clear whether this is an empirically relevant channel. Lastly, this work's model makes a sharp prediction on the shape of the pricing function, which other papers have to impose by assumption instead.[7] In conclusion, the above model of a competitive college market approximates reality sufficiently well for the purposes of this paper.

Not being explicit about colleges' objectives comes at a cost. How do colleges respond when government education policies adjust? This work will maintain the assumption that they

[7] Recent work by Cai and Heathcote (2018) is an important exception. Cai and Heathcote also model a competitive market for colleges, resulting in an endogenous distribution of colleges. Going beyond this paper, they treat colleges as a 'club good' so that there is a strategic aspect to college choice in their model.

simply do not. Objectives from which this would result are thinkable, although they would be non-standard.

At last, students can, in practice, choose to enter public colleges. Public colleges in the United States largely function like private ones, with the qualification that the government sets pricing schedules and determines how much money a public college has to spend on education. The availability of public colleges often depends on the place of residence. To capture all this, I model one representative public college with its own pricing schedule (also consisting of a sticker price and institutional grants) that is set by the government, to which all students have access. In practice, there is heterogeneity in public colleges, although offerings are set by local governments. This likely makes availability less responsive to demand than private college offerings (as well as dependent on geography). Modeling a single representative public college also greatly reduces computational complexity.

3.2 The Labor Market

This paper considers stationary equilibria. There is a continuum of agents with a unit mass. Each agent spawns a new agent with a mass identical to its own. We refer to the former as parents and the latter as children. The timing of the life-cycle is deterministic and equal for all households. Each agent goes through a life-cycle from age 0 to age T, representing his working life. The symbol $'$ is used to denote variables pertaining to an agent's children. The below consciously mixes explicit notation (to denote the timing of an individual life cycle) and recursive notation (to account for links between generations that are part of an infinitely lived dynasty). Appendix A.3 sketches the corresponding solution method.

There are two special phases in this life cycle: In the first period of his life, the agent has access to colleges and, if he chooses to enroll in a college, a system of student grants and loans. At a later point in life (age t^I) he makes an inter-vivos transfer to his child, who begins 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, so that these parents are said to be imperfectly altruistic). This is in line with the literature on inter-vivos transfers, which finds that transfers depend on need or effectiveness (Gale and Scholz, 1994). Modeling the entire life-cycle is useful for at least two reasons: First, several measurements in the empirical literature are taken at specific ages, so that having a model counterpart to these ages is important. Second, the model is then able to capture the life cycle of earnings as in the data, thereby ensuring that returns to education are adequately captured.

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-Huggett-Aiyagari 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. Individual gross earnings are a combination of human capital and its price, hours worked, and the realization of idiosyncratic wage uncertainty.

Compared to most models of college choice, human capital is continuous in this paper. 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. In college, human capital growth is formed by a constant elasticity production function in ability, goods, and time investment. I later show that this functional form captures returns to college well. Those who choose not to go to college or are no longer in college accumulate human capital by a function that is of the Ben-Porath (1967) type, taking only time as an input. This functional form has proven to be successful at capturing the life-cycle of earnings, as well as its heterogeneity across the earnings distribution.

Throughout, agents are assumed to be fully aware of their own ability, which is in line with the finding in the literature that students' uncertainty about their own learning ability is small (cf. Hendricks and Leukhina, 2017). Only one period of fixed length (which will later be set to four years in the data) is used to represent the entire higher education career. Human capital accumulates at the end of that period. This is somewhat restrictive with regard to the time taken to complete college. In reality, some students go to two-year colleges, and some engage in graduate studies. However, it deserves emphasis to say that these different sizes of educational investment are not ruled out: during the period, students can still spend different levels of money and time. The issue is treated with care when connecting the model to data. A similar point holds with regards to drop-outs: these are not modeled explicitly, but that does not undo the empirical strategy of this paper. All relevant data used are conditioned on college entry only.

Each agent in the model economy is linked to their parents in three ways. First, agents' ability to accumulate human capital is correlated with that of their parents. Second, parents endogenously decide how many financial resources to make available to their children as they make initial decisions on human capital investment. Third, government education policies are dependent on parental income. These mechanisms are important in assessing the impact of education policies on human capital investment decisions: when policies change and make more or fewer 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 A.2 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 the estimation of these items directly from the data. Finally, students can also choose to go to a representative public college.

3.2.1 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: $s^t = [s_t, s_{t-1}, \dots, s_1, s_0]$. 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 investment in human capital, a assets, b student loans, and v inter-vivos transfers. k denotes college choice (*work*: $k = 1$; *study at a private college*: $k = 2$; *study at a public college*: $k = 3$). For a generic variable x, $I_{[x]}$ is an indicator function that equals one when x is true and zero otherwise. q denotes gross parental wages, which is described in further detail below. The same goes for student loan repayment functions $\pi(b, q)$ and borrowing constraints. $\mathbb{E}$ is the usual expectations operator. Denote a vector of control variables as follows:

$$\mathbf{z}_t = [c_t, l_t, e_t, a_{t+1}].$$

The initial problem now consists of a college choice, meaning an individual can choose to go to college or not. If the individual does go to a private college, there is an additional choice of the level of educational investment d (which is available at any positive level). If the individual goes to a public college, educational investment d^g is set by the government (as is its price $\tilde{d}^g$). The individual's choice will depend on a fixed learning ability α, parental wages q (to be discussed below), and their initial asset holdings a_0. Formally:

$$V(\alpha, q, a_0) = \max\left\{W_0(\alpha, h_0(\alpha_0), 0, a_0), \max\left\{C^g(\alpha, q, a_0), \max_{d\geq 0} C_d(\alpha, q, a_0)\right\}\right\}$$

College enrollment lasts for one period of the model, during which the problem of an individual who goes to college d looks as follows:

$$C_d(\alpha, q, a_0) = \max_{\mathbf{z}_0(\bar{s}), b_1} \left\{ \frac{(c_0^{\nu} l_0^{1-\nu})^{1-\sigma}}{1-\sigma} + W_1(\alpha, h_1(d, e_0, \alpha), \pi_1(b_1, q), b_1) \right\}$$

subject to:

$$c_0(1+\tau_c) \leq (1 - l_0 - e_0) w h_0 x(\bar{s})(1 - \tau_n(\cdot)) - \tilde{d}(d, q, \alpha)$$
$$+ a_0(1 + r(1 - \tau_a)) - a_1 - b_1$$

$$c_0 \geq 0, \quad 0 \leq l_0, e_0 \leq 1, \quad l_0 + e_0 \leq 1,$$
$$a_1 \geq 0, \quad 0 \geq b_1 \geq -\underline{b_0}, \quad a_1 b_1 \geq 0, \quad s_0 = \bar{s}.$$

Leisure and consumption enter periodic utility multiplicatively. The utility function is tied down by parameters σ and ν.[8] Consumption and consumption taxes are paid for by what remains of net labor earnings, assets, and student loans after paying for college. Labor earnings are composed of hours worked $(1 - l - e)$, wage rate w, human capital h, and idiosyncratic shock $x(s)$. The idiosyncratic shock $x(s)$ is a function of stochastic state s, where $\bar{s}$ is the starting state that is common to all agents. The initial level of human capital is denoted by h_0.

The problem for those who enter the representative public college is the same, only that their education now costs $\tilde{d}^g(d, q, \alpha)$ and yields an investment of d^g.

An individual who does not go to college or has finished studying enters the labor market.

[8]With this functional form, the elasticity of inter-temporal substitution is given by $\frac{1}{1-\nu(1-\sigma)}$, and the Frisch elasticity by $\frac{1-\nu(1-\sigma)}{\sigma} \frac{l}{(1-l)}$.

The problem of working life is the following:

$$W_j(\alpha, h_j, \pi_j, a_j) =$$

$$\max_{\substack{\left\{\{\mathbf{z}(s^t)\}_{s^t}\right\}_{t=j}^{t^I-1}, \left\{\{\mathbf{z}(s^t)\}_{s^t}\right\}_{t=t^I+1}^{T-1}, \\ \left\{\left\{\mathbf{z}(s^{t^I},\alpha')\right\}_{s^t}\right\}_{\alpha'}, \left\{\left\{v(s^{t^I},\alpha')\right\}_{s^{t^I}}\right\}_{\alpha'}}} \mathbb{E}\left\{\sum_{t=j}^{T-1} \beta^t \frac{(c_t^\nu l_t^{1-\nu})^{1-\sigma}}{1-\sigma} + \omega\beta^{t^I} V(\alpha', q', v)\right\}$$

subject to $\forall t \in \{j, \ldots, T-1\}$:

$$c_t(1+\tau_c) \leq (1 - l_t - e_t) w h_t(d_{t-1}, e_{t-1}, h_{t-1}, \alpha) x(s_t)(1 - \tau_n(\cdot))$$
$$- vI_{[t=t^I]} + a_t(1 + r(1-\tau_a)) - a_{t+1} - \pi_j$$

$$c_t, v \geq 0, \quad 0 \leq l_t, e_t \leq 1, \quad l_t + e_t \leq 1, \quad a_{T-1} \geq 0, \quad a_1 \geq 0,$$
$$q' = wh_{t^I} x_{t^I}, \quad \alpha' \sim \Gamma_\alpha(\alpha, \alpha'), \quad s_0 = \bar{s}, \quad s_{t+1} \sim \Gamma_s(s_t, s_{t+1}).$$

This is a typical life-cycle problem, where next-period utility is discounted by β. The parameter ω discounts the value function of the child's adult life at t^I, which starts at t^{I+1}. At t^I, parents can make an inter-vivos transfer v that affects their child's initial value function. Consumption is paid for using net labor earnings and assets after student loan repayment π.

Human capital So what does an individual gain from college or time spent learning? Both increase human capital, but in different ways. Out of college, human capital production similarly follows from a Ben-Porath (1967) function. This functional form, which has been of much use in the macroeconomics literature, can match the life cycle of earnings well given the right parameterization. Key is that the time input is measured in *human capital hours*, which ensures that hours spent learning or earning are always in a direct trade-off:

$$h_{t+1} = h_t(1 - \delta_h) + \alpha(e_t h_t)^{\beta^W}. \tag{3}$$

In college, the post-depreciation gain in human capital (denoted $\Delta_\delta h_1 \equiv h_1 - h_0(1 - \delta_h)$) is assumed to have a constant elasticity in both goods and hours of human capital invested, as well as in ability. Combined with the assumption that a zero investment of either goods or time results in zero creation of human capital ($h^C(0, e_0 h_0, \alpha) = h^C(d, 0, \alpha) = 0$), this immediately yields the following:

$$\log(\Delta_\delta h_1) = \log \beta_0^C + \beta_1^C \log \alpha + \beta_2^C \log(e_0 h_0) + \beta_3^C \log d, \tag{4}$$

or in levels:

$$h_1 = h_0(1 - \delta_h) + \beta_0^C \alpha^{\beta_1^C} (e_0 h_0)^{\beta_2^C} (d)^{\beta_3^C}. \tag{5}$$

Thus, the same ability that helps learning during working life also determines learning ability in college. I assume that ability is more effective in college, meaning $\beta_1^C > 1$. The log-constant term in the above is important, because the role of ability has to nevertheless be rescaled versus that in working life, where the distribution of α is parameterized. Initial human capital, h_0, is simply assumed to be a linear function of ability, which makes the two perfectly correlated. This is further discussed in Section 4 below. Because zero goods spending in college results in an ineffective function, the model will endogenously generate a minimum level of spending among college students.

Cost of college As follows from the model of colleges above, the monetary input d 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 d also depends on ability α and on q, the gross parental wage rate, which is determined by the previous generation: $q' = wh_{t^I}x_{t^I}$.[9] The discussion of the parameterization of the model elaborates these points further.

Information structure The agent is uncertain about the next realization of his idiosyncratic earnings state $s_t \in S$. All agents start out from the same state: $s_0 = \bar{s}$. In the next period (when all individuals work) s_1 is drawn from $\Gamma_{s_1}(q)$, where the parental gross wage rate influences the probability of starting out in a good state. This allows the model to capture the importance of parental networks and influence. Thereafter s_t follows a first-order discrete Markov process with transition matrix $\Gamma_s(s_t, s_{t+1})$. 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 generally uncertain about his child's ability $\alpha' \in A$, but gets to know the child's ability right before he makes an inter-vivos transfer. (All choice variables at t^I therefore also depend on α'.) Ability is discrete and drawn from the joint distribution of parents' and children's ability $\Gamma_\alpha(\alpha, \alpha')$.

Taxation A government charges taxes on consumption τ_c, labor income $\tau_n((1 - l_t - e_t)wh_tx(s_t))$, and capital income τ_a. Labor income taxes are non-linear. The government's

[9] 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 gross parental wage rate 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 policymakers indeed attempt to rule out).

budget, after consideration of education policies, is balanced by neutral (or wasteful) spending G that does not influence any choices.[10]

Student loans The student loan system mimics the 2003 Stafford loan system as follows. At age 0, college-going students fall into one of two eligibility categories on the basis of their parents' wages at the time they become independent decision-makers. If parental wages (q) are not higher than y^*, the student qualifies for subsidized loans up to $\underline{b}^s$ as well as unsubsidized loans up to $\underline{b}^u$. If parental wages are above y^*, the student can only borrow at the unsubsidized rate up to $\underline{b}^s + \underline{b}^u$. 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 $a_1 b_1 \geq 0$. This structure follows and simplifies Abbott, Gallipoli, Meghir, and Violante (2013). After the college-going period, the natural borrowing constraint applies: all loans must be repaid by the end of working life.

