



Institute for Empirical Research in Economics
University of Zurich

Working Paper Series
ISSN 1424-0459

Working Paper No. 402

**Accounting for the Changing Role of Family Income
in Determining College Entry**

Christoph Winter

Revised version, December 2011

Accounting for the Changing Role of Family Income in Determining College Entry*

Christoph Winter[†]

September 15, 2011

Abstract

In recent decades, the US has experienced a widening of the college enrolment gap between rich and poor families. This is commonly interpreted as evidence for a tightening of borrowing constraints. This paper asks whether this is indeed the case.

I present an incomplete-markets overlapping-generations model with college enrolment, in which altruistic parents provide transfers to their children. In the model the rise in earnings inequality observed between 1980 and 2000 acts as the driving force for generating the trends in the data. With the help of counterfactual experiments, I find that fraction of constrained households is much higher (24 instead of 8 percent) than indicated by the narrow enrolment gap in 1980. Contrary to what the development of the enrolment gap in the data suggests, the share of constrained households actually fell (to 18 percent) between 1980 and 2000. I show that altruism is important for explaining these findings.

Keywords: Dynamic General Equilibrium Models with Overlapping Generations, Parental Transfers, College Enrolment and Borrowing Constraints, Earnings Inequality
JEL classification: D11, D31, D58, D91, I2

*This is a substantially revised version of the first chapter of my Ph.D. dissertation, written at the European University Institute. This version supersedes EUI working paper 2007/49 and IEW working paper Nr. 402. I would like to thank my advisor Morten Ravn and my thesis committee, consisting of Alex Michaelides, Claudio Michelacci and Salvador Ortigueira, for comments on earlier drafts. I have also received many useful comments from Ken Judd, Dirk Krueger, John Knowles, Ramon Marimon, Costas Meghir, Franck Portier, Morten Ravn, Víctor Ríos-Rull, David Scherrer, Kjetil Storesletten, Mariya Teteryatnikova, Gianluca Violante, Ken Wolpin and Fabrizio Zilibotti as well as from seminar participants at various institutions. The hospitality of UPenn economics is gratefully acknowledged, where part of this paper was written. All errors are mine.

[†]Address: Department of Economics, University of Zurich, Mühlebachstr. 86, CH-8008 Zurich. E-Mail: christoph.winter@econ.uzh.ch

1 Introduction

Over the last three decades, the US has experienced a widening of the college enrolment gap between rich and poor families. This was documented by Ellwood and Kane (2000) and Belley and Lochner (2007). In this paper, I study the driving forces behind this development. I am particularly interested in understanding whether an increase in enrolment gaps implies that borrowing constraints have become binding for a larger fraction of population.

Answering this question is relevant for wide range of policy issues, like the design of college subsidies, or questions concerning social mobility. Following the seminal work of Becker and Tomes (1979), many papers have analyzed whether constraints are binding, with mixed evidence.¹ A popular approach of gauging the importance of borrowing constraints is to compute the college enrolment gaps between children of rich parents and children of poor parents (e.g. Ellwood and Kane (2000), Carneiro and Heckman (2002), Kane (2006)). The implicit assumption made here is that children from rich families are not financially constrained in their college choice because they receive enough support from their parents while children from poorer parents do not have access to such funds. Indeed, parental transfers constitute an important source of college financing in the data (see e.g. Gale and Scholz (1994) and Keane and Wolpin (2001)). Accordingly, a stronger impact of family income on college attendance, which has been observed for the US, suggests that borrowing constraints play a more prominent role (Belley and Lochner (2007), Lochner and Monge-Naranjo (2010)).

In this paper, I develop a quantitative theory of college enrolment which is designed to shed light on the complex interaction between family resources, parental transfers, children's academic ability and borrowing constraints underlying the enrolment pattern in the data. The model features incomplete markets and idiosyncratic productivity shocks, as in the seminal contribution of Aiyagari (1994). Households go through a life cycle of working and retirement, as in Huggett (1996).²

The key innovation of my paper with respect to the previous empirical literature is the structural approach of modeling parental transfers. This approach allows me to trace out the behavioral responses of parental transfers with respect to changes in the economic environment, e.g. changes in the return or in the cost of education, which the US economy has experienced recently (see e.g. Krueger and Perri (2006), Heathcote et al. (2010) and Collegeboard (2005)). My model can also be used to assess whether the enrolment gaps observed in the data are informative about binding borrowing constraints. This is an important contribution to an ongoing debate in the literature (Ellwood and Kane (2000), Carneiro and Heckman (2002), Kane (2006), and more recently Brown et al. (2011)).

I show that the fraction of households that are borrowing constrained in their college decision is much higher than the average enrolment gap. The average enrolment gap in the model is just below 7 percent and thus in line with the data from the beginning of the

¹Kane (2006) and Brown et al. (2011) contain recent overviews. Kane (2006, p.1396) concludes that "[...] it is difficult to find a definite test of the existence of borrowing constraints in the literature".

²Huggett and Ventura (2000), Storesletten et al. (2004a), among others, have shown that models of this class are consistent with many empirical facts of savings and consumption over the life cycle.

1980s (Carneiro and Heckman (2002)). The fraction of borrowing constrained households is 24 percent and thus more than three times as high. Perhaps surprisingly, my model predicts that the fraction of constrained households declined over time, from 24 percent at the beginning of the 1980s to 18 percent at the beginning of the 2000s. This fall occurred in the model despite the fact that it is - broadly - consistent with the increase in the enrolment gap observed between 1980 and 2000 (Belley and Lochner (2007)).

In order to understand this apparent contradiction, it is important to consider the different college financing opportunities in the model first. I assume that college education is financed with the help of parental transfers and loans.³ When modeling parental transfers, I assume that parental transfers occur because parents are altruistic towards their offspring, following the seminal work of Barro (1974) and Becker (1974). Parental altruism can be imperfect, as in Laitner (2001).

With respect to loans, Keane and Wolpin (2001) find that borrowing opportunities for college financing are limited in the US. I model this constraint by assuming that the amount households can borrow against their own future earnings is restricted. This is a standard life cycle liquidity constraint, see e.g. Cunha and Heckman (2007). An additional constraint prevents parents from borrowing against their children's future income to finance investments in them. This is the intergenerational borrowing constraint, see Brown et al. (2011). Because altruism is only one-sided and intergenerational borrowing is prohibited, parents face a trade-off between saving for their own future consumption and providing transfers to their children. Parents may find it optimal to underinvest in their children's college education, in which case the intergenerational borrowing constraint is binding. Whether lifetime liquidity constraints prevent young households from attending college thus depends on the willingness of young households to invest in college as well as on their parents' willingness to provide support.

In my model, the decision of children to invest in college is a function of the cost and the (expected) return, which in turn depend on academic ability. Academic ability is defined as the set of all skills relevant for college education. It is partly transmitted from parents to children. Academic ability influences the probability of dropping out from college before graduating. More able students are more likely to complete their education (Chatterjee and Ionescu (2010)). Only college graduates receive an educational premium on their earnings. More able students also need to pay less for tuition, as documented by McPherson and Schapiro (2006). Consistent with the US system of college financing, I also assume that the direct cost of tuition a student faces are negatively correlated with family resources: students from rich families pay more tuition than students from poor families.

The extent to which the willingness of children to attend college exceeds their parents' willingness to provide support also hinges on specific parameter values. I thus choose the key model parameters, such as the discount factor, the tightness of the borrowing constraint, the degree of intergenerational discounting, the transmission of ability across generations, and the cost of college education by calibrating the model separately to moments from US

³According to Keane and Wolpin (2001), students finance between 20 and 60 percent of their college expenses by parental transfers. In total, parental support for their children's college education is substantial. Gale and Scholz (1994) document that parental payments amounted to 35 billion dollar in 1986.

economy at the beginning of the 1980s (the benchmark calibration) and the 2000s.

The results that follow from the benchmark calibration imply that the fraction of households restricted in their college decision by the presence of lifetime liquidity constraints is much higher than indicated by the average enrolment gap. This suggests that the enrolment gap is not a good measure for the actual fraction of constrained households. There are two reasons that lead to this conclusion. The first is that the enrolment gap does not capture the fact that the intergenerational borrowing constraint may be binding even if parents are rich. Because altruism is one-sided and imperfect, even parents that are well-off may face a trade-off between providing transfers and own savings. Also not captured by the average enrolment gap is the fact that the cost of attending college are lower for children of poor families, all other things equal. This is because this group receives subsidies. In a world without borrowing constraints, children from poor families are thus more likely to enrol in college compared to children from rich families, all other things equal. I show that these two mechanisms are able to explain the stark difference between the results from the previous literature which are based on the average enrolment gap and the actual share of constrained households in my model.

In order to mimic the situation at the beginning of the 2000s, I raise earnings inequality and tuition fees. The rise of the college premium makes attending college more attractive. This is partly outweighed by the increase in tuition fees which makes attending college more expensive. Children from rich families can afford to pay the increased tuition. As a result, enrolment rates for children with rich parents rise. Since enrolment rates from children from poor families do not change substantially between 1980 and 2000, we observe an increase in the enrolment gap between children from rich and poor families. However, this does not mean that borrowing constraints became binding for a larger fraction of the population. Borrowing constraints became less binding for both children from rich and poor families. By definition, the group of the income-rich parents benefited the most from the increase in inequality, as their income increased with respect to the average. Therefore, these parents are now affluent enough to invest optimally in their children's education. Because of the prominent role the rise in earnings inequality plays for my results, my paper makes an important contribution to the recent literature that studies the macroeconomic consequences of rising wage inequality in the US (see Krueger and Perri (2006), Heathcote et al. (2010)).

Instead, children from poor families finance college mainly by borrowing. In my model, college dropouts cannot discharge their debt, consistent with the rules that are applied in the US (Chatterjee and Ionescu (2010)). That is, college loans need to be repaid, independently of whether college is completed successfully or not. Since the college wage premium is only earned by those who actually graduate, the repayment rules for college loans decrease the willingness of risk-averse households to borrow to finance their college education. In the context of my model, this implies that the fraction of borrowing constrained households from poor families decreased over time, since these households are not willing to accept the additional risk that would be associated with borrowing more in order to finance the increase in tuition fees.

The question whether education gaps are due to inefficiencies is of first-order importance

for public policy. My findings imply that borrowing constraints have a strong influence on college attendance but that the connection between observable enrolment gaps and binding borrowing constraints turns out to be rather loose. This suggests that enrolment gaps should be interpreted with caution if one is interested in detecting inefficient education outcomes.

Empirical support for the existence and the bindingness of the intergenerational borrowing constraint has been provided in an important recent piece of work by Brown et al. (2011). Both my model as well as theirs share the assumption that parental altruism is one-sided. Based on this assumption, they identify families which invest inefficiently in their children's college education. They find that the fraction of constrained families accounts for up to 50 percent of the relevant population. My paper complements the work of Brown et al. (2011) by using a quantitative model. This approach allows me to consider a rich life cycle model where earnings are uncertain, ability is transmitted between generations and the cost and return of education differ between students in various dimension. Considering a quantitative model also allows me to compute a wide set of moments and compare them to the data, such as the development of the enrolment gaps over the last decades.⁴ Importantly, both my findings as well as the results from Brown et al. (2011) point to the fact that borrowing constraints affect a substantially larger proportion of the population in their college decision compared to what the preceding literature has suggested.

An increasing number of papers implement altruism in computable life cycle models with endogenous education choice. Most recently, Gallipoli, Meghir and Violante (Gallipoli et al. (2010), henceforth GMV (2010)) propose an OLG model with one-sided altruism and sequential education choice.⁵ GMV (2010) introduce an aggregate production function where different types of human capital are not (necessarily) perfectly substitutable. They allow explicitly for changes in life cycle earnings and wealth profiles. When estimating the earnings process, they also distinguish between permanent ability and idiosyncratic labor shocks. For their estimation, GMV (2010) use data from the Panel Study of Income Dynamics (PSID), the Current Population Survey (CPS) and the National Survey of Young (NLSY). With the help of their model, the authors compute the effect of different policy interventions on optimal education decisions, inequality, and output.⁶ Compared to GMV (2010), my paper sheds light on the interaction between one-sided altruism and borrowing constraints as a determinant of college enrolment. Moreover, I show that parental transfers are important for explaining the changes in the relationship between family income and college enrolment that have occurred over time.

This paper is also related to a number of important papers in the literature that incorporate altruism to study education policies, such as Caucutt and Lochner (2004), Restuccia and Urrutia (2004), Cunha (2007), Bohacek and Kapicka (2010) as well as Holter (2011). Some

⁴Moreover, my approach allows me to consider various forms of financial college aid, such as need-based and merit-based aid, as well as the interaction between altruism and life cycle saving motives.

⁵Early papers analyzing the general equilibrium implications of education policies include Heckman et al. (1998) and Abrahám (2004), who examines wage inequality and education policy in a general equilibrium OLG model with skill biased technological change.

⁶Garriga and Keightley (2007) are also interested in optimal education policies. They do not endogenize parental transfers. Instead, they explicitly model the dropout decision and labor supply during college.

of these contributions focus on skill formation (as in Cunha (2007)) or on the relative importance of early versus late credit constraints, as in Caucutt and Lochner (2004). Bohacek and Kapicka (2010) study the welfare effects of educational reforms. Both Restuccia and Urrutia (2004) and Holter (2011) are interested in the determinants of earnings persistence. Cunha (2007) as well as Bohacek and Kapicka (2010) assume two-sided altruism: families care both about their predecessors and their descendants. Under this assumption, parents and children pool their resources and solve the same maximization problem. Two-sided altruism thus implies that children provide transfers to their parents as well. However, there is little evidence for this in the data, as argued by Gale and Scholz (1994) and Brown et al. (2011). Caucutt and Lochner (2004), Restuccia and Urrutia (2004) and Holter (2011) also assume one-sided altruism. Compared to these papers, I focus on explaining the development of enrolment gaps over time and on analyzing the impact of borrowing constraints.

Lochner and Monge-Naranjo (2010) argue that borrowing constraints became binding for a larger fraction of the population between 1980 and 2000. Lochner and Monge-Naranjo (2010) derive borrowing constraints endogenously from the design of government student loan programs and from limited repayment incentives in private lending markets. They use their setup to analyze the impact of an increase in the tuition fees and the college premium. They find that the minimum level of pre-college wealth that is required to guarantee efficient investment in college education rises significantly, despite the fact that an increase in the college premium raises the amount of credit that is provided by private lending markets. In this paper, I endogenize the initial wealth distribution by assuming parental altruism. This enables me to compare my model to the patterns of family income and college attendance that are observable in the data. Importantly, my results show that when the initial wealth distribution is generated by one-sided altruism, increasing earnings inequality and raising tuition actually reduces the fraction of constrained households over time.