After the college-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 $b_1 < 0$:[11]

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

Otherwise, $\pi_t = 0$. For those who do not enter college, $b_1 = 0$. Finally, in the above $\underline{b}_0 = \underline{b}^s + \underline{b}^u$.

3.2.2 Stationary Equilibrium

The production function takes the following functional form:

$$F(K, H) = K^\theta H^{1-\theta}. \tag{6}$$

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

[10] The model could also include a tax on inter-vivos transfers, but a reasonable parameterization would set that tax to zero: educational investments for children are exempt from taxation in the US tax code, which additionally includes sizable annual and lifetime exclusion levels. Inter-vivos transfers are also hard to observe, so that significant evasion of any remaining tax burden is to be expected.

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

Labor, capital and goods markets are perfectly competitive. We model the economy as closed to labor, and open to capital and goods. 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 r. A share δ of capital is lost to depreciation, which firms reinvest from production. 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. The government also issues and collects student debt, pays for public college subsidies, and collects taxes on labor earnings, capital income, and consumption. The government also pays for government expenses and is assumed not to hold any government debt or assets other than those mentioned. The government's budget constraint is shown in equation 14 below.

Let $x^*_\tau(\iota_t)$ denote a decision rule given states $\iota_t \in \mathcal{I}_t$ for a generic 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;
college pricing schedules $\tilde{d}(d, q, \alpha)$;
allocations K, H;
government spending G;
net exports NX, net foreign asset position NA;
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 $d(\iota_0)$ and time $\{e_t(\iota_t)\}_{t=0}^{T-1}$ investment in human capital, college choice $k(\iota_0)$, student loan borrowing $b(\iota_0)$, 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 r, tax functions τ_c, τ_n, τ_a, sets S and A, transition matrices $\Gamma_{s_1}(q)$, $\Gamma_s(s_t, s_{t+1})$, and $\Gamma_\alpha(\hat{\alpha}, \alpha)$, grant schedules $g^I(q, \alpha) \geq 0$, repayment function $\pi_t(b_1, q)$, as well as the parameters of the model,* ***the following holds****:*

- *the decision rules solve the households' problem as described in subsection 3.2.1;*

- *college pricing schedules solve the colleges' problem; as a result, college pricing schedule* *take the following form:*

$$\tilde{d}(d,q,\alpha) = d - g^I(q,\alpha), \tag{7}$$

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

$$r = F_1(K,H) - \delta_a, \tag{8}$$

$$w = F_2(K,H); \tag{9}$$

- $\lambda_t(\mathbb{I}_t)$ *are age-dependent fixed points of the 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* $\Gamma_{s_1}(q)$ *and* $\Gamma_s(s_t, s_{t+1})$*,*

  – *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:*

$$H = \sum_{t=0}^{T-1} \int_{\mathcal{I}_t} (1 - l_t - e_t) x_t h_t \,\mathrm{d}\lambda_t; \tag{10}$$

- *the market for capital clears:*

$$K = \sum_{t=0}^{T-1} \int_{\mathcal{I}_t} a_t \,\mathrm{d}\lambda_t - NA; \tag{11}$$

- *the balance of payments with respect to the rest of the world holds:*

$$rNA = -NX; \tag{12}$$

- *the market for goods clears (aggregate investment in assets equals depreciation since the equilibrium is stationary):*

$$F(K,H) = \sum_{t=0}^{T-1} \int_{\mathcal{I}_t} c_t \,\mathrm{d}\lambda_t + G + \delta_a K + \int_{\mathcal{I}_0} I_{[k=2]} d \,\mathrm{d}\lambda_0 + \int_{\mathcal{I}_0} I_{[k=3]} d^g \,\mathrm{d}\lambda_0 + NX; \tag{13}$$

- *and the government balances its budget (where the term involving $(d - \tilde{d})$ captures all grants for private colleges, and the term involving $(d^g - \tilde{d}^g)$ captures public college subsidies and grants):*

$$
\begin{aligned}
G + \int_{\mathcal{I}_t} -I_{[k>1]} b_1 \,\mathrm{d}\lambda_0 + \int_{\mathcal{I}_0} I_{[k=2]} (d - \tilde{d}) \,\mathrm{d}\lambda_0 + \int_{\mathcal{I}_0} I_{[k=3]} (d^g - \tilde{d}^g) \,\mathrm{d}\lambda_0 \qquad (14) \\
= \sum_{t=0}^{T-1} \int_{\mathcal{I}_t} (c_t \tau_c + a_t r \tau_a) \,\mathrm{d}\lambda_t + \sum_{t=0}^{T-1} \int_{\mathcal{I}_t} (n_t w(h_t) x(s_t) \tau_n(\cdot)) \,\mathrm{d}\lambda_t + \sum_{t=1}^{m} \int_{\mathcal{I}_t} \pi_t \,\mathrm{d}\lambda_t.
\end{aligned}
$$

Appendix A.3 sketches the solution method and computational procedure.

4 Parameterization

I now proceed to discuss the parameterization of the model. The parameterization targets the year 2003 or the closest possible. The reason for targeting 2003 is data availability: college enrollment in the datasets by Chetty et al. and Hoxby (2016a), ability tests of children in the NLSY dataset, as well as a number of other measurements used in the below all take place close to that year.

The parameter space consists of three parts: Some parameters are estimated outside of the model. These are described in subsection 4.1. 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 the model to their counterparts in the data. These two types of parameters are both described in subsection 4.2.

4.1 Estimation

A number of important drivers are estimated outside of the model using microeconomic data. These are 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, I do not choose a functional form. Instead, I 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 establishing a connection between the ability of mothers and their children. The test I use to assess the ability of children is the PIAT Math test, who were between 14 and 16 years old (for the sample I select) when taking the test. I then sort both mothers and children into quintiles on their respective scores and determine a transition matrix. Figure 2 displays the results graphically. Test scores are persistent yet mean-reverting, with a stronger persistence in the tails than in the middle. The overall correlation between mothers and their children's test scores is 0.38.

Because the AFQT score is constructed to generate percentiles, I assume a linear transformation of a discretized standard normal distribution of ability. Specifically, each state is assigned the expected value of an observation in the corresponding quintile of a standard normal distribution. Denoting the discretized standard normal distribution by $\tilde{\alpha}$, and its lowest entry by $\underline{\tilde{\alpha}}$, the distribution of α is formed as follows:

$$\alpha = \tilde{\alpha}\gamma + \rho. \tag{15}$$

The parameters ρ and γ are needed because ability has no natural unit. A description of how these two parameters are set follows further below.

Ex ante, there are two issues with the approach. First, these test scores may not actually be a good measure of ability transmission. As we will see, the model (as a positive prediction) produces realistic values of intergenerational persistence by a number of measures. This is perhaps the argument that provides the most comfort. In addition, these test scores are commonly used in the literature as measures of ability. In fact, our procedure essentially follows Abbott et al. (2013). Finally, it is important to keep in mind that the procedure is only used to tie down the transition of ability, but not its distribution, which follows from a common functional form assumption.

Second, since ability at the start of adult life is taken as given, one may argue that the Lucas critique applies: policy changes may lead to changes in the behavior of parents and children at earlier ages. This, in turn, would potentially alter the ability distribution at the start of adult life that we take as given. In that sense, what is called ability here should be interpreted very strictly: it is the transition and distribution of ability as currently measured. In practice, policy changes that make learning ability more worthwhile give parents incentives to invest in their children's earlier education. Parental education may also itself have causal effects on the education of children, because parents have more resources or because they teach their children differently. These channels would likely strengthen the behavioral mechanisms considered in this paper: lifting constraints on educational investment in adult life will make earlier investments more valuable as well (assuming different stages of education are complements). From that perspective, the effect sizes reported in this paper will be conservative.

Figure 2: Ability transition

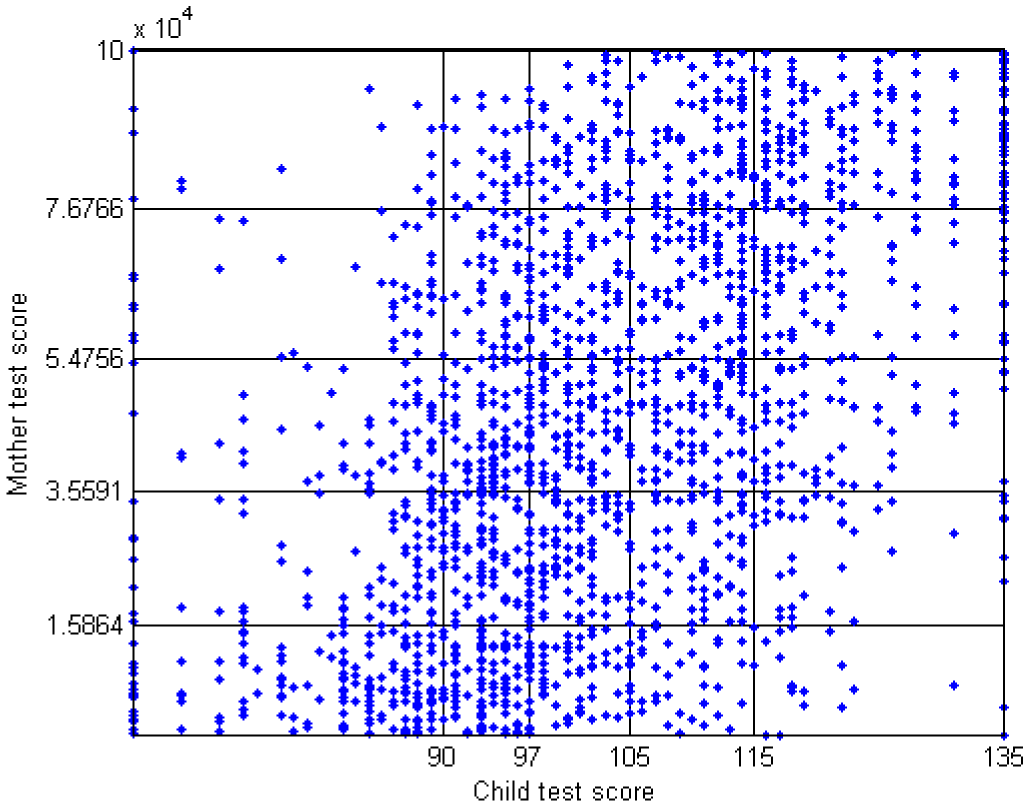

Grid lines show quantile bounds for each of the tests. The density of observations in each rectangle of the grid illuminates persistence in test scores across generations.

Earnings uncertainty For most of the life cycle, the model setup of this paper restricts idiosyncratic earnings uncertainty to be of a first-order Markov form so that only one state is required to track the idiosyncratic component of 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 parameterized period length 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 symmetric first-order Markov process with two states, which has the same persistence and unconditional variance as the estimated AR(1) model. The final $\Gamma_s(s_t, s_{t+1})$ and $x(s_t)$ are shown in Table 2. The initial state of s in the model is fixed and denoted $\bar{s}$, and is set to 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

The careful approach above is chosen because these shocks are an important element of this paper. The yearly persistence implied by the four-year probability of remaining in the same state is 0.92. This is in line with other findings in the literature. For example, Storesletten, Telmer, and Yaron (2004) report a yearly autocorrelation of 0.95.

Student grants and college subsidies Appendix A.2 provides an overview of the landscape of US education policy around 2003. In part, student aid consisted of student loans and subsidies to public colleges, which are modeled explicitly in this paper and parameterized further 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 costs for college ($\tilde{d}$) and actual investment in human capital (d). I now lay out the mapping between these two variables, and then parameterize it by estimation from data.

First, let us call the sticker prices observed in the data (a college's headline figure for tuition and fees) s^D (where the superscript D will refer to data). 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. I consider a linear relationship as follows:

$$g^D = \beta_0 + \beta_1 q^D + \beta_2 s^D + \beta_3 \alpha^D. \tag{16}$$

Given data on grants, we can estimate the parameters in this equation.

Because sticker prices are paid either through private expenditures or from aid (which we have defined broadly), we have $s^D = g^D + \tilde{d}^D$. In addition, competitive pricing schedules guarantee that $d^D = g^D + \tilde{d}^D$, so that $s^D = d^D$. Concluding all of this, investment in college is determined as follows:

$$d^D(\tilde{d}^D, q^D, \alpha^D) = \frac{1}{(1-\beta_2)}[\beta_0 + \beta_1 q^D + \tilde{d}^D + \beta_3 \alpha^D]. \tag{17}$$

We still need to connect the data variables to those in the model, since the units of account are different (a dollar in the model does not equal a dollar in the data). I resolve this by rewriting equation 17 to include only relative terms:

$$\frac{d^D}{\bar{y}^D} = \frac{\beta_0}{(1-\beta_2)\bar{y}^D} + \frac{\beta_1}{(1-\beta_2)}\frac{q^D}{\bar{y}^D} + \frac{1}{(1-\beta_2)}\frac{\tilde{d}^D}{\bar{y}^D} + \frac{\beta_3}{(1-\beta_2)\bar{y}^D}\bar{\alpha}^D + \frac{\beta_3\sigma_\alpha^D}{(1-\beta_2)\bar{y}^D}\frac{\alpha^D - \bar{\alpha}^D}{\sigma_\alpha^D}. \quad (18)$$

Here, $\bar{y}^D$ are average earnings as measured in the data. $\bar{\alpha}^D$ and σ_α^D are also assumed measurable in the data, and represent the mean and standard deviation of α^D. Now, note that this is an equation relating normalized instructional expenditure $\frac{d^D}{\bar{y}^D}$ to normalized parental income $\frac{q^D}{\bar{y}^D}$, normalized personal education expenditure $\frac{\tilde{d}^D}{\bar{y}^D}$, and normalized ability $\tilde{\alpha}^D = \frac{\alpha^D - \bar{\alpha}^D}{\sigma_\alpha^D}$. All of these terms have clear model counterparts, while the coefficients are measurable in the data.[12]

Rewriting for the model counterpart of equation 18, we get (with the superscript M referring to model variables):

$$d^M = a_0 + a_1 q^M + a_2\tilde{d}^M + a_3\tilde{\alpha}^M. \quad (19)$$

Here, $a_0 = \frac{\beta_0}{(1-\beta_2)}\frac{\bar{y}^M}{\bar{y}^D} + \frac{\beta_3}{(1-\beta_2)}\frac{\bar{y}^M}{\bar{y}^D}\bar{\alpha}^D$, $a_1 = \frac{\beta_1\bar{n}}{(1-\beta_2)}$, $a_2 = \frac{1}{(1-\beta_2)}$, and $a_3 = \frac{\beta_3\sigma_\alpha^D}{(1-\beta_2)}\frac{\bar{y}^M}{\bar{y}^D}$. All inputs underlying these terms can be estimated from data.