The remainder of the paper is structured as follows. The model is presented in Section 2. Section 3 gives the household's problem in recursive notation, while Section 4 introduces the equilibrium definition. The calibration of the model's parameters is presented in Section 5. I discuss my results in Section 6. Finally, Section 7 concludes.

2 Model

Overview. I consider a life cycle economy with altruistic parents. Parents provide transfers to their children. Children take parental transfers as given, and decide about whether to attend college or not. Children can also borrow against their future earnings. All other credit markets are closed, in particular, parents are not allowed to borrow against the future income of their descendants. I allow for idiosyncratic productivity shocks during working life. These assumptions allow me to study the effects of an endogenously generated initial distribution of assets on college enrolment, and to analyze the determinants of the initial asset distribution in a realistic life cycle setting.

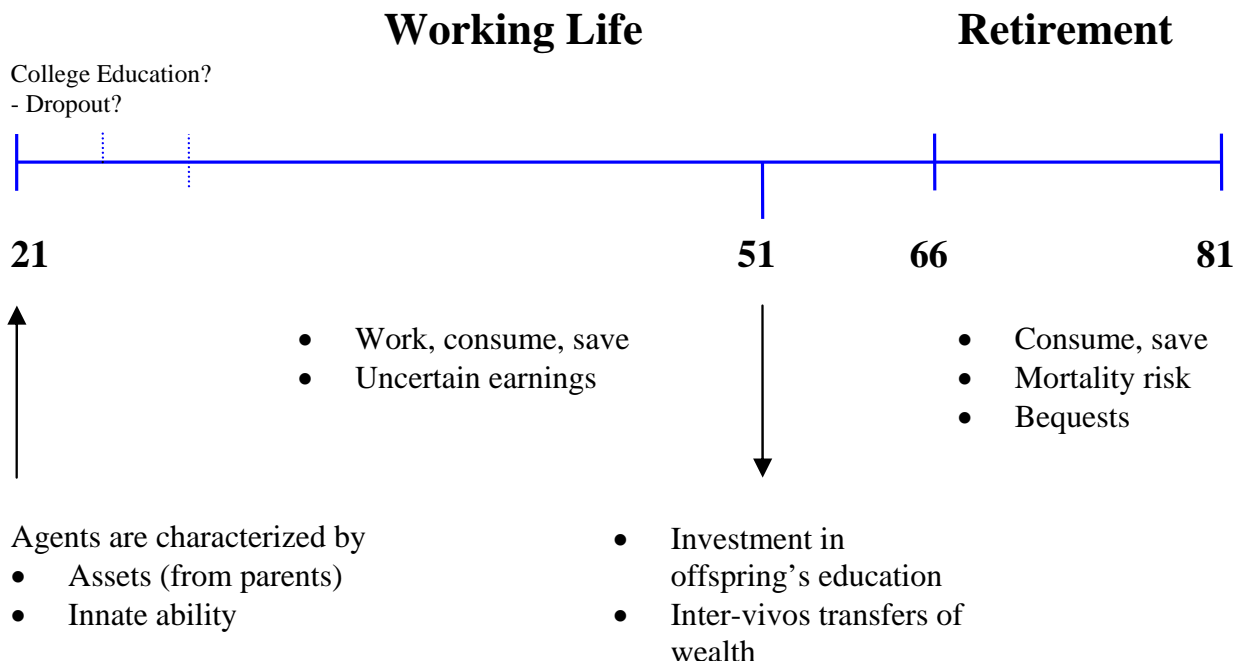


Figure 1: Life cycle and Generation Structure

The Life Cycle of a Household. There is a continuum of agents with total measure one. I assume that the size of the population is constant over time. Let j denote the age of an agent, $j \in J = \{1, 2, \dots, J^{\max}\}$. Agents enter the economy when they turn 21 (model period $j = 1$). Before this age, they belong to their parent household and depend on its economic decisions. During the first 45 years of their 'economic' life, agents work. This implies that the agents work up to age 65 (model period $J^{\text{work}} = 45$). Retirement takes place at the age of 66 ($j = 46$), which is mandatory. When agents turn 51 ($j = 31$), their children of age 21 form their own household. This implies a generational age gap of 30 years. It is assumed that there is one child household for each parent household. Agents face a declining survival probability after their children leave home. Terminal age is 81 ($J^{\max} = 60$). Since annuity markets are closed by assumption, agents may leave some wealth upon the event of death. The remaining wealth of a deceased parent household is passed on to its child household.

The assumptions regarding the life cycle and the transfer behavior are summarized in Figure 1.

Labor Income Process. During each of the 45 periods of their working life, agents supply one unit of labor inelastically.⁷ The productivity of this labor unit of an j -year old agent is measured by $\varepsilon_j^e \eta^{j,e}$, where $\{\varepsilon_j^e\}_{j=1}^{J^w}$ is a deterministic age profile of average labor

⁷College graduates work for fewer periods, see below.

productivity of an agent with education level e :⁸

$$e \in E = \{hs, col\}$$

where hs denotes high school education and col college education. For retired households and for students attending college, $\varepsilon_j^e = 0$.

$\eta^{j,e}$ describes the stochastic labor productivity status of a j -year old agent with education level e . Given the level of education e , I assume that the labor productivity process is identical and independent across agents (no aggregate productivity shocks) and that it follows a finite-state Markov process with stationary transition probabilities over time. More specifically,

$$Q(\eta^e, N^e) = \Pr(\eta^{j+1,e} \in N^e | \eta^{j,e} = \eta^e)$$

for high-school graduates. $N^e = \{\eta_1^e, \eta_2^e, \dots, \eta_n^e\}$ is the set of possible realizations of the productivity shock η^e .

College Investment, College Attendance and College Completion. Upon entering the economy, all households possess a high school degree.

I assume that the time it takes to complete a college degree is four years. In the US, as well as in other OECD countries, there is a significant fraction of students who enter college, but actually leave without having obtained a degree, see e.g. Restuccia and Urrutia (2004) or Akyol and Athreya (2005).

Upon entering college, students are required to own enough resources to finance their studies. Financial conditions thus do not matter for college completion, which is consistent with evidence provided by Stinebrickner and Stinebrickner (2007). In my model, students drop out because they fail to achieve the requirements that are necessary to obtain a degree. I assume that more able students are also more likely to graduate. Light and Strayer (2000) as well as Chatterjee and Ionescu (2010) report that there is a strong positive correlation between college completion and performance in scholastic tests. I denote the college dropout probability by λ .

In line with evidence provided by Stinebrickner and Stinebrickner (2007), I assume that the time until dropout is two years. Akyol and Athreya (2005) cite evidence for the fact that the return to the later years of college is substantially higher compared to the return to the first two years. Based on this finding, I assume that college dropouts face the same earnings process as high school graduates. As a consequence, college investment is indivisible.

In the following, I discuss the formation of academic ability, which is the underlying force of behind college success, in greater detail.

Ability and its Formation. I define academic ability as the set of skills (both cognitive and non-cognitive) that is relevant for success in college. Academic ability is thus denoted by the college success probability $(1 - \lambda)$. I assume that academic ability has no direct effect

⁸I do not consider high school dropouts. The share of high school dropouts is small in the data, see Rodriguez et al. (2002) who measure a share of 17 percent in the 1998 SCF.

on earnings, other than through education.⁹ Cawley et al. (2001) argue that it is hard to find an impact of cognitive and non-cognitive skills on earnings, after controlling for schooling attainment.¹⁰

The formation of skills is an active research area. So far, it appears to be consensus that the family "plays a powerful role in shaping abilities through genetics, parental investments and through choice of child environments", as Cunha and Heckman (2007, p.1) put it. In order to capture the role of genes, I assume that at the beginning of their life cycle, households are endowed with innate ability \mathbf{f} . Innate ability is partly transmitted from their parents, following a transition matrix Γ .

Following Keane and Wolpin (2001), I assume that other determinants such as parental investments in early education and parental choices of child environments can be approximated by the educational achievements of parents, in the following denoted by e^p . In the data, educational achievement is highly correlated across generations.¹¹ Moreover, Carneiro et al. (2007) find that an additional year of mother's schooling increases the child's performance on a standardized math test by almost 0.1 of a standard deviation.¹²

Parental education e^p and innate ability \mathbf{f} interact in determining the academic ability and thus also the probability of college success $(1 - \lambda)$, respectively. Evidence presented by Cunha and Heckman (2007) suggest that genes - innate ability in my context - and environmental factors (approximated by parental education) interact in multiple ways during childhood and adolescence in the formation of skills. Given these non-linearities, I do not impose any parametric structure on the process of skill formation. Instead, I will calibrate the values of λ that are associated with different combinations of \mathbf{f} and e^p endogenously.

As mentioned above, there is strong positive correlation between the performance in standardized test scores, such as the Armed Forces Qualification Test (AFQT) or the Scholastic Aptitude Test (SAT), and college completion rates.¹³ Since a single test score is unlikely to capture the full set of skills that are necessary for college success (Carneiro and Heckman (2002)), I distinguish between 'academic' ability and 'observable' ability in the following.¹⁴

Let \mathcal{K} denote a random variable with realization κ_a . κ_a denotes the performance of a young household in standardized test scores. Both κ_a and λ are known to children (and their

⁹Also see Heathcote et al. (2010)

¹⁰Cawley et al. (2001) show that the fraction of wage variance explained by measures of cognitive ability after controlling for human capital measures such as education and work experience is low. They argue the correlation between measured cognitive ability and schooling is so high that it is not possible to separate the two unless one is willing to make strong assumptions the parametric structure (e.g. log-linearity and separability), which Cawley et al. (2001) test and reject. The authors also provide evidence for the fact that non-cognitive skills (such as having self-discipline to follow the rules) also impact earnings mainly through schooling attainment.

¹¹See Black and Devereux (2010) for a comprehensive and recent review of the empirical literature

¹²With respect to the channels that transmit the effect of maternal education to the child, they find a substantial role played by income effects, delayed childbearing and assortative mating. Other potential channels that are mentioned in the literature are neighborhood effects, family stability and preferences for education (Haveman and Wolfe (1995))

¹³SAT scores and AFQT scores are highly correlated, see footnote 7 in Light and Strayer (2002).

¹⁴Moreover, as noted by Heckman et al. (2006), standardized tests are affected by a person's schooling and family background at the time tests are taken, which makes test scores a noisy measure for true ability.

parents) at the time the college decision is made, while only κ_a is observable by the public (e.g. colleges). I will also refer to κ_a as 'observable' ability. More specifically, I assume that $\mathcal{K} \sim \mathcal{N}(\mu_{\mathbf{f},e^p}, \sigma_{\mathbf{f},e^p})$, where the mean and the standard deviation depend on \mathbf{f} and e^p , to account for the empirical fact that test scores and college completion probabilities are correlated. Let $f(\mu_{\mathbf{f},e^p}, \sigma_{\mathbf{f},e^p})$ be the corresponding probability density function. In the next section, I discuss the relationship between observable ability and the cost of college.

In sum, my model of education and skill formation is consistent with the following empirical patterns: (i) a positive correlation between college attendance/completion and measures of ability, (ii) a positive correlation between college attendance/completion and parental education, (iii) a positive correlation between measures of ability and earnings.

Direct Cost of College Education. I assume that the cost of college are perfectly negatively correlated with observable ability. That is, the cost of college are given by $\kappa = -\kappa_a$. Two remarks are in order. First, since κ_a can be positive or negative, κ can be positive or negative as well. Negative cost of education should be interpreted as a stipend which covers part of the living expenses.¹⁵ Second, because of the tight relationship between observable ability and the cost of education, I will use the terms observable ability and cost of college as synonyms in the remainder of the paper.

A negative link between observable ability and tuition arises through the admission policy of US college and universities. These institutions compete for the best students (i.e. those with the highest test scores) with the help of financial aid packages (Linsenmeier et al. (2002)). As a result, better students need to pay less in order to enrol in college.¹⁶

Individual aid packages can be explicitly based on academic promise and achievement (so called "merit aid"). Only a small fraction of students are recipients of merit aid. Most students receive so called "need based aid", which is officially based on their "expected family contribution to pay" (further details follow below). Here, it is important to notice that most financial aid packages are implicitly based on student's merit, even if they are labeled differently.¹⁷

The following figures make this point clear. A 100 point difference in the SAT increases (need-based) grant aid between 500 and 2300 Dollars (in prices of 1996) over the course of a college career (McPherson and Schapiro (2006)). Private institutions typically have

¹⁵As I will outline in the calibration procedure (see Section 5) the share of students with negative cost of college turns out to be below 1 percent, which is in line with the fraction of students that receive generous grants.

¹⁶More talented students are more likely to generate positive peer effects, thus enhancing the education of all students. Attracting students with higher levels of measured quality is also important for an institution's reputation, see McPherson and Schapiro (2006).

¹⁷McPherson and Schapiro (2006) write: "Normal practice at American colleges is to present a prospective student with a "package" of aid, generally including some combination of federal, state and institutional grant, a recommended loan, and a work-study job. [...] By the same token, two students at the same college, both receiving only need-based aid, may receive quite different aid packages. The more desirable student may receive either a larger total aid package or a similar total aid package with a larger component of grant aid and lower amounts of loan and work. And this can happen without any of the dollars being labeled "merit" dollars. (McPherson and Schapiro (2006, pp. 1412-1413).

more resources at their disposal and thus offer more generous aid packages (Long (2004) and Collegeboard (2010)). These findings imply that a student from the lower end of the ability distribution has to pay between 3000 and 14000 Dollar more for his college education than a student at the upper end.¹⁸

The fact that the direct cost of college education are lower for young households with higher levels of observable ability generates heterogeneity with respect to the total cost of college education. All other things equal, more able students will be more likely to attend college, a pattern that is also observable in the data. Other dimensions of heterogeneity of the cost and return of attending college are generated by assuming that financial aid is higher for students from poorer families, which is discussed in the next section. The (expected) return of attending college differs because more able students are more likely to complete their education.¹⁹

Subsidies and Need-Based Financial Aid. Need-based financial assistance is based on the ability of the student's family to pay, the so-called 'expected family contribution'. It is calculated on the bases on parental income and wealth (Feldstein (1995)). A student with poor parents can expect to receive more financial assistance than a student whose parents are well off. I denote the fraction of direct college expenses (κ) that is covered by financial aid by ν . The net direct costs of college attendance, after financial help is taken into account, are given by $(1 - \nu)\kappa(I(\kappa) > 0)$, where $I(\kappa)$ is an indicator function taking the value 1 if college expenses are positive.²⁰ ν is a function of parental income and parental assets, $\nu(y^p, a^p)$. Because parental assets enter the calculation of the expected family contribution, there is an "education tax rate" on capital income, which has a powerful adverse effect on capital accumulation, according to Feldstein (1995). Two remarks regarding the expected family contribution are in order. First of all, as pointed out by Brown et al. (2011), the

¹⁸These figures are based on the assumption that students at the lower end of the distribution have a SAT scores below 700, while students at the upper end have scores above 1300. According to Chatterjee and Ionescu (2010), 8 percent of the students who took the SAT score below 700, while about 15 percent score higher than 1300.