Next, we turn to measurement. 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), gross parental income, as well as SAT scores (combined scores), which function as a proxy for human capital.[13] I use these data to estimate equation 16 by Ordinary Least Squares, restricting the sample to 4-year colleges. The regression is done separately for private and public colleges. Observations containing zeros are excluded, except for grants. The regression is weighted by the NCES's full sample weights.[14]

Table 3 contains the estimates, as well as a measure of the explanatory power of the linear model and the number of observations used. Finally, the resulting parameters of equation 19, which are directly fed into the model, are displayed as well. From the NPSAS, we have that

[12] The counterpart to q^D, gross parental income, is q^M, a gross wage rate. Thus, we use $q^M\bar{n}$ as the relevant counterpart to turn the model wage rate into model earnings. We set $\bar{n}$ to 0.35, based on a daily time endowment of 16 hours (for each of 7 days) and a reported weekly 39.53 hours of total market work in 2003 Aguiar and Hurst (2007).

[13] Later waves of the NPSAS do not include all of this information, so that the year 1995-1996 is closest to the calibration's target year.

[14] Access to the NPSAS is restricted, so that the estimation of non-linear functional forms was not possible.

$\sigma^D_\alpha = 226.1$ and $\bar{\alpha}^D = 930.0$ 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 (standard errors in brackets)

	(16) g^D Public	Private			(19) d^M Public	Private
Constant	1127.10	376.88	Constant:	a_0/y^M	0.05	0.11
	(331.59)	(815.61)				
q^D	-0.01	-0.02	q^M :	a_1	-0.01	-0.02
	(0.01)	(0.02)				
s^D	0.15	0.24	$\tilde{d}^M$:	a_2	1.18	1.32
	(0.02)	(0.04)				
α^D	0.19	2.39	$\tilde{\alpha}^M$:	a_3/y^M	0.00	0.02
	(0.38)	(0.72)				
R^2	0.09	0.12				
Observations	~5,600	~4,800				

The regression results show that grants for private colleges depend more on merit and need compared to those for public ones, but the latter have a larger constant component. This conclusion regarding the intercept changes slightly when translating the regression results to model parameters in the right half of the table. Because sticker prices are higher in private colleges, grants tend to be higher as well, even when disregarding merit and need. In both types of colleges, not spending anything results in a positive grant when parental income is zero. Also, spending more results in more investment (since that is a one-to-one relationship) but also in more need and thus more grants, making the coefficient of spending larger than one. Negative grants could technically occur in this linear relationship for some combinations of inputs, but do not actually occur in the calibrated model.

4.2 Moment Matching

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 ($T = 12$).

Childbirth occurs at age 28, which is the average age of mothers at child birth[15], so that children start their working life when the parent is aged 46. Inter-vivos transfers are made during the period before that ($t^I = 6$).

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 rental rate of physical capital is set such that the post-depreciation yearly rate r is 1% per year. This results in an interest rate slightly below that of an equivalent closed complete market economy ($\frac{1}{\beta} - 1$). θ is set equal to the capital share of total factor income in the data (0.33).[16]

Preferences ν and σ are set to match average hours worked and the elasticity of intertemporal substitution. The former is taken to be 35%, based on a daily time endowment of 16 hours and a reported weekly 39.53 hours of total market work in 2003 Aguiar and Hurst (2007). For the latter, we rely on a meta-study by Havránek (2015), who finds that the literature's best estimate for this elasticity is 0.3-0.4 after correcting for publication bias. I use the midpoint of that range.

Inter-vivos transfers Abbott et al. (2013) do extensive empirical work on inter-vivos transfers using survey data from the NLSY97.[17] They estimate average total inter-vivos transfers between age 16 and 22 to be \$30,566 in 2000 dollars (79% of the 2000 average wage, or 20% of 4 years of average wage when accounting for the model period), and we set ω to match this figure with our one-off inter-vivos transfer.

Human capital The initial distribution of human capital (h_0) is assumed to be a linear transformation of the distribution of ability, and thereby perfectly correlated with ability. Here, the paper essentially takes the view that it is the ability to learn that is, together with actual knowledge, built earlier in life. Once the child matures, the two are then separate entities: underinvestment can lead to a level of knowledge that is low versus learning ability, and vice versa. If we were to let go of the link at an earlier age, catch-up effects might occur where an undertrained but able child, given the same educational, outperforms peers who are more knowledgeable, to begin with. While this may certainly occur in practice, we choose to

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

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

[17]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.

ignore the effect here: First, the empirical literature points in another direction, suggesting that there are strong complementaries between early and later education. Indeed it seems that the purpose of training in early childhood is, in large part 'learning to learn' what is taught in tertiary education and at work. Second, related papers that separate ability and initial human capital early in life, such as Huggett, Ventura, and Yaron (2011), find the two to be strongly correlated. Other papers have therefore proceeded in the same way as I do, notably Guvenen, Kuruscu, and Ozkan (2014).

Quantities of human capital are yet to be normalized, which is done as follows:

$$h_0 = h_{norm} + (\tilde{\alpha} - \underline{\tilde{\alpha}})\psi. \tag{20}$$

Thus, the lowest level of initial human capital in the economy is normalized to h_{norm}. The resulting normal distribution (approximate due to discretization) has mean $(h_{norm} - \underline{\tilde{\alpha}}\psi)$ and standard deviation ψ. These results are used to implement equation 19.

Summing up, the parameters γ, ρ (from equation 15) regulate the distribution of ability, the parameter ψ (from equation 20) regulates the distribution of initial human capital (while h_{norm} can be set to any computationally convenient value), and β^w and δ_h (from equation 3) regulate the build-up of human capital while at work. I set all of these parameters to capture features of the distribution of age-earnings profiles.

Huggett, Ventura, and Yaron (2011) do empirical work to establish the distribution of patterns of life-cycle earnings, taking into account time fixed effects. These data are displayed in Figure 7 (Appendix A.4). The sample consists of men who are attached to the labor force. They show: (i) that earnings increase and then decrease over the life cycle, (ii) how large this movement is versus what is given at the beginning of the cycle, (iii) that inequality grows with age, and (iv) how much inequality there is in the system overall. The model equivalents of these patterns are driven by the distributional parameters above. I take the following moments from the data that capture these patterns:

1. Average earnings at age 32 over average earnings at age 24. (1.37)
2. Average earnings at age 48 over average earnings at age 24. (1.57)
3. Average earnings at age 60 over average earnings at age 24. (1.32)
4. The variance of log earnings at age 32. (0.34)
5. The variance of log earnings at age 48. (0.42)

College effectiveness The effectiveness of college, together with the life-cycle of earnings, is informative of the extent to which human capital is determined before college. The constant elasticity functional form in equation 5 leaves the following parameters to be determined: β_0^C, β_1^C, β_2^C, and β_3^C.

Key to some of the questions this work is after is the relative importance of financial resources in the college production function of human capital. Hoxby (2016a) identifies the effectiveness of money across the distribution of colleges in a setting where financial investment is approximately exogenous, meaning that ability is controlled for. I target these results, which are described further below.

β_0^C determines how effective ability is in college versus at work so that the share of the population that decides to go to college is informative. Using data from Chetty et al. (2017), which is on the relevant cohorts, I find that 75% of individuals in the relevant cohort enroll in some sort of college. This is the relevant empirical counterpart for the model.

It is generally worth noting at this point that the paper takes a broad view of human capital: human capital is continuous, and I do not explicitly deal with dropouts, 2-year colleges, professional degrees, etcetera.

Combining all this, average inputs of time and money then imply the remaining parameters:

1. Data on time use by students are hard to come by and do not generally paint a consistent picture. Perhaps most important here is to capture the ability of students to finance their education by work time. I use the 2003 American Time Use Survey, and restrict the sample to those enrolled in college and spending at least some time attending class. I then calculate how much time these students spend on education (including education-related travel) versus work (including work-related travel), aggregating individuals using 'ATUS final weights'. The ratio of the former category versus the latter is 2.02: active students spend about twice as much time studying as they spend working. I then halve this ratio twice: once to account for the time during which colleges are out of session[18], and once to account for time actually spent in college during the 4-year period[19].

2. To tie down spending on education, the parameterization targets the share of GDP spent on tertiary education from private sources. Because private spending in our

[18]The assumption here is that time worked stays constant over the year, and time spent studying goes to zero during half the year. The exercise remains approximate due to the lack of appropriate aggregate data.

[19]This is to correct for dropouts from college, 2-year-colleges, etcetera. Again, the exercise remains approximate due to the lack of appropriate aggregate data.

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 0.86% of the 2003 average wage.

Hoxby (2015, 2016a, 2016b) measures causal returns of a marginal dollar investment upon college entry from discounted lifetime income. Her method is similar to that of earlier work by Dale and Krueger (2002) and produces consistent results, but for a wider range of colleges. Combining administrative data on incomes, clearinghouse data on college applications, and data on college expenditures, she compares students who are 'on the bubble' of getting admitted to a college. Student SAT scores help to identify students who are close to being admitted or rejected. The assumption that identifies causal effects is the following: For students whose credentials are close to the typical cut-off, admission can be thought of as a random event. Paired comparison methods then establish the extra monetary investment caused by admission and the subsequent returns to that investment. In doing so, the least selective college is normalized to add zero value.

Results show a marginal dollar return of around 3.5 after discounting for colleges that are at least somewhat selective, and these returns increase slightly in college selectivity. The results are relevant to college entry, so that college dropouts, 2-year or 4-year colleges, and all other such issues, are averaged out. Hoxby's results have a model equivalent: I simulate the effect of an exogenous extra dollar investment in college on discounted lifetime income (using the same discounting method as Hoxby). As in the quasi-experimental setting, agents can adjust their subsequent choices in response to this extra investment. The average of the resulting returns is a model moment that can be set to match the typical return reported in Hoxby (2016b).

Public college Two choices are required regarding the representative public college. What is the cost of attending before any grants (the sticker price), and by how much does the government directly subsidize the college? According to Johnson (2014), the subsidy rate for an average public college is about 53%, which leaves some $5,640 a year (or 10.48% of average earnings) of an average $12,000 in spending per student per year to be paid for by students and grants (data for 2011-2012). Therefore, $5,640 a year is the average sticker price.

Student loans I follow the structure of Abbott, Gallipoli, Meghir, and Violante (2013) to model the student loan system around 2003, but with some simplifications. The main simplification is that I do not model private student loans. As detailed in Appendix A.2 and in Abbott, Gallipoli, Meghir, and Violante (2013), these were a small source of financing, and mostly available to students whose parents were sufficiently credit-worthy (while including them comes at a significant computational cost).

The parameters y^*, $\underline{b}^s$, $\underline{b}^u$, r^s, and r^u are informed by the following moments:

1. Stafford aid was 0.33% of GDP or 0.31% of the average wage (College Board, 2013).
2. Subsidized Stafford aid was 54.54% of total Stafford aid (College Board, 2013).
3. The subsidized loan limit over the unsubsidized loan limit, which was 0.95.[20]
4. Interest rates for either type of student loan was around 4% (or 17% on a 4-year basis) in 2003. [21]

The repayment period length m is set to 20 years. While the initial repayment period has typically been ten years, this is easily extended in practice.

The above moments with regards to the Stafford loan system are chosen to accurately represent the generosity of the program overall, as well as to specific family income groups. Cumulative loan limits for the two types of Stafford loans exist, and we use these to tie down the relative generosity of the two programs in terms of available funds. However, whether students can borrow up to these limits depends on a number of other factors (for example, their class level, dependency status, cost of attendance, and financial need), so that we instead focus on matching overall amounts of borrowing. Costs of student loans (interest rates) are taken from the data, ensuring an accurate representation of that aspect.

Tax policies Guvenen, Kuruscu, and Ozkan (2014) collect data on US earnings taxes at different income levels for the year 2003 from the OECD. Using these data, I directly estimate the two parameters of the much-used tax function described in Heathcote, Storesletten, and Violante (2017). All of that results in the function below, where $\bar{y}$ are the average United States earnings (and the same parameter that is used in the implementation of equation

[20] 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, while the aggregate unsubsidized loan limit in 2003 was $18,000, which was 45.37% of GDP per capita at the time.

[21] According to FinAid (2018), interest rates for Subsidized and Unsubsidized Stafford loans were the same in the early 2000s, the rate being 4.06% in 2002-2003 and 3.42% in 2003-2004.

19).[22]

$$\tau_n(\cdot, \bar{y}) = 1 - \frac{1}{1.3434}\left(\frac{n_t w h x_t}{\bar{y}}\right)^{-0.11867}.$$

I take the consumption and capital income tax rates from McDaniel (2007): $\tau_c = 0.075$, and $\tau_a = 0.232$ for 2003. Government expenditures G are calibrated to clear the government budget.

Overview Table 4 provides an overview of parameters set outside of the model and their values. Table 5 lists parameters that were set to match moments. It displays the final parameter values (columns one and two), and the set of moments as measured in the model and in the data (columns three and four). Because these parameters are set jointly, parameters and moments do not necessarily correspond. Nevertheless, moments are presented next to the parameters they are considered to be informative of. Percentages refer to either average wage or average wage per capita.

Table 4: Parameters set outside of the model

	Value	Moment
β	0.949	Discount rate
δ_a	0.047	Asset depreciation rate
$F_1(K,H)$	0.088	Rental rate of physical capital
θ	0.314	Capital share of income
r^s	0.170	Subsidized Stafford loan rate
r^u	0.170	Unsubsidized Stafford loan rate

Few of the parameters in Table 5 have a natural interpretation. The value for ω suggests that parents count their children's value function for a fifth of their own at the age where the children mature. Depreciation of human capital is 12% during a four-year period. The elasticity of human capital growth in college is largest in ability, followed by money and time invested.

[22] I 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 5: Parameters set to match moments

Parameter	Value	Moment	Model	Data
σ	3.32	Elasticity of intertemporal substitution	0.35	0.35
ν	0.80	Average labor supply	0.41	0.39
ω	0.20	Average inter-vivos transfer	18%	20%
γ	0.90	Variance of log earnings: age 32	0.34	0.34
ρ	1.90	Average earnings: age 48 vs. age 24	1.58	1.57
β^W	0.10	Average earnings: age 32 vs. age 24	1.39	1.37
δ_h	0.12	Average earnings: age 60 vs. age 24	1.45	1.32
ψ	0.40	Variance of log earnings: age 48	0.48	0.42
β_0^C	0.70	Share with some college	76.11%	74.65%
β_1^C	1.58	Share of average wage spent on tuition	0.86%	0.83%
β_2^C	0.11	Time spent on education versus work	0.46	0.51
β_3^C	0.66	Average causal dollar return	3.18	3.50
$\tilde{d^g}(d^g, 0, 0)$	0.73	Average sticker price tuition	10.48%	10.48%
d^g	1.56	Average subsidization rate	0.53%	0.53%
y^*	25.00	Subsidized versus Unsubsidized Stafford aid	50.16%	54.54%
$\underline{b}^s$	0.27	Subsidized versus unsubsidized loan limit	0.95	0.95
$\underline{b}^u$	0.28	Overall Stafford aid	0.36%	0.31%

5 Implications

This section and Appendix A.4 describe positive implications of the model that have not been targeted in parameterizing the model.

5.1 Intergenerational Mobility

No measure of intergenerational persistence of economic outcomes has been targeted in the model's parameterization. This subsection compares model predictions against actual measurements of persistence. The success of the model summarized below provides confidence in the ability transition matrix that was based on test scores.

Table 6 contains several measures of intergenerational persistence, first for the model and then for the data. A range of estimates of the IGE exists in the literature. In their reviews of the literature, Lee and Seshadri (2014) and Landersø and Heckman (2017) respectively arrive at ranges of 0.4–0.6 and 0.3–0.5. The range of estimates is wide due to differences in sample selection and treatment: one can restrict the ages at which earnings are measured, the labor market attachment of individuals, their gender, etcetera. I report the model IGE using the entire model population (i.e. all age groups) and then again controlling for age.

Table 6: Measures of Intergenerational Mobility

	Model	Data
IGE	0.34	0.3–0.6
IGE with controls for age	0.31	
Intergenerational rank correlation of earnings	0.30	0.34
Correlation in educational attainment	0.21	0.11–0.45

Rank correlation measures are typically a bit lower. One such measure, by Chetty, Hendren, Kline, and Saez (2014) comes out at 0.34, in line with our model (again using the entire model population). Finally, the intergenerational correlation of educational attainment can also be measured. For the model, I do so simply using indicators of college entry. As expected, the outcome is a bit below the other measures and well within the range reported in the literature (see Mulligan (1999, Table 1)).

5.2 Heterogeneity of Intergenerational Mobility

Part of the motivation for this paper is the large heterogeneity in IM by college enrollment. Figure 3 reports the model equivalents to the measures used in Chetty et al. (2017), which are displayed in Figure 1. Instead of grouping students by the college they go to, I form quintiles of students by college spending (since spending is what identifies a college in the model).[23] I then measure average ability by college quintile (where individual abilities have been normalized to have mean zero and a standard deviation of one.). The patterns we see are qualitatively and quantitatively in line with their empirical counterparts. The share of low-income students (defined as a family income in the bottom 20%) amongst those that go to college ('access') falls as the average ability of students in a spending quintile rises. At the same time, the likelihood that a student from a low-income family reaches the top 20% of the income distribution ('success') rises in the average ability of students. The product of the two, the share of college students that go from the bottom to the top of the distribution ('mobility'), is flat across ability. The same patterns hold for the spending quintiles themselves, where spending has been normalized by the average labor earnings in a model period.

5.3 College Entry and Heterogeneity

Lochner and Monge-Naranjo (2011b) show the gradient of college enrollment by measured ability and family income empirically, based on NLSY97 data for the early 2000s. Belley and Lochner (2007) obtain qualitatively similar results. Enrollment rises strongly in ability,

[23]There are few distinguishable quintiles due to bunching in the public college.

Figure 3: Access, Success, and Mobility in the Model

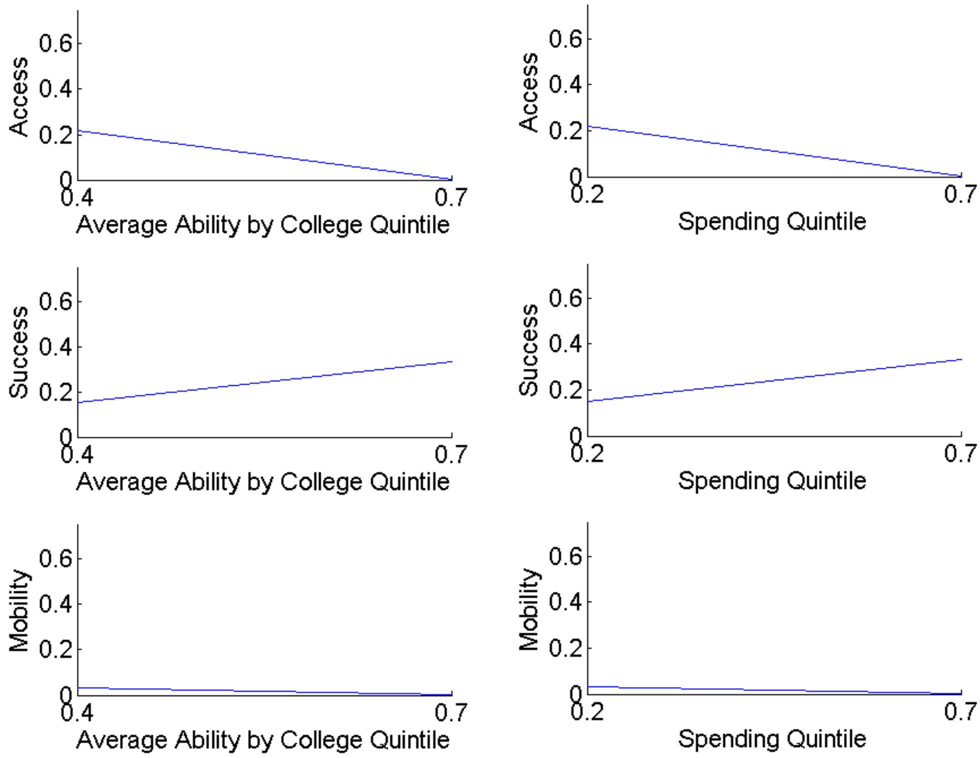

and also (but less strongly) in family income. Figure 4 displays the model equivalent.[24] The top three ability quintiles are essentially unconstrained at the extensive margin (to enroll in college or not).[25]

Figure 4 also splits entry between public and private colleges and shows that the majority of students enroll in the former. In the data (e.g., the cohorts analyzed in Chetty et al. (2017)), almost 80% of students go to 'public' colleges, although this category also includes 2-year private, not-for-profit institutions. As we would expect, children of lower-income families tend to go to public colleges.

A similar pattern is visible in Figure 5, which shows goods invested in education by the same split. Investment is expressed in terms of 2003 average earnings, which is $168,780 for

[24]It is worth noting that family earnings here are not equivalent to q from before, due to labor supply and a difference in timing of measurement.

[25]The model generates a surprising enrollment pattern for those in the second ability quintile. Members of this quintile that do not enroll in college are exclusively agents who have not received any transfers from their parents. At the same time, their access to grants reflects their parents' income position, leading to a negative gradient in family income.

Figure 4: College Entry by Ability and Family Income

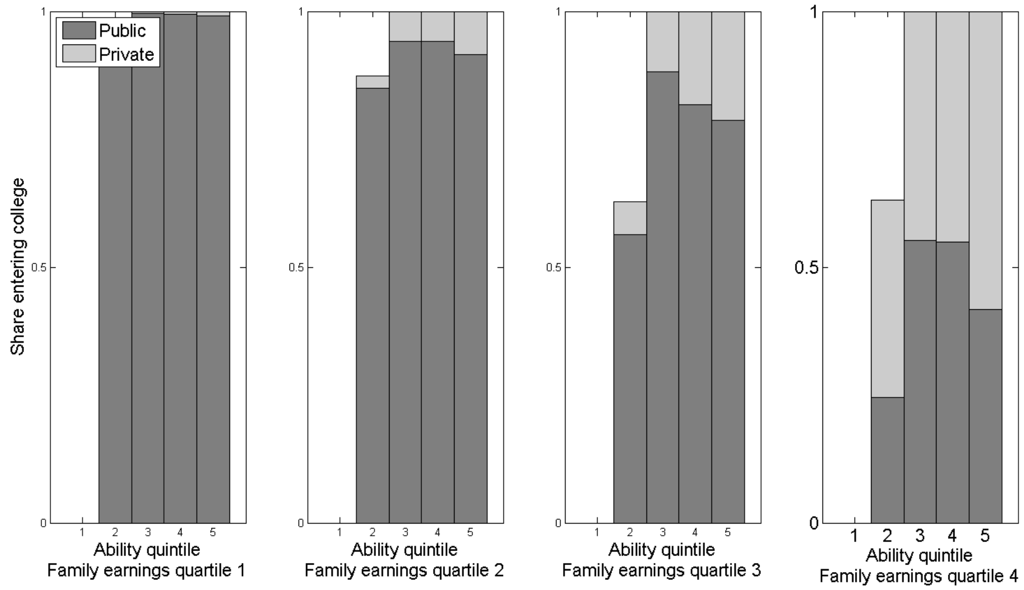

Figure 5: Average Investment in Education by Ability and Family Income

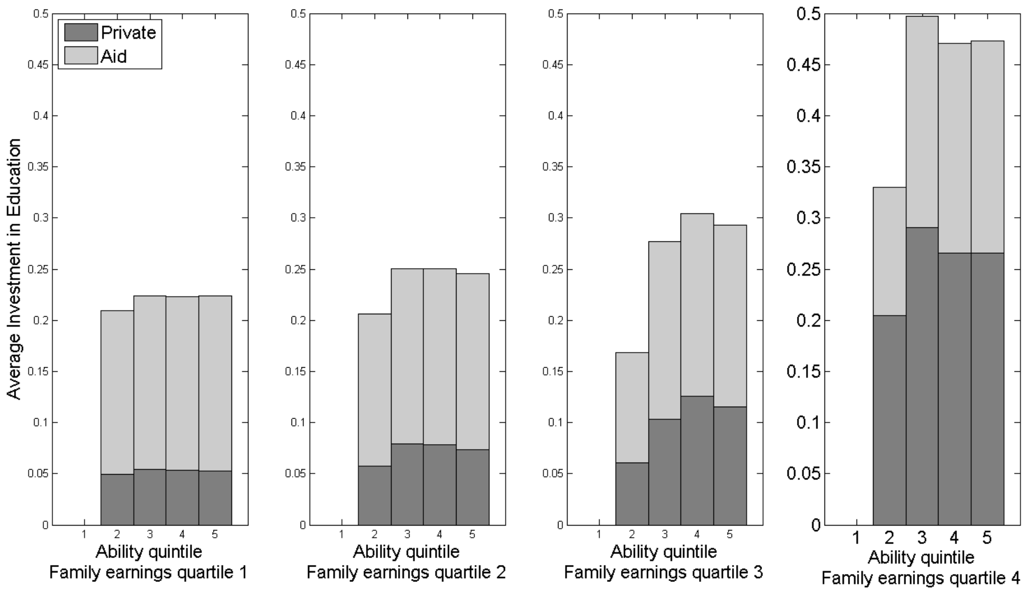

a four year period. Investment is split in private spending (including spending financed by loans) and aid, which consists of both grants and subsidies in the case of public colleges. The pattern is apparent: private spending is mostly driven by family resources and much less by ability.

Aid is remarkably stable over both gradients. This is due to three factors. First, there is only one representative public college, so that subsidies are the same for all that go there. Second, grants for private colleges are much larger, bringing aid for the students that go there into the same range. Third, for both types of colleges, aid schedules are dominated by the intercept, meaning they are roughly the same for all their students.