¹⁹There are other ways to generate heterogeneity. In important contributions, GMV (2010) and Heathcote et al. (2010) assume that tuition is the same for everybody. Instead, they assume that college education is associated with utility costs ('psychic' costs), which are higher for the less able. Heckman et al. (2005) argue that the direct costs of college education are the sum of tuition and psychic costs. To what extend the direct costs are determined by tuition or by psychic costs remains an open question. The reason is that it is not possible to measure psychic costs directly in the data (Heckman et al. (2005), footnote 32). According to Heckman et al. (2005), psychic costs are needed to explain the observation that the take up rate is low, despite the fact that (ex-post) monetary return from attending college is high. Recently, Ozdagli and Trachter (2009) show that this so-called "returns to education puzzle" can be explained if one allows for risk-averse decision makers, who face the uncertainty about the outcome of college. Ozdagli and Trachter (2009) conclude that this provides "[...] a less controversial explanation for the once obscure psychic costs" (p. 25). Moreover, estimates provided by Dynarski (2003) suggest a quantitatively significant role of tuition costs in explaining college enrolment. She finds that increasing tuition subsidies by 1000 Dollar (in prices of 1998) raises college enrolment by 4 percentage points. In the appendix, I show that my modeling strategy is consistent with the empirical elasticities found by Dynarski (2003).

²⁰Recall that negative direct costs from college may occur if the measured ability of the student is very high and the student receives a scholarship.

expected family contribution that is calculated from the family’s ability to pay is not legally guaranteed. Put differently, there is no way children can actually force parents to give what they are expected to pay. Second of all, as noted by Dick and Edlin (1997), the financial aid that the student actually receives does not necessarily cover the difference between the expected family contribution and cost of college. Federal programs do not provide enough subsidized aid to meet the need of all students, and most colleges are not committed to cover the entire residual (Dick and Edlin (1997)).

Assets and Loans. Household accumulate savings in form of assets a_j . They borrow by taking out loans χ . Consequently, $a_j \geq 0$ and $\chi \geq 0$. Borrowing can only take place at the beginning of the life cycle, either to finance college education or for non-college related expenses or both.

Loans are assumed to be closed-end installment loans, which are characterized by fixed payments and a fixed term. This assumption is motivated by the fact that college loans are commonly closed-end loans, see e.g. GMV (2010). Moreover, installment loans are the most sizable component of unsecured consumer debt.²¹ Consequently, I abstract from ‘open-ended’ credit (such as credit card debt). The assumption that all borrowing takes place at the beginning reflects the findings that this is the stage of the life cycle where credit is most needed.²²

I assume that the terms and conditions for χ follow the ones for US college loans. In the following, I outline some institutional details regarding the market for college loans. The bulk of loans are provided under government sponsored loan programs (GSL), such as Perkins or Stafford.²³ All loan programs have in common that participants need to repay their loans in full, independently of whether they completed college successfully (Chatterjee and Ionescu (2010)). This implies that financing college by loans is associated with a substantial financial risk.

Participating students can defer loan payments until six (Stafford) or nine (Perkins) months after leaving school (Lochner and Monge-Naranjo (2010)).²⁴ Under the standard

²¹Using the 1983 Survey of Consumer Finances (SCF), Kennickell and Shack-Marquez (1992) report that 42 percent of all households whose head is younger than 35 years have taken out installment loans (other than car loans) in 1983. At the beginning of the 1980s, installment loans were thus the most important form of unsecured consumer debt. Installment loans appeared not only to be more frequent, but also to be more sizable. The median amount of installment debt carried by households was 1600 Dollars, whereas it was just 600 Dollars in form of credit card debt (in 1989 Dollars). More recent waves of the SCF confirm this pattern.

²²Kaplan and Violante (2010) compute that more than 40 percent of households borrow at the beginning of the life cycle. The desire to borrow vanishes as households grow older, and approaches zero around the age of 50. This corresponds roughly to the pattern that we observe in the SCF, see Kennickell et al. (2009). In my benchmark calibration, which I outline below, the fraction of households that have zero financial wealth following age 22 is only 4 percent.

²³At least until the mid-1990s, only a few private lenders offered student loans outside the government sponsored loan programs (Lochner and Monge-Naranjo (2010)). Lochner and Monge-Naranjo (2010) report that amount of student loans coming from private sources has risen since then, although private loans appear to be most prevalent among graduate students in professional schools and undergraduates at high-cost private universities.

²⁴Unsubsidized loans accrue interest over the deferment period. Accrued interest payments are added to

repayment plan, borrowers have up to ten years to repay their loan in full. However, there are several circumstances under which borrowers can postpone their repayments, for example in times of economic hardship. At the maximum, borrowers can extend the repayment period up to 25 years.²⁵ As an approximation, I thus assume that all borrowers repay their loans in 25 years, starting from the year in which college students graduate.²⁶

I denote the per-period installment by l_j . College students are exempted from debt service while enrolled in college, which implies that $l_j = 0$ if $1 \leq j \leq 4$ and for all $j \geq 31$. Accrued interest is accumulated and added to the principal. After having finished college, the total amount due is thus given by $\chi(1+r^*)^4$, where r^* is the subsidized interest rate on college loans, $r^* < r$. Loans are subsidized by the government. For $5 \leq j \leq 30$, l_j is calculated as follows. First, I assume that the repayment scheme is fixed, as it is common for installment loans. In addition, I also assume that the (per-period) loan redemption ι is constant and given by $\frac{\chi(1+r^*)^4}{25}$. This implies that residual debt declines at a linear rate over the repayment period. Because of falling interest payments, l_j falls from $\iota + r^*\chi(1+r)^4$ to 0 between $j = 5$ and $j = 31$.

I assume that the interest rate, net of capital taxes or subsidies, is the same for assets and loans. This requires that $r^* = (1 - \tau_k)r$. Together with the repayment structure, this implies that I can assume that all households take out loans up to the upper limit $\bar{\chi}$. The amount of the loan resources that households do not spend on college or non-college related expenses can be saved in financial assets. By doing so, households can exactly replicate the payment stream resulting from the loan. This simplifies the households decision problem considerably, since I do not need to keep track of the debt holdings of each individual household. I further assume that the loan system is managed by a financial intermediary, and that the subsidy on the interest rate are covered by the government.

With respect to the existence of borrowing limits, Keane and Wolpin (2001) document that borrowing limits exist and are tight. They report that the maximum annual amount of loans that can be taken out of the GSL Program is about 25 percent of the average undergraduate tuition, room and board expenses across two and four year colleges in the academic year 1997-98. They conclude that it is impossible to finance even one year of college using uncollateralized loans.²⁷

the loan principal. For subsidized loans, the government covers the interest on loans while students are enrolled (Lochner and Monge-Naranjo (2010)).

²⁵<http://studentaid.ed.gov/PORTALSWebApp/students/english/OtherFormsOfRepay.jsp> (retrieved on September 10, 2011)

²⁶Since death does not occur before the age of 53 ($j=33$), no agent dies in negative net worth.