Because ability is perfectly observable, the college enrollment gradient in ability is much stronger than in the data. It is fairly obvious that a pattern of increasing college entry in ability could be achieved by introducing 'preference shocks' that depend on the parental background, as is typical in some parts of the literature. I do not pursue this for three reasons: First, it makes the model less parsimonious. Second, it is not obvious to what extent constraints, measurement error (since SAT scores and the like are an imperfect measure of ability), and preferences are each responsible for the pattern observed in the data. Third and most importantly, the results presented here bear out one of the results of the paper: even when there are no constraints to overall college enrollment (the choice to go to college or not, which I call the *extensive* margin of college choice), students may still be significantly constrained in their choice of a particular college (which I call the *intensive* margin of college choice). After all, investment in education still depends on family earnings.

For completeness, Figure 6 displays the distribution of private colleges by their total educational investment. The investment level of the representative public college is highlighted as well (normalized by the average labor earnings in a model period). Private colleges invest more in students than public ones. Their distribution is increasing at lower levels of spending, which is in part due to students at the lower end that crowd into the public college.

Finally, returns to education are also heterogeneous by college. The approximate average of the marginal per-dollar returns reported in Hoxby (2016b) was a target in the parameterization of the model. That paper also shows some heterogeneity of returns to college, with the per-dollar return growing in the average SAT scores of incoming students. The same pattern is present in model-generated data: marginal per-dollar returns grow in students' ability.

5.4 Life Cycle

Labor earnings are front and center in the model presented here. How do these look in the model versus in the data? Some of the model's parameters have been tied down targeting age-earnings profiles from Huggett, Ventura, and Yaron (2011), as discussed above. Figure 7 shows these profiles in full, together with their model equivalents. The model matches the earnings life-cycle overall, although with a shortage of curvature. This is also the case in the work by Huggett et al., who take a similar approach. The model also generates too strong an increase in the variance of log labor earnings, but matches the data well around the age

Figure 6: Distribution of Colleges

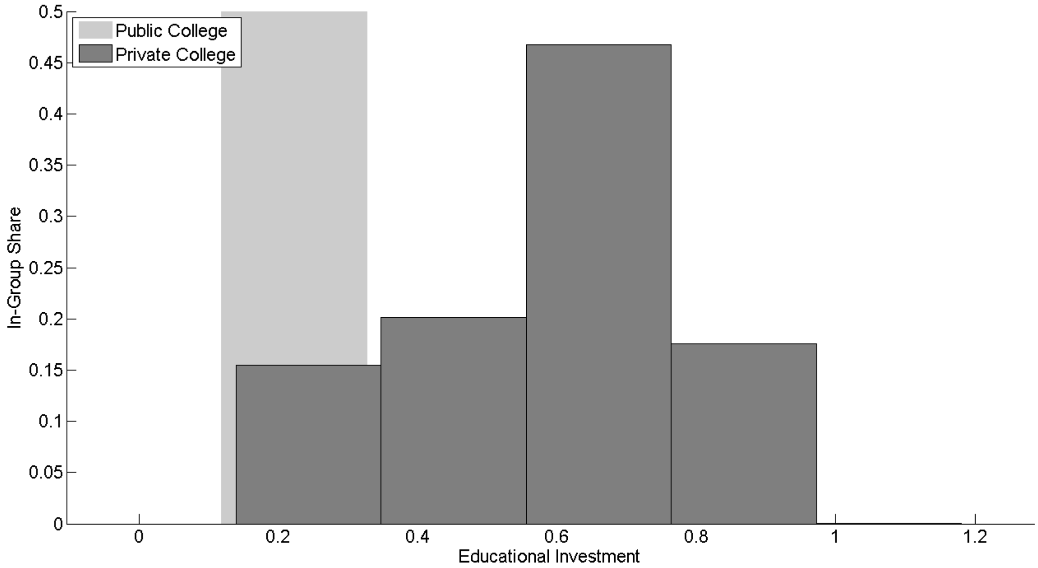

IM is typically measured.

Figure 7: Age-Earnings Profiles

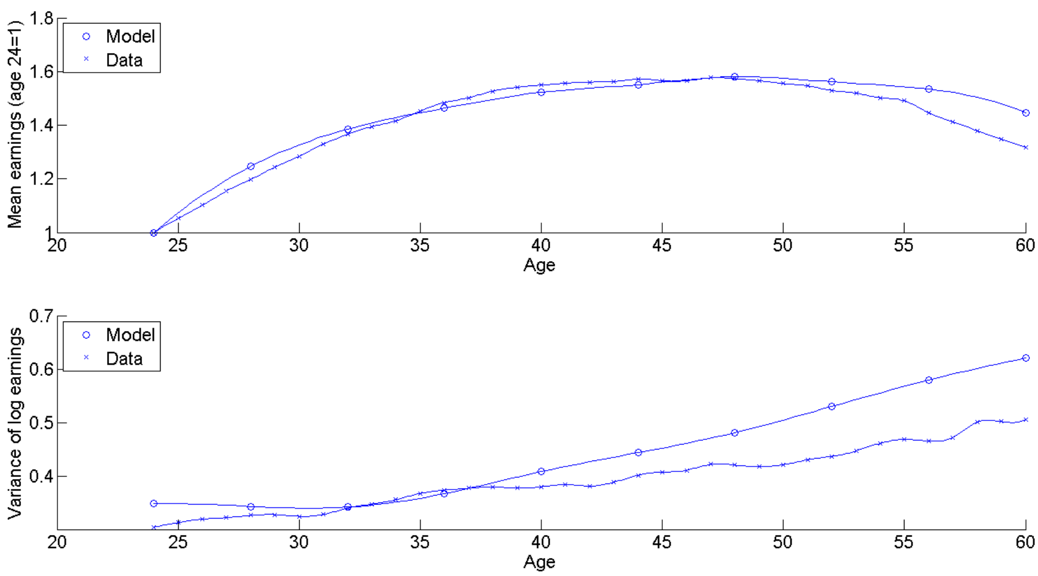

Appendix A.4 discusses further model implications regarding labor earnings, sources of college funding, and on IM when controlling for college entry.

6 Counterfactuals

This section conducts policy experiments within the model, by comparing stationary equilibria of the model economy. It begins with a decomposition of the persistence of earnings over generations. Next, it considers the effects of different higher education policies. Then, it considers the transition from one steady state to another. Finally, it discusses the robustness of the results to a key assumption.

6.1 A Decomposition of Earnings Persistence

I begin with a decomposition of the persistence of earnings over generations. There are three intergenerational links in the model: ability is correlated over generations, parents can transfer money to their children, and education policies take into account parental wages. Education policies, in turn, consist of a subsidized public college, grants, and (Subsidized and Unsubsidized) Stafford student loans.[26]

I now decompose the intergenerational persistence of earnings into components by removing each of these links. First, I remove all persistence from the ability transition matrix by assigning an equal probability to each destination level of ability (for each starting level). Second, I set ω, the parameter that determines the extent to which parents internalize the well-being of their children, to zero. As a result, parents will no longer make any inter-vivos transfers. Third, subsidies to the public college, student grants, and loan limits are all set to zero. I then undo each of these changes in the same order. The results are displayed in Table 7. I display both the IGE and the rank correlation measure of intergenerational persistence and also provide a break-down in percentages of the baseline model. Earnings are measured at age 36 for both generations, so that the measures are as clean as possible.

Table 7: Decomposition of Intergenerational Persistence

IGE		Rank Corr.		Ability Persistent	IV-transfers Full	Policies Full
0.0	(0%)	0.0	(0%)			
0.06	(38%)	0.08	(53%)	✓		
0.10	(63%)	0.10	(67%)	✓	✓	
0.16	(100%)	0.15	(100%)	✓	✓	✓

Roughly half of the intergenerational persistence in the baseline model is driven by ability (as measured at age 18). The remainder is split between inter-vivos transfers and education

[26] Access to public college is not conditional on parental earnings, but I include this policy in my decomposition nevertheless.

policies, with the latter being twice as important as the former. Thus, education policies account for about a third of the overall persistence. Importantly, inter-vivos transfers are allowed to adjust endogenously to changes in education policies. That adjustment mechanism is important when measuring the impact of policies, as has previously been analyzed in work by Abbott, Gallipoli, Meghir, and Violante (2013).

Is the effect of transfers due to their overall level or their distribution? These two components can be separated by an extra experiment in which all agents are given the average transfer (measured in the experiment with transfers but without policies). Unreported results show that this yields roughly the same persistence as the experiment without any transfers at all. In short, it is not the level of the transfers that matters, but who gets them. Transfers matter because they help pay for college.

How important are different education policies? We can investigate this by re-activating them one-by-one, starting with the public college, then grants, and then subsidized and unsubsidized loans. Table 8 presents the results. Again, inter-vivos transfers and other choices are alloed to adjust endogenously. It turns out that all policies increase persistence, with loans having the strongest impact. This is in line with the theoretical ambiguity discussed in Section 2, to which we will return below.

Again we can ask whether it is the level of these policies or their distribution that matters. Similar to the case of transfers, I run two extra experiments. The first gives the average grant to all that enroll in college (in the experiment which includes grants but not loans). The second gives the average loan to all that enroll in college (in the baseline model). Unreported results show that giving average grants actually reduces persistence compared to the experiment without grants. The same applies to loans: giving average loans reduces persistence compared to the experiment without loans. Again, it is the distribution that matters.

A look back at the distribution of aid in Figure 5 clarifies why. It is students from high-income families that receive most aid. This, in turn, is a consequence of their higher spending on education: in the data, grants are strongly related to college sticker prices. Evenly redistributing aid results in a policy that is more targeted at low-income students, which is where the positive effect on IM comes from. This brings us to the next lesson from these extra experiments: while policies modeled after actual policies increase persistence, there are some policies that would reduce persistence. The effect of education policies on IM depends on the shape the policies take. We will return to this below.

Regarding the sources of persistence, it remains to note that the effects of education policies are significant in size: they are of the same order of magnitude as cross-country differences in

Table 8: Decomposition by Specific Policies

IGE		Rank Corr.		Public College	Grants	Loans Subs.	Loans Unsubs.
0.10	(63%)	0.10	(67%)				
0.10	(63%)	0.11	(73%)	✓			
0.12	(75%)	0.13	(87%)	✓	✓		
0.15	(94%)	0.14	(93%)	✓	✓	✓	
0.16	(100%)	0.15	(100%)	✓	✓	✓	✓

persistence (cf. Corak, 2013). They are also comparable to the effect of significant reductions of tax progressivity, which have been reported as important drivers of persistence elsewhere in the literature (cf. Holter, 2015).

6.2 The Effects of Higher Education Policies

I now proceed to explore the theoretical ambiguity from Section 2 in depth. Thereafter, I will analyze the set of policies currently in place.

6.2.1 Removing Borrowing Constraints

Typically, the stated goal of government student loan schemes is to alleviate borrowing constraints that students would otherwise face.[27] I will now study how intergenerational persistence would look if all borrowing constrains on students were indeed removed. Specifically, I remove all education policies (including the public college) and then let all students borrow up to the natural borrowing limit in period zero. I interpret this experiment as if the government is providing loans at market rates. Any net loss to the government from this change comes out of government expenditure.

Table 9 reports the results. Both measures of persistence fall, but the effects are small. This is not because removing borrowing limits has no effects: education investment is much larger than in the baseline model. Instead, as we will see now, opposing effects on mobility cancel each other out.

[27] There are at least two reasons why the First Welfare Theorem breaks down in the environment this paper studies (apart from its overlapping generations structure), both resulting in a rationale for government intervention. First, government taxation drives a wedge between the private and social returns of labor, and therefore between the private and social returns to education. This creates potentially complex optimal (education) policies that have filled an extensive literature (see Subsection 1.2 for an introduction). Second, markets are incomplete: wages are subject to idiosyncratic shocks, and potential students have limited access to borrowing. Student loans address this latter issue.

Table 9: Removing the Borrowing Limit

IGE		Rank Corr.		Baseline	Unconstrained
0.16	(100%)	0.15	(100%)	✓	
0.15	(94%)	0.14	(93%)	✓	✓

Let us return to the discussion in Section 2. There, equation 2 established the following decomposition of the IGE (of wages, ignoring labor supply):

$$\beta^{IGE} = \underbrace{\frac{Var(\log h)}{Var(\log h) + Var(\log x)}}_{\text{Weight of } h} \underbrace{\frac{Cov(\log h, \log h')}{Var(\log h)}}_{\text{LPC of } h' \text{ on } h} + \underbrace{\frac{Var(\log x)}{Var(\log h) + Var(\log x)}}_{\text{Weight of } x} \underbrace{\frac{Cov(\log x, \log h')}{Var(\log x)}}_{\text{LPC of } h' \text{ on } x}$$

The point was the following: the impact of education policies is theoretically ambiguous because they have two opposing effects. Relieving financial constraints makes children less dependent on their parents. This may reduce both projection coefficients in the above formula. However, the IGE is a weighted average of the two, so that changes in the weights are crucial. If policies increase the variance of log human capital, emphasis will shift to the first projection coefficient, which is typically the larger of the two. As a result, measured persistence may go up rather than down.

Table 10 contains a full numerical decomposition. All effects go in the direction we expected: human capital becomes more mobile over generations. Going against this channel, the variance of human capital increases. On the whole, intergenerational persistence falls.

Table 10: The effect of financial constraints on the IGE

	Baseline	Unconstrained
LPC of h' on h	0.53	0.46
LPC of h' on x	0.05	0.00
$Var(\log h)$	0.11	0.16
$Var(\log x)$	0.08	0.08
β^{IGE} (wage rates)	0.31	0.29
β^{IGE}	0.16	0.15

Moving on to a different comparison, Table 11 contrasts the two sets of policies that are the focus of Section 2: a constrained and an unconstrained policy. The constrained case refers to the model without any education policies. The unconstrained case is the same as before: students can borrow at market rates up to the natural borrowing constraint. Again, the unconstrained policy results in a lower persistence of human capital. But now the variance of human capital more than doubles when the constraints are removed. As a result, removing borrowing constraints reduces IM in this comparison.

Table 11: The effect of financial constraints on the IGE

	Constrained	Unconstrained
LPC of h' on h	0.48	0.46
LPC of h' on x	0.03	0.00
$Var(\log h)$	0.07	0.16
$Var(\log x)$	0.08	0.08
β^{IGE} (wage rates)	0.24	0.29
β^{IGE}	0.10	0.15

It is perhaps in Table 11 where the paper's main point is borne out best. We observe that there is a direct trade-off between IM and efficiency: when we consider the effect of a policy that removes borrowing constraints, earnings become more persistent across generations.

I now investigate in more detail why the variance in log human capital increases. The variance in log human capital can change for two reasons: Agents with different abilities may naturally differ in their human capital. We will call this between-ability dispersion. But agents may also differ in human capital when they have the same ability to learn but receive a different education. This is within-ability dispersion. By the law of total variance, we have the following decomposition:

$$Var(\log h) = \underbrace{Var(\mathbb{E}[\log h|\alpha])}_{\text{between-ability dispersion}} + \underbrace{\mathbb{E}[Var(\log h|\alpha)]}_{\text{within-ability dispersion}}$$

Removing borrowing constraints will increase between-ability dispersion if high ability types are more constrained or benefit more from additional investment. But within-ability dispersion will decrease as financial constraints are removed, because agents with the same ability will make more similar educational choices. Table 12 shows numerical results for this decomposition for the constrained and unconstrained policies. As it turns out, the within-ability component is generally small (zero up to two decimals) in comparison to the between-ability component. As a result, the variance of log human capital increases when financial constraints are removed.

In much of the public economics literature on education (discussed in Subsection 1.2), education policies are regressive rather than progressive policies because they benefit the future rich. Very similarly, education policies in this paper's model tend to increase inequality.

What happens to educational choices when there are no constraints to borrowing? All potential students in the top four ability quintiles go to college, raising overall college enrollment to 80%. None of these students go to the public college, because their optimal investment level turns out to be much larger than the investment the public college makes. Investment

Table 12: Decomposing the variance in log human capital

	Constrained	Unconstrained
Between-ability dispersion	0.07	0.16
Within-ability dispersion	0.00	0.00
$Var(\log h)$ (wage rates)	0.07	0.16

quadruples for those that spent the least before. Differences in college spending by ability almost disappear. The remaining college heterogeneity is much smaller, indicating that most college heterogeneity is due to financial constraints at the intensive margin. The fact that heterogeneity in human capital increases while investment heterogeneity actually shrinks, can be attributed to the greater effectiveness of investment in human capital for higher-ability students.

Other model variables change as well. Average labor earnings increase by 14%. The college premium (measured at age 32) increases to 2.23 (a 39% increase). The composition of college financing changes: IV-transfers are now only used to pass assets on to children and not to finance college. As a result, they are now lower for those that go to college than or those that do not (altruism towards the latter group is more effective). Hours worked in college are almost zero. At the same time, hours studied (which are a complement to monetary investment) increase threefold. Finally, the patterns in Figure 3 ('access', 'success', and 'mobility') remain qualitatively similar, but access to high-investment colleges for students from low-income families improves significantly.

6.2.2 Current Policies

Let us now return to the set of current policies. How does the theoretical ambiguity work here? Table 13 again shows the effects of adding education policies one-by-one (as in Table 8), but now provides a model-based measurement of each of the terms in equation 2.

Table 13: Double decomposition

	LPC of h' on h	LPC of h' on x	$Var(\log h)$	$Var(\log x)$
No Policies	0.48	0.03	0.08	0.08
Public College	0.47	0.02	0.09	0.08
Grants	0.50	0.03	0.09	0.08
Subsidized Loans	0.52	0.05	0.10	0.08
All Loans	0.53	0.05	0.11	0.08

From the decomposition in Table 8, we learned that each of these policies add persistence.

We now see why: the relative importance of human capital ($Var(\log h)$) in the weighted sum increases with each step. The weight of idiosyncratic shocks ($Var(\log x)$) remains constant. As a result, the IGE goes up.

But how about the other terms in the decomposition? The intergenerational LPC of human capital actually increases with policy. This is counter to the expectancy that policies help disadvantaged children move up. The reason once again turns out to be the increased role of human capital: these policies increase income inequality (by increasing human capital inequality) so that the income differences between advantaged and disadvantaged families actually grow. Because these policies do not alleviate all constraints (at the intensive margin), increased differences in income then still have an effect across generations. High-income parents can now transfer even more resources to their children, allowing them to make even greater investments in college. The same line of reasoning explains the remaining term of the decomposition: increased inequality in human capital increases the effect of parental luck on the educational outcomes of children.

This finding is of some importance. It implies that even if we were able to measure IM in terms of human capital rather than (noisy) earnings, welfare-improving policies may still reduce IM. Thus, the caveats around the interpretation of IM go beyond mere measurement issues.

To understand this better, let us consider the effect of introducing student loans (in addition to grants and public colleges) on educational investment. Figure 8 shows investment by the same groups as before. Investment increases for all groups, but the increase is much stronger for students from high-income backgrounds.

As long as students are still constrained, parental transfers matter. Parental transfers actually become larger rather than smaller under these policies (conditional on receiving any transfer). While parental transfers are in principle a substitute for policies, higher income inequality in the presence of constraints can still mean that transfers become more unequal.[28] As we saw before, this effect is not universal. When removing all borrowing constraints, parental transfers no longer increase in education spending, and the intergenerational LPC of human capital decreases.

[28] Parental transfers also complicate welfare comparisons at an individual level. Students from advantaged families may, in some cases, prefer an absence of student loans, because it induces their parents to transfer more money. Parental transfers are preferable to student loans from the viewpoint of the child because the latter must be repaid. For that reason, I do not delve into welfare comparisons between policies at the individual level. As one would expect, numerical results show clear aggregate steady-state welfare gains for the education policies included in the model (when using, for example, a Utilitarian criterion to aggregate).

Figure 8: Average Investment in Education: No student loans versus Baseline

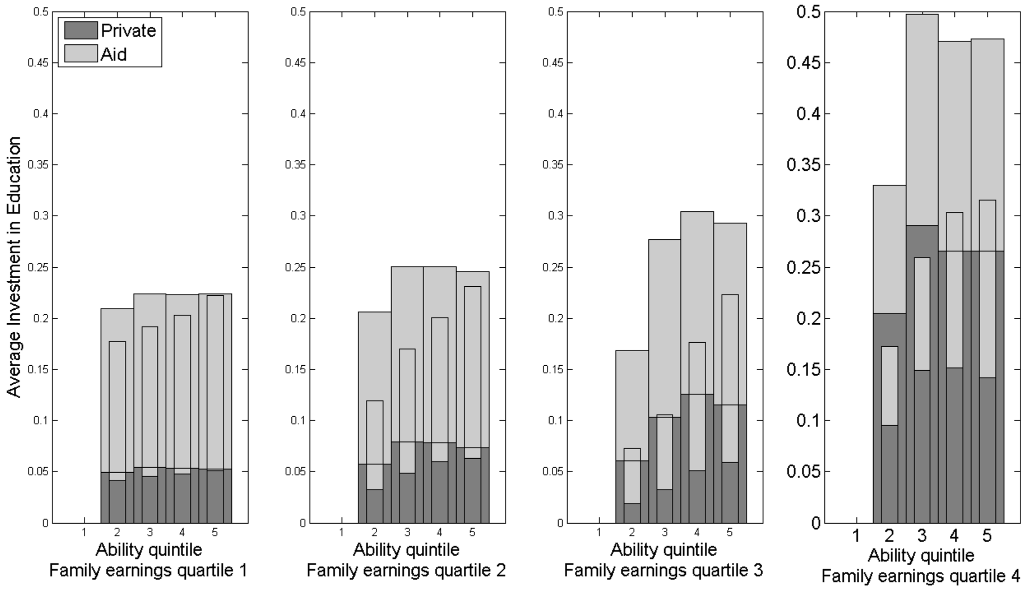

The thinner bars represent the case without student loans.
The wider bars refer to the baseline model, and display the same data as Figure 5.

6.3 Transition

So far we have compared steady states only. That raises the question: how long does it take to move from one steady state to another, and is the adjustment linear? The model setup allows for a simple way to assess this issue. A transition analysis can be done by simulating new generations under some set of policies, using the steady state of another set of policies to draw parental characteristics. This is like having a set of policies introduced when children become of college-age. It comes at the cost of some realism, as a policy introduced at one point in time would typically impact people of different ages at the same time. But the benefit of this approach is that the results are not smoothed by having a mix of policies applied to any cohort.

Figure 9 shows the IGE and it's components. Generations up to generation 0 live under the Constrained policy. For generation zero and beyond these constraints are removed, and they live under the Unconstrained policy. The timing is as follows: measures for generation zero in the figure refer to parents who lived under the old policy and children who lived under the new one.

The transition takes place rapidly. The adjustment of the IGE is monotone, and has mostly taken place by the time the children have gone to college under the new policy (even though it didn't apply to their parents yet). Going by components, almost all of the effect dissipates by generation 1, meaning the adjustment towards steady state has mostly taken place after

Figure 9: Transition from Constrained to Unconstrained case

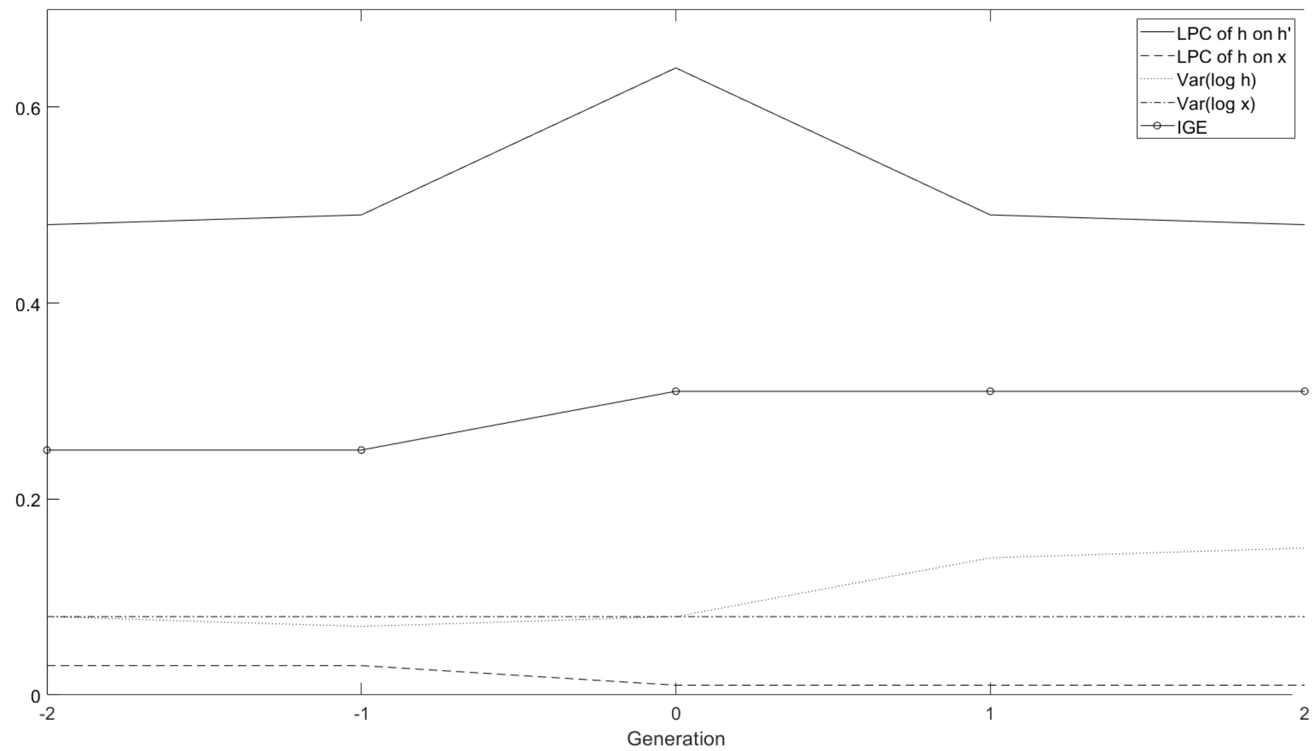

both parents and children have gone to college under the new policy.

For generation zero, the LPC of parental human capital on childrens' human capital ($\frac{Cov(\log h, \log h')}{Var(\log h)}$) spikes. Their correlation stays roughly the same,[29] but the variance of log children's human capital rises immediately which in turn increases the variance. But the variance of log parental human capital only follows one period later. As a result, the IGE increases. For the next generation, the variance of parental human capital increases as the new generation of parents has lived under the unconstrained policy. The LPC drops back down, but it now also receives extra weight in the calculation of the IGE. As a result, the IGE does not fall back down.

6.4 Robustness

Idiosyncratic shocks are a key component to the results presented above. Throughout, it is assumed that all individuals are subject to the same idiosyncratic shocks. Crucially, if education also increases the 'luck' component in earnings, then our model would overestimate the positive effect education policies have on persistence. A review of the literature on earnings shocks suggests that this is not the case: more educated individuals are subject

[29] As we would expect, removing the constraints does cause this correlation to dip slightly for generation zero: high-ability children from low income families look less like their parents. But quantitatively, that effect is small versus the effect on the variance of log human capital.

to similar or even somewhat smaller idiosyncratic income risks than are less educated individuals. For example, Meghir and Pistaferri (2004) report that the variance of unexplained earnings growth in their setup falls in education. Abbott, Gallipoli, Meghir, and Violante (2013) provide further evidence and find little difference over education groups for persistent shocks and a variance of non-persistent shocks that is smaller for the more educated. Finally, the findings of Blundell, Graber, and Mogstad (2015) for Norwegian administrative data do not contradict these conclusions. In short, the key modeling assumption on earnings shocks in this paper is conservative with respect to the main result.