²⁷Keane and Wolpin (2001) identify borrowing constraints by measuring net worth. The proportion of the sample with negative net worth increases from 11.5 percent at age 20 to 16.3 percent at age 25 and then falls to 9.1 percent at age 30. Keane and Wolpin (2001) further report that average net debt, conditioning on having negative net worth is generally on the order of 5000 US- Dollar. At age 25, i.e. shortly after college, 16 percent of the group with negative net worth held debt of more than 10,000 US- Dollar and 20 percent less than 1000 US- Dollar, which are small amounts, given that the average undergraduate tuition, room and board expenses across two and four year colleges in the academic year 1997-98 amounted to almost 10,000 US-Dollar, according to Keane and Wolpin (2001)

Altruism and the Timing of Parental Transfers. Besides loans, parental transfers are the other source of financing at the beginning of the life cycle. In the model, there are intended and unintended transfers flowing from parents to their children. Intended transfers are generated by one-sided altruism, as in Laitner (2001), Nishiyama (2002), GMV (2010) and Brown et al. (2011).²⁸ One-sided altruism implies that parents care about the lifetime well-being of their mature children, but not the other way round. I assume that a parent household decides about intended transfers at the age of 51, when the child household is 21 and forms an independent household. Gale and Scholz (1994) find that the mean age of transfer-givers is 55 years in the 1983-1986 Survey of Consumer Finances.

Given the initial endowment received from the parent, the child can then decide whether to consume the resources received from the parent, save it for future consumption and/or spend it for college investment. It is important to note that because parents are altruistic, they anticipate the division of their transfers that is optimal from the children's point of view. The resulting allocation is the same as if parents decided themselves about the optimal split between education investment and pure financial transfers.²⁹

Unintended transfers take the form of end-of-life bequests and arise because lifetime is uncertain and annuity markets are absent. Upon the event of death, the remaining parental wealth is passed on to the children.

A remark regarding the timing of the transfers is in order. As noted by Laitner (2001), if children are not borrowing constrained, the timing of transfers is indeterminate. In this case, parents are indifferent between giving transfers at the beginning of the child's life cycle in form of inter-vivos transfers, at the end of their own life in form of bequests, or by making a sequence of gifts in-between. This holds as parents can commit to a specific sequence of transfers, implying that parents and children do not strategically interact. If children however face (potentially) binding borrowing constraints, it is optimal for an altruistic parent to make transfers as early as possible, i.e. at the beginning of the life cycle. The result that binding liquidity constraints trigger transfers holds even if there is no commitment and parents and children play a dynamic game (Barczyk and Kredler (2010)). There is indeed ample empirical evidence for the fact that households that are subject to binding borrowing constraints are more likely to receive transfers (see e.g. Cox (1990)). In line with this evidence, I thus assume that all intended transfers resulting from parental altruism are made at the beginning of the child's life cycle, where borrowing constraints are most likely to be binding.

Technology. A representative firm produces a final output good Y using aggregate physical capital K and aggregate labor measured in efficiency units L as inputs. The production technology $F(K, L)$ obeys constant returns to scale.

Government. The government collects taxes from labor and asset income at rates (τ_w, τ_k) . Tax revenues are used to finance pension benefits pen , college subsidies ν as well as

²⁸I allow for the fact that altruism may be imperfect.

²⁹In the following, I will use the terms 'financial transfers', 'inter vivos transfers' and 'financial inter vivos transfers' interchangeably.

subsidies on the interest rate r . The government adjusts τ_w in order to balance its budget in every period.

3 Recursive Problem

It is convenient to describe the recursive problem by going backward from retirement age. I thus first consider the optimization problem of a parent household, afterwards the problem of a child household.

3.1 Parent households

A parent household is of age j such that $31 \leq j \leq J^{\max}$. A parent household works during the first 15 years and is retired afterwards. The household faces a declining survival probability, $\psi_j < 1$ if $j \geq 33$.

Parent Households, After Retirement. When retired ($j \in \{46, \dots, J^{\max}\}$), the household receives social security benefits, pen , and chooses consumption c_j and its end-of-period wealth level a_{j+1} . The optimization problem of this household can be written in recursive formulation as follows:

$$V_{p,r}(s_{p,r}^j) = \max_{c_j, a_{j+1}} \{u(c_j) + \beta\psi_j V_{p,r}(s_{p,r}^{j+1})\} \quad (1)$$

where β is the discount factor. $V_{p,r}(s_{p,r})$ is the value function of a retired household facing a state vector $s_{p,r}$, given by

$$s_{p,r}^j = (a_j, j)$$

The household maximizes (1) subject to

$$\begin{aligned} a_{j+1} &= (1 + r(1 - \tau_k))a_j + pen - c_j, \\ a_{j+1} &\geq 0 \text{ and } c_j \geq 0 \end{aligned}$$

r is the interest rate on capital and taxes on capital income are denoted by τ_k . In the terminal period $J^{\max} = 60$, the continuation value is zero. Households consume their remaining wealth and $a_{J^{\max}+1} = 0$.

Parent Households, Working, After Transfers Have Been Made. A parent household who is working and who has provided transfers to its child household is j years old, where $32 \leq j \leq 45$. This household earns w per (efficient) unit of labor, which is supplied inelastically. Total labor productivity depends on the education level e as well as on the realization of the idiosyncratic productivity shock $\eta^{j,e}$. The parents' problem then reads as follows:

$$V_{p,w}(s_{p,w}^j) = \max_{c_j, a_{j+1}} \left\{ u(c_j) + \beta\psi_j \sum_{\eta^{j+1,e} \in N^e} V_{p,w}(s_{p,w}^{j+1}) Q(\eta^{j,e}, \eta^{j+1,e}) \right\} \quad (2)$$

where $Q(\eta^{j,e}, \eta^{j+1,e})$ is the law of motion for productivity shock and $s_{p,w}^j$ is the vector of state variables at age j , given by

$$s_{p,w}^j = (a_j, e, \eta^{j,e}, j)$$

Agents maximize (2) subject to the budget constraint

$$a_{j+1} = (1 + r(1 - \tau_k))a_j + (1 - \tau_w)w\varepsilon_j^e \eta^{j,e} - c_j,$$

where τ_w denotes a linear tax on labor income. Again, financial assets are required to be positive:

$$a_{j+1} \geq 0 \text{ and } c_j \geq 0$$

In the last period before retirement ($j = 45$), the continuation value is replaced by $\beta\psi_j V_{p,r}(s_{p,r}^{j+1})$.

Parent Household, First Period: I now describe the parental problem in the first period of parenthood $j = 31$. At this stage, parents choose their own savings a_{j+1} and the transfers to their child household in such a way that their total utility is maximized. Transfers are denoted by tra .

Expressed in terms of a Bellman equation, the decision problem of a parent household at $j = 31$ reads as

$$V_{p,w}(s_{p,w}^{31}) = \max_{c_{31}, a_{32}, tra} \left\{ u(c_{31}) + \beta\psi_{31} \sum_{\eta^{32} \in N^e} V_{p,w}(s_{p,w}^{32}) Q(\eta^{31}, \eta^{32}) + \varsigma V_0(s_0) \right\} \quad (3)$$

where ς is the intergenerational discount factor. I allow for imperfect altruism, that is, $0 \leq \varsigma \leq 1$. If $\varsigma = 0$, parents care only about their own utility. The model thus nests a pure life cycle economy ($\varsigma = 0$) and a dynastic model ($\varsigma = 1$) as extreme cases. Both Laitner (2001) and Nishiyama (2002) show that the observable flow of transfers is consistent with an intermediate case.

$V_0(s_0)$ denotes the discounted lifetime utility of a young household at the beginning of his economic life. It depends on $s_0 = (tra, \kappa, \mathbf{f}, e^p)$. $V_0(s_0)$ is described in greater detail in the next section.

Parent households face the following state variables at $j = 31$:

$$s_{p,w}^{31} = (a_{31}, e, \eta^{31,e}, \mathbf{f}, \kappa)$$

Notice that the innate ability level of the child \mathbf{f} as well as the level of κ are part of the parent household's state space because together with the parental education level e , these variables determine the cost of college as well as the likelihood of dropping out of college.

The parental budget constraint is given by

$$a_{32} = (1 + r(1 - \tau_k))a_{31} + (1 - \tau_w)\varepsilon_{31}^e \eta^{31,e} w - tra - c_{31}$$

$$a_{j+1} \geq 0 \text{ and } c_{p,1} \geq 0$$

In the following, I describe the problem of a young (child) household.

3.2 Young households

Young household are of age $1 \leq j \leq 30$. The problem of a young household depends on the status of his parents. Children with parents who are alive expect to receive bequests in the future, and thus need to keep track of their parents' wealth holdings. In the following, I thus distinguish between young households with deceased parents and young households with parents who are alive. Since parent household do not die before age $j = 33$, the problem of a young household with deceased parents starts at the age of $j = 3$.³⁰

Young households with deceased parents: I first describe the recursive problem of a young households whose parents are dead. The optimization problem of this household reads as

$$V_{y,d}(s_{y,d}^j) = \max_{c_j, a_{j+1}} \left\{ u(c_j) + \beta \sum_{\eta^{j+1} \in N^e} V_{y,d}(s_{y,d}^{j+1}) Q(\eta^j, \eta^{j+1}) \right\} \quad (4)$$

where $V_{y,d}(\cdot)$ is the value function of a young household with deceased parents and $s_{y,d}$ is the vector of state variables in period j , which is given by

$$s_{y,d}^j = (a_j, \mathbf{f}, e, \eta^{j,e}, j)$$

Agents maximize (4) subject to the constraints

$$\begin{aligned} a_{j+1} &= (1 + r(1 - \tau_k))a_j + (1 - \tau_w)w\varepsilon_j^e \eta^{j,e} - c_j - l_j \\ a_{j+1} &\geq 0 \text{ and } c_j \geq 0 \end{aligned}$$

College students ($e=\text{col}$) do not receive income while studying, that is, $\varepsilon_j^{e=\text{col}} = 0$ for $1 \leq j \leq 4$. Moreover, by assumption, debt does not need to be serviced in the first four years ($l_j = 0$), independently of the education level.

At $j = 30$, in the last period before young households become parents, the continuation value is replaced by

$$\beta \sum_{\mathbf{f}^c} \sum_{\eta^{j+1} \in N^e} \int_{\kappa} V_{p,w}(s_{p,w}^{31}) f(\kappa; \mu_{\mathbf{f},e^p}, \sigma_{\mathbf{f},e^p}) d\kappa Q(\eta^j, \eta^{j+1}) \Gamma(\mathbf{f}, \mathbf{f}^c)$$

where the transmission of household specific types between (becoming) parents and their children needs to be taken into account.

Young households whose parents are alive. I now consider the problem of a young household whose parent household is alive. I denote the asset holdings of the parents of a household at age j by a_j^p .

³⁰This is for simplicity as parental education is an important determinant of children's dropout probability. Recall that students who leave college without a degree do this at the end of model period $j = 2$.

The child household does not know when the parent household dies. As a consequence, the value function is a weighted sum of the utility it receives if the parent household dies and the utility which is obtained if the parent continues to live for another period. Since the age gap between parent and child household is 30 years, the survival probability of the parent household is given by ψ_{j+30} . The optimization problem can thus be described by the following functional equation:

$$V_{y,a}(s_{y,a}^j) = \max_{c_j, a_{j+1}} \left\{ \begin{array}{l} u(c_j) + \beta(1 - \psi_{j+30}) \sum_{\eta^{j+1} \in N^e} V_{y,a}(s_{y,a}^{j+1}) Q(\eta, \eta^{j+1}) \\ + \beta \psi_{j+30} \sum_{\eta^{j+1} \in N^e} V_{y,a}(s_{y,a}^{j+1}) Q(\eta, \eta^{j+1}) \end{array} \right\} \quad (5)$$

The state vector $s_{y,a}^j$ is given by

$$s_{y,a}^j = (a_j, a_j^p, \mathbf{f}, e, e^p, \eta^{j,e}, \eta^{j^p, e^p}, j)$$

Rational expectations imply that children are perfectly able to forecast their parents' asset holdings. Therefore, children know the law of motion of parents' asset holdings, and a_j^p , e^p , η^{j^p, e^p} become part of the children's information set. When parents are retired, this reduces to a_j^p .

College students deserve special attention, as they face a certain probability of leaving college without graduating. In line with empirical evidence discussed before, I assume that dropping out occurs after two years of college. Hence, at $j = 2$, the recursive problem of college students reads as follows:

$$V_{y,a}(s_{y,a}^{j=2} |_{e=col}) = \max_{c_j, a_{j+1}} \left\{ \begin{array}{l} u(c_j) + \beta(1 - \lambda_{\mathbf{f}, e^p}) \sum_{\eta^{j+1} \in N^e} V_{y,a}(s_{y,a}^{j=3} |_{e=col}) Q(\eta, \eta^{j+1}) \\ + \beta \lambda_{\mathbf{f}, e^p} \sum_{\eta^{j+1} \in N^e} V_{y,a}(s_{y,a}^{j=3} |_{e=hs}) Q(\eta, \eta^{j+1}) \end{array} \right\} \quad (6)$$

where $\lambda_{\mathbf{f}, e^p}$ is the probability of dropping out from college without a degree. λ depends the innate ability \mathbf{f} and the level of parental education e^p .

At the age of $j = 30$, the child household knows that its parent household will die for sure in the current period. The continuation value is as stated above for the young households with deceased parents.

The budget constraint depends on whether parents died in the previous period or not. If not, the constraints are given by

$$\begin{aligned} a_{j+1} &= (1 + r(1 - \tau_k))a_j + (1 - \tau_w)w\varepsilon_j^e \eta^{j,e} - c_j - l_j, \\ a_{j+1} &\geq 0 \text{ and } c_j \geq 0 \end{aligned}$$

Again, college students ($e=col$) do not receive income while studying and $\varepsilon_j^{e=col} = 0$ for $1 \leq j \leq 4$. Moreover, $l_j = 0$ if $1 \leq j \leq 4$, independently of the education level.

If the parent household dies in period $j - 1$, child households inherit the residual (end-of-period) wealth holdings from their parents, and the per-period flow budget constraint reads as

$$\begin{aligned} a_{j+1} &= (1 + r(1 - \tau_k))(a_j + a_j^p) + (1 - \tau_w)w\varepsilon_j^e \eta^{j,e} - c_j - l_j \\ a_{j+1} &\geq 0 \text{ and } c_j \geq 0 \end{aligned}$$

3.2.1 Young Households, College Investment Decision

At the beginning of their economic life, young households receive transfers tra from their parents, decide about attending college and how much to borrow. I assume that this takes place before the realization earnings shock and before consumption and saving decision are made in the first period.