7 Conclusion

Using a combination of theory and data, this paper has attempted to explore the connection between higher education policies and IM. It has shown that a human capital-based model can explain the persistence of earnings across generations, as well as its relation to college choice. We now know that the relationship between higher education policies and IM is theoretically ambiguous. Going one step further, the parameterized model of this paper suggests that common higher education policies actually decrease IM.

This surprising finding is due to the fact that earnings do not just originate from human capital, but are at least also due to luck. Education policies that increase the mobility of human capital also decrease the importance of luck, thereby making earnings more persistent over generations overall. Policymakers should be careful not to interpret lower levels of IM as signs of inefficiency.

Several directions may be worth pursuing in further research. The paper introduces a model of colleges in a competitive landscape that other researchers may find useful. The topic of college heterogeneity, which this work has shown to be important in a number of respects, may generally warrant further exploration. In particular, the findings in this paper suggest that college heterogeneity is important for the study of financial constraints: even when all potential students can afford to go to *some* college (the *extensive* margin of college choice), they may not be able to enroll in the college that is best for them (the *intensive* margin of college choice). But most of the empirical literature on financial constraints looks at whether students are constrained from entering some college (see Lochner and Monge-Naranjo (2011a) for an overview), or are constrained when already in college (e.g. Stinebrickner and Stinebrickner, 2008). To what extent are students constrained in their choice between colleges? Further research may provide an answer to this question.

A Appendix

A.1 The Empirics of Colleges and Intergenerational Mobility

Recent empirical work suggests that higher education is key to understanding the causes of IM. Chetty, Friedman, Saez, Turner, and Yagan (2017) find that there is a strong correlation between parental earnings and child earnings for the United States as a whole (a rank-rank correlation of 0.288). But this correlation is much smaller when including college fixed effects (0.100), i.e., a child's earnings are almost unrelated to that of their parent once we know what college the child goes to. This is found to be true irrespective of the type of college. Figure 10 illustrates the analysis.[30]

Chetty et al. also document IM by type of college. Specifically, they calculate the proportion of students in each college that comes from a low-income background, defined as a family income in the bottom 20%. This is then seen as a measure of *access* to the college. Next, they find the proportion of students from a low-income family that reach the top 20% of the income distribution, which they take as a measure of a college's *success* in generating mobility. Finally, the product of these two is the share of students in a college that comes from a bottom 20% family and reaches the top 20%. This is a measure of *mobility.*

Figure 1 displays these measures per college, first by a measure of average student ability on the left, and then by a measure of educational investment on the right.[31] In order to emphasize the absolute value of the probability measures, all graphs have a common vertical axis. Access decreases as a college enrolls students with higher average SAT scores, or as it spends more on instructional expenditures. At the same time, the success of those students that do attend increases. As a result, mobility across colleges is remarkably flat and low. In short, college heterogeneity appears to be an important part of the story.

[30]Chetty et al. 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 the longest counts as the college that the student attended. Their preferred measure of intergenerational persistence is rank-rank regressions of parental and child earnings. (Both parents and children are assigned a rank within their own distribution. Then, the child rank is regressed on parental rank. This procedure is the same as the Spearman correlation, but additionally allows for the inclusion of controls, such as college attended.)

[31]Data from Chetty et al. (2017). Only 4-year private colleges are included, and those who report instructional expenditures per student above $40,000 are excluded as outliers.

The findings by Chetty et al. add to an established literature that suggests human capital is the main culprit in IM, and that once we understand educational outcomes, we will largely understand intergenerational earnings persistence. These findings also suggest that higher education policies may have an important role to play: policies help low-income students pay for college, and college expenditures seem to propel them upward in the earnings distribution.

Figure 10: Controlling for Colleges

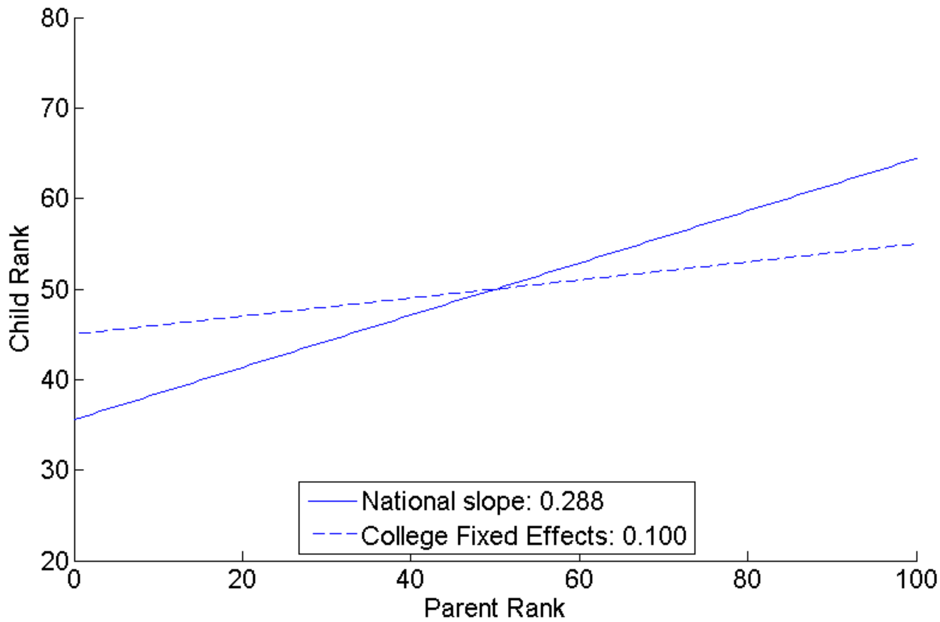

Lines are created by fitting a straight line with slope equal to the IGE estimate through the point (50,50).

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

Table 14 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. In order 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). Because of their limited size, as well as their limited relevance to those students that are likely to face financial constraints in the absence of any loans, this paper does not model private loans.

Default on student loans is not part of this paper. 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.

A.3 Solution and Computation

A.3.1 Solution Method

Given the assumptions underlying the above definition of stationary equilibrium, prices have analytical solutions. (The institutional grant component of college pricing schedules is exogenously given.) This leaves the individual's problem to be solved.

The individual problem is a simple life-cycle problem that can be solved by backward induction. At the same time, generations are linked through imperfect altruism. This complicates matters, but not by much: simple rewriting of the problem yields a single recursive equation, which is a relatively standard problem in macroeconomics.

$$V(\alpha, q, a_0) = \max_{\substack{\left\{k(s^0), d(s^0)\right\}_{s^{t0}}, \\ \left\{\left\{\mathbf{z}(s^t)\right\}_{s^t}\right\}_{t=j}^{t^I-1}, \left\{\left\{\mathbf{z}(s^t)\right\}_{s^t}\right\}_{t=t^I+1}^{T-1}, \\ \left\{\left\{\mathbf{z}(s^{t^I}, \alpha')\right\}_{s^t}\right\}_{\alpha'}, \left\{v(s^{t^I})\right\}_{s^{t^I}}}} \mathbb{E}\left\{\sum_{t=0}^{T-1} \beta^t \frac{(c_t^{\nu} l_t^{1-\nu})^{1-\sigma}}{1-\sigma} + \omega \beta^{t^I} V(\alpha', q', v)\right\}$$

Constraints and transitions are suppressed in the above for parsimony, but are unchanged except that they now depend on $k(s^0)$. The problem can be solved by iterating on an initial guess of V.

The recursive structure combined with individual life cycles increases computational demands, but when solving the problem by iteration on an initial guess the additional burden is reduced by the possibility of introducing *Howard Improvement* steps: one does not need to redo maximization on every iteration, which saves time when the maximization problem is 'large', as is the case here.

A.3.2 Computational Procedure

The computational procedure, for given parameter values, we adopt in this paper proceeds as follows.

1. Guess the initial value function at independence V (because we interpolate between grid points, guesses at grid points are sufficient).

2. Solve the individual's problem. I elaborate on this below. This results in an updated function V.

3. Update V (i.e. repeat from step 1) until convergence.

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.

The individual problem is solved by backward induction. Because each life-cycle starts with a discrete choice, value functions will have kinks and (through inter-generational links that reach back many generations) so-called 'second-order kinks'. Thus, inter-temporal first-order conditions do not apply. Instead, the optimization to solve individual problems is done over time using robust multi-level grid methods at each step of the life cycle. Leisure is assumed to always be interior, deviations from which are treated as a numerical imprecision.[32]

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

[32] The fraction of individuals that chooses zero work time is tiny, which is due to the absence of bequests in the model.

A.4 Further Model Implications

Further to Section 5, the below describes model implications and compares them to their empirical counterparts.

A.4.1 Paying for College

Inter-vivos transfers are an important source of college financing. Gale and Scholz (1994) discuss how these transfers are distributed empirically. The model distribution of these transfers, displayed in Figure 11 (where transfers are normalized by average labor earnings in a model period), resembles the data: it is heavily right-skewed, with a mass point at zero, and transfers are generally larger for those that go to college than for those that do not. Substantial grants also go to those who do not go to college: this is due to the model structure, in which there is only one moment for parents to act on their altruistic preferences.

Figure 11: Distributions of Inter-Vivos Transfers

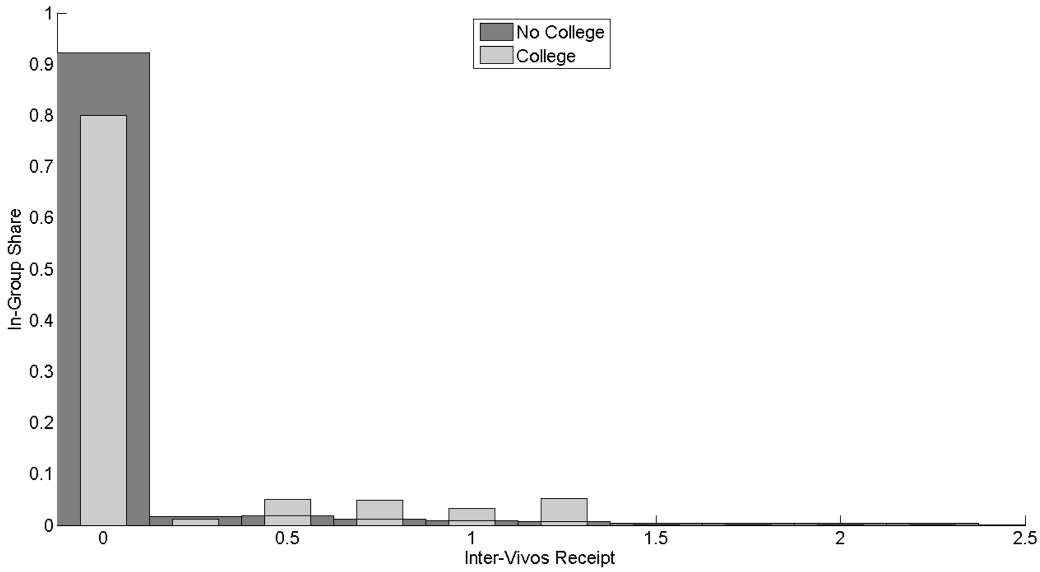

Grants and subsidies have been discussed above. Student loans are the remaining form of student aid. The model's parameters that regulate Stafford loans were set to match overall subsidized and unsubsidized loan amounts, as well as their relative limits. The model also makes predictions on the fraction of students that take up loans. The fraction of students with subsidized Stafford loans was 37.3% in 2000, and the same fraction for unsubsidized Stafford loans was 21.2% (Abbott et al., 2013). These moments and their model equivalent are displayed in Table 15. Clearly, the model does not do a good job at matching loan uptake at the extensive margin. As a result, the model has too many students taking up loans, and too little loans per student, compared to the data. This would be a result of missing heterogeneity of eligibility, for example, because eligibility, in reality, is tied to additional conditions, such as actual college expenses. Given the complexity of modeling such conditions, I have focused on matching overall loan availability at the expense of generating a realistic cross-sectional pattern of loan uptake. A similar comment applies to the modeling of inter-vivos transfers: the model presumably misses some sources of heterogeneity here too, for example in preferences, that would make financing needs more heterogeneous.

Table 15: Remaining Moments

	Model	Data
Share of students with Subsidized Stafford Loans	75%	37%
Share of students with Unsubsidized Stafford Loans	98%	21%
Public spending on education as % of GDP	0.77%	1.07%
Government expenditures as % of GDP	20.8%	37%
Weekly hours worked in college	33	>12
Weekly hours studied in college	15	<25
Frisch elasticity	1.20	>0.75
Wage premium	1.61	1.61

The overall role of the government in this model is in line with the data. First, public spending on education, which in our model includes institutional grants, is about 1% to 1.5% of GDP depending on sources.[33] Second, government expenditures amounted to 36.6% of GDP in 2003 according to the OECD. The model economy includes numerous sources of taxation, but not all so that it underestimates the size of the government somewhat. Model counterparts to both figures are reported in Table 15.

The final source of college financing is time use: students can choose how much time to spend working instead of studying or enjoying leisure. The ratio of time spent on education versus at work has already been targeted. As mentioned, quality data points on time use by students are hard to come by. Comparing the model to the data on this issue is difficult for a further reason: in the model, students are identified by enrollment, and therefore in principle include drop-outs, part-time students, students in two-year programs, and so forth, for an entire four year period. In spite of these issues, I produce model predictions of time use levels in Table 15. They are in line with sources other than those already reported from: Data from the National Center for Education Statistics (2017) lead to an estimate of 12.24

[33]For example, NIPA reports public expenditures of 0.91% of GDP for 2003, while institutional grants amounted to .16% of GDP (making for a total of 1.07%).

hours worked per week for a full-time student (a lower bound for the model). According to the calculations by the Bureau of Labor Statistics based on data from the Amerian Time Use Survey[34], an average full-time student spends 3.5 hours a day (or 24.5 hours a week) on educational activities (an upper bound for the model).

A.4.2 Labor Earnings

Figure 12 displays the model distribution of earnings, with average earnings normalized to one. The model distribution has much in common with the well-known empirical distribution, specifically that it is right-skewed and has a long right tail. Unreported results show a distribution of wealth with a significant population of indebted agents, and a large mass of agents with close to zero asset holdings. However, the wealth distribution does not produce the large right tail that is observed in data. This is due to the fact that the model does not include inheritances, nor inter-vivos transfers at ages other than the start of adult life. It also does not include a retirement period. Age-savings profiles also reflect the model structure: assets gradually deplete after an initial receipt of parental transfers, and the average agent starts building up savings again towards age 40. That build-up of assets is temporarily interrupted by inter-vivos transfers to children.

Figure 12: Distribution of Labor Earnings

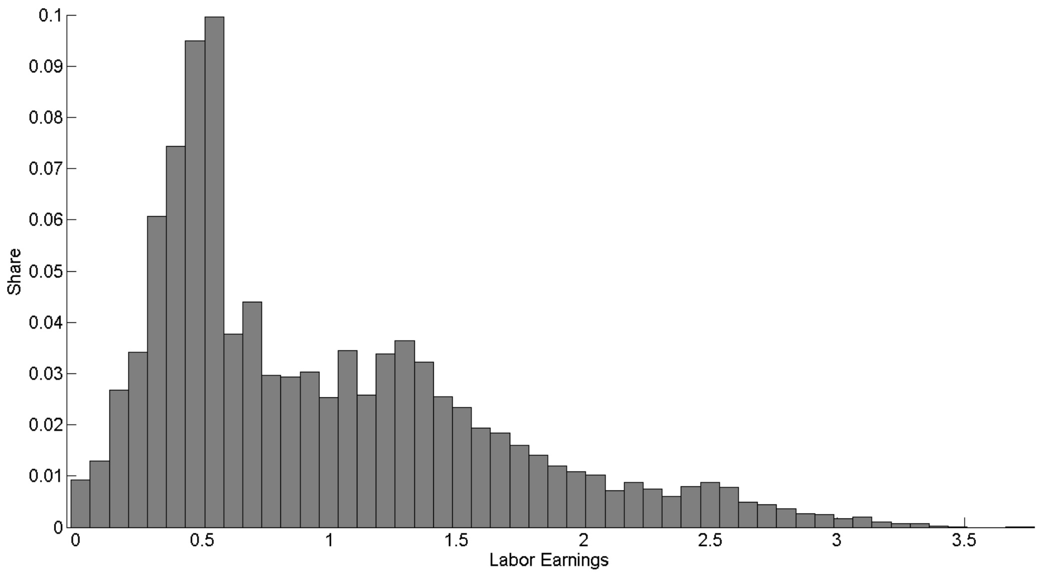

Frisch elasticities are typically used to measure the responsiveness of labor supply in macro

[34]`https://www.bls.gov/tus/charts/students.htm`, data are from the period 2011-2015.

models. Chetty, Guren, Manoli, and Weber (2011) argues for Frisch elasticities of 0.75 in macro models, but Keane and Rogerson (2012) argue that values well over one are more in line with the data once human capital is taken into account. The model's implied average Frisch elasticity is in that region, see Table 15.

Finally, the raw college premium at age 32, the ratio between the average wages of those with 16 years of education or more over the average wages of those with less, is 1.61. This is taken from a sample of the 2000 US Census, obtained from IPUMS.[35] The model generates a figure that is in line with this measure (see Table 15).

A.4.3 Controlling for colleges and parental income

Chetty et al. (2017) discuss how the IGE changes when one controls for the college a student enters. In doing so, they reduce the sample to those who enter any college. Intuitively, if college choice is a perfect measure of human capital, and human capital is all there is to earnings, one would expect the IGE to be reduced to zero. If financial constraints lead able students to enter worse colleges, then in a given college, disadvantaged children might even out-earn others, making the IGE negative. Chetty et al. report a national IGE (based on rank-rank regressions) of 0.29, which is reduced by two-thirds to 0.1 when including college fixed effects. (See Figure 10 and the discussion there.) Should this result be taken to indicate that human capital cannot explain all of the earnings persistence? This paper's model suggests another explanation.

Figure 13 illustrates a similar procedure, but on model-generated data. College choices in the model are very granular. So instead of using college fixed effects, educational investment is included as a control. Again, one would expect to find a flat or even declining line, since, in the model, all persistence is due to human capital. Interestingly, the coefficient also remains positive in the case with controls, just as in the data.

Key to understanding this is the following: college spending and ability are not perfectly assortative, even in the absence of constraints. While college spending is more effective for the more able, their optimal level of investment is nevertheless not necessarily higher. This is because time and money can be transformed through wages in the model so that the optimal investment level also depends on demands for leisure time. Substitution effects dominate income effects in the model overall so that the more able people earn and work more later in life. At the same time, it is optimal for them to enjoy more leisure relatively early in life, as their expected wages grow steeply. These increased demands for leisure in college can undo

[35] The bottom 1% of incomes are dropped, as are those who work less than 40 weeks a year or less than 35 hours a week. The Census' own top-coding corrections are accepted as is.

the higher effectiveness of spending. In short, smarter students sometimes study less and enjoy more leisure, even when they can afford to go to a more expensive college. As a result of the imperfect assortativeness between ability and investment, controlling for colleges does not actually control for human capital. The initial premise that one should expect the slope to be zero when human capital alone explains persistence is false: it can be positive, even when constraints are present.

Figure 13: IGEs with and without controls

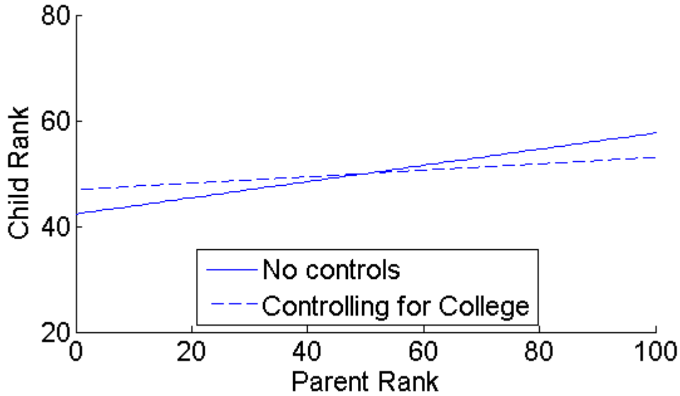

Graphs are created by fitting a straight line with slope equal to the IGE estimate through the point (50,50).

Landersø and Heckman (2017) estimate the IGE non-linearly, and find that persistence is larger for those with higher-income parents. A simple quadratic regression on model-generated log earnings indeed produces a convex relationship between the earnings of two generations: earnings persistence is stronger when parents have high earnings. Such findings are entirely in line with one of the main messages of this paper: at the top of the earnings distribution, human capital is more important than other components of earnings. At the same time, human capital itself is quite persistent. Thus, we quite naturally find higher persistence at the top of the distribution.

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.

Becker, G. S. and N. Tomes (1979). An equilibrium theory of the distribution of income and intergenerational mobility. *Journal of Political Economy 87*(6), 1153–1189.

Belley, P. and L. Lochner (2007). The changing role of family income and ability in determining educational achievement. *Journal of Human capital 1*(1), 37–89.

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

Blundell, R., M. Graber, and M. Mogstad (2015). Labor income dynamics and the insurance from taxes, transfers, and the family. *Journal of Public Economics 127*, 58–73.

Bovenberg, A. L. and B. Jacobs (2005). Redistribution and education subsidies are Siamese twins. *Journal of Public Economics 89*(11), 2005–2035.

Cai, Z. and J. Heathcote (2018). College tuition and income inequality.

Caucutt, E. M. and L. Lochner (2012). Early and late human capital investments, borrowing constraints, and the family. Technical report, National Bureau of Economic Research.

Caucutt, E. M., L. Lochner, and Y. Park (2017). Correlation, consumption, confusion, or constraints: Why do poor children perform so poorly? *The Scandinavian Journal of Economics 119*(1), 102–147.

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

Chetty, R., A. Guren, D. Manoli, and A. Weber (2011). Are micro and macro labor supply elasticities consistent? A review of evidence on the intensive and extensive margins. *The American Economic Review 101*(3), 471–475.

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.

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.

Cunha, F. and J. Heckman (2007). The technology of skill formation. *American Economic Review 97*(2), 31–47.

Cunha, F., J. J. Heckman, L. Lochner, and D. V. Masterov (2006). Interpreting the evidence on life cycle skill formation. *Handbook of the Economics of Education 1*, 697–812.

Dale, S. B. and A. B. Krueger (2002). Estimating the payoff to attending a more selective college: An application of selection on observables and unobservables. *The Quarterly Journal of Economics 117*(4), 1491–1527.

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.

Fernandez, R. and R. Rogerson (1995). On the political economy of education subsidies. *The Review of Economic Studies 62*(2), 249–262.

Fernandez, R. and R. Rogerson (1996). Income distribution, communities, and the quality of public education. *The Quarterly Journal of Economics 111*(1), 135–164.

Fernández, R. and R. Rogerson (1998). Public education and income distribution: A dynamic quantitative evaluation of education-finance reform. *American Economic Review 88*(4), 813–833.

Fernandez, R. and R. Rogerson (2003). Equity and resources: An analysis of education finance systems. *Journal of Political Economy 111*(4), 858–897.

FinAid (2016, November). Historical loan limits.

FinAid (2018, February). Historical interest rates.

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

Findeisen, S. and D. Sachs (2016b). Optimal need-based financial aid.

Fuller, M. B. (2014). A history of financial aid to students. *Journal of Student Financial Aid 44*(1), 42–68.

Gale, W. G. and J. K. Scholz (1994). Intergenerational Transfers and the Accumulation of Wealth. *The Journal of Economic Perspectives 8*(4), pp. 145–160.

Grawe, N. D. and C. B. Mulligan (2002). Economic interpretations of intergenerational correlations. *Journal of Economic Perspectives 16*(3), 45–58.

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.

Han, S. and C. B. Mulligan (2001). Human capital, heterogeneity and estimated degrees of intergenerational mobility. *The Economic Journal 111*(470), 207–243.

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.

Havránek, T. (2015). Measuring intertemporal substitution: The importance of method choices and selective reporting. *Journal of the European Economic Association 13*(6), 1180–1204.

Heathcote, J., K. Storesletten, and G. L. Violante (2017). Optimal tax progressivity: An analytical framework*. *The Quarterly Journal of Economics 132*(4), 1693–1754.

Heckman, J. J., L. Lochner, and C. Taber (1998a). Explaining rising wage inequality: Explorations with a dynamic general equilibrium model of labor earnings with heterogeneous agents. *Review of economic dynamics 1*(1), 1–58.

Heckman, J. J., L. Lochner, and C. Taber (1998b). General-equilibrium treatment effects: A study of tuition policy. *The American Economic Review 88*(2), 381–386.

Hendricks, L. and O. Leukhina (2017). How risky is college investment? *Review of Economic Dynamics 26*, 140–163.

Herrington, C. M. (2015). Public education financing, earnings inequality, and intergenerational mobility. *Review of Economic Dynamics 18*(4), 822–842.

Holmlund, H., M. Lindahl, and E. Plug (2011). The causal effect of parents' schooling on children's schooling: A comparison of estimation methods. *Journal of Economic Literature 49*(3), 615–51.

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. (2016a). The dramatic economics of the u.s. market for higher education. *NBER Reporter 3*.

Hoxby, C. M. (2016b). The productivity of u.s. postsecondary institutions: Implications for markets and policies. The Fellows Lecture, The 2016 Annual Meeting of the Society of Labor Economists.

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.

Jacobs, B. and A. L. Bovenberg (2011). Optimal taxation of human capital and the earnings function. *Journal of Public Economic Theory 13*(6), 957–971.

Johnson, N. (2014). College costs and prices: Some key facts for policymakers. *Lumina Issue Papers*.

Keane, M. and R. Rogerson (2012). Micro and macro labor supply elasticities: A reassessment of conventional wisdom. *Journal of Economic Literature 50*(2), 464–76.

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.

Meghir, C. and L. Pistaferri (2004). Income variance dynamics and heterogeneity. *Econometrica 72*(1), 1–32.

Mulligan, C. B. (1999). Galton versus the human capital approach to inheritance. *Journal of Political Economy 107*(S6), S184–S224.

National Center for Education Statistics (2017). The condition of education.

Restuccia, D. and C. Urrutia (2004). Intergenerational persistence of earnings: The role of early and college education. *American Economic Review*, 1354–1378.

Stantcheva, S. (2017). Optimal taxation and human capital policies over the life cycle. *Journal of Political Economy 125*(6), 1931–199.

Stinebrickner, R. and T. Stinebrickner (2008). The effect of credit constraints on the college drop-out decision: A direct approach using a new panel study. *The American Economic Review 98*(5), 2163–2184.

Storesletten, K., C. I. Telmer, and A. Yaron (2004). Cyclical dynamics in idiosyncratic labor market risk. *Journal of Political Economy 112*(3), 695–717.