The college investment decision is based on the cost of college κ as well as the innate ability f and parental education e^p . Because the productivity shocks $\eta^{j,e}$ depend on the education level, the decision about attending college needs to be taken *before* the first period's productivity shock $\eta^{1,e}$ materializes.

$$V_0(s_0) = \max_{e \in \{hs, col\}, a_1} \{E[V_{y,a}(s_{y,a}^{j=1} | e=col, a_1)], E[V_{y,a}(s_{y,a}^{j=1} | e=hs, a_1)]\} \quad (7)$$

where the subscript 0 indicates that the college enrolment decision takes place before young households haven taken other economic actions. The expectation is taken with respect to the set of possible productivity shocks in the first period.³¹

The constraints of this problem are as follows:

$$\begin{aligned} a_1 &= tra + \bar{\chi} - \kappa|_{e=col} \nu(a_{p,1}, \varepsilon_j^e \eta^{31,e} w) \\ a_1 &\geq 0 \end{aligned}$$

The state vector that households faces when making this decision is given by

$$s_0 = (tra, \kappa, \mathbf{f}, e^p)$$

Young households take parental transfers tra as given. All young households borrow up to the maximum debt limit, which is $\bar{\chi}$.³² Based on the cost of college (net of subsidies) $\kappa\nu$, households decide about attending college. The optimal education choice for an individual that faces the state vector s_0 is described by $edu(s_0)$. The residual financial resources a_1 are kept for future consumption.

4 Definition of a Stationary Competitive Equilibrium

Define S as the state space corresponding to the vector of state variables in the household problems (1-7) with generic element s . Let Σ_S be the sigma algebra on S and denote the corresponding measurable space by (S, Σ_S) . The measure of households on (S, Σ_S) is denoted by Φ .

I define a stationary recursive equilibrium in the economy I study as follows

³¹I assume that productivity shocks in the first period are equally likely.

³²As outlined above, the assumption that all households borrow up the limit is inconsequential, since the interest rates on borrowing and lending are the same.

Definition 1 Given a government policy $\{\text{pen}, \tau_k\}$ and a college subsidy rule ν , a Stationary Recursive Competitive Equilibrium is a set of value functions $V(s)$ and a set of policy functions $\{\text{edu}(s), c(s), a(s), \text{tra}(s)\}$, non-negative prices of physical capital and of effective labor $\{r, w\}$, and a measure of household Φ such that hold:

1. Given prices and policies, the value functions $V(s)$ are the solutions to problems (1)-(7). The functions $\{\text{edu}(s), c(s), a(s), \text{tra}(s)\}$ are the associated policy functions.
2. The prices r and w are consistent with profit maximization of the firm, i.e.

$$r + \delta = F_K(K, L)$$

$$w = F_L(K, L)$$

3. The labor tax rate τ_w adjusts such that the government's budget is balanced:

$$\tau_w = \frac{\text{pen} \int_{S_{p,r}} d\Phi + \Xi + \Lambda - \tau_k r K}{wL}$$

where the total amount of college subsidies Ξ is given by

$$\Xi = \int_{S_{y,0}} (1 - \nu(a_{p,1}, \varepsilon_j^e \eta^{31,e,w})) \kappa \text{edu}(s_{y,0}) d\Phi$$

and the total amount of subsidies on the interest rate for loans Λ is given by

$$\Lambda = (r - r^*) \int_{S_{y,0}} \bar{\chi} d\Phi$$

4. The financial intermediary runs a balanced budget:

$$\int_{S_{y,0}} \bar{\chi} d\Phi = \int_{S^{j \leq 30}} \iota d\Phi$$

5. The asset market, the labor market and the final good market clear:

$$K = \int_S a(s) d\Phi$$

$$L = \int_{S^{j \leq J^{\text{Work}}}} \varepsilon_j^{hs} \eta^{j,hs} d\Phi + \int_{S^{j \leq J^{\text{Work}}}} \varepsilon_j^{col} \eta^{j,col} d\Phi$$

$$C + [K - (1 - \delta)K] + T + I = F(K, L)$$

where

$$\begin{aligned}
C &= \int_S c(s) d\Phi \\
T &= \int_{S_{p,w}^{j=31}} tra(s) d\Phi \\
I &= \kappa \int_{S_{y,0}} \nu(a_{p,1}, \varepsilon_j^e \eta^{31,e} w) \kappa edu(s) d\Phi
\end{aligned}$$

6. *The Aggregate Law of Motion is stationary:*

$$\Phi = H(\Phi)$$

The function H is generated by the policy functions $a(s)$, $c(s)$, $tra(s)$, $edu(s)$, the transition matrix of productivity shocks, $Q(\eta^j, \eta^{j+1})$, the distribution of \mathcal{K} , $f(\kappa; \mu_{\mathbf{f}, e^p}, \sigma_{\mathbf{f}, e^p})$ and transition matrix of innate ability $\Gamma(\mathbf{f}, \mathbf{f}^e)$.

Notice that the stationarity condition requires that child households are (on average) 'identical' to their parents in the sense that they reproduce their parent household's distribution once they become parents themselves. This in turn implies that the distribution of transfers and inheritances that child households receive is consistent with the distribution of transfers that is actually left by parent households. I present more details about the computational procedure in the appendix.

5 Parameterization and Calibration

Experiment Design. To measure the extent to which borrowing constraints are binding and whether they became more binding over time, I compute the strength of borrowing constraints in two steady-state equilibria. The first equilibrium allocation is calibrated to be consistent with the US economy at the beginning of the 1980s. More specifically, I target the size of parental transfers, the college enrolment rate, the correlation of college education across generations and dropout rates. The resulting allocation is labeled as 'economy 1980'. In the appendix, I show that the model is also consistent with a series of other empirical regularities related to college enrolment behavior in the US, which are not used as targets in the calibration procedure. Between 1980 and 2000, the US economy was characterized by an increase in the skill premium, larger residual earnings inequality as well as a rise in tuition fees. I incorporate these developments by changing the respective parameters in the model. I label the resulting equilibrium allocation as 'economy 2000'. In the following, I describe the calibration procedure in greater detail.³³

³³In all my experiments, it is assumed that households have perfect foresight and that the regime change occurs unexpectedly. These assumptions are standard in the literature, see e.g. Heathcote et al. (2010). Moreover, the economies 1980 and 2000 are studied at steady-state, the transition between the two steady-states is not considered. This approach is valid if the changes that occurred between 1980 and 2000 had

5.1 Economy 1980

I distinguish between parameters that are set outside the model and others that are calibrated internally.

5.2 Parameters Set Outside of the Model

Technology, Demographics and Preferences. Utility from consumption in each period is given by $u(c) = \frac{c^{1-\gamma}}{1-\gamma}$. Production follows the aggregate production function $F(K, L) = K^\alpha L^{1-\alpha}$, where L is aggregate labor measured in efficiency units (see equation 5). By assumption, college graduates have a skill premium and supply more labor in efficiency units compared to high school graduates. The skill-premium is independent of the fraction of high skilled labor.³⁴ I set the capital share in income (α) equal to 0.36, as estimated by Prescott (1986). Following Imrohoroglu et al. (1995) and Heer (2001), I assume that capital depreciates at an annual rate of 8 percent. The conditional survival probability ψ_j is taken from the National Vital Statistics Report, Vol. 53, No. 6 (2004) and refers to the conditional survival probability for the US population. Only values between age 53 and age 80 are used. I assume that the survival probability is zero for agents at the age of 81. The survival probability for households younger than 53 years is set equal to 1.³⁵ The preference parameter γ determines the relative risk aversion and is the inverse of the intertemporal elasticity of substitution. I follow Attanasio et al. (1999) and Gourinchas and Parker (2002) who estimate γ using consumption data and find a value of 1.5. This value is well in the interval of 1 to 3 commonly used in the literature.

Earnings Process. I assume that the process that governs the productivity shocks $\eta^{j,e}$ follows an AR(1) process with persistence parameter ρ^{hs} for high school graduates and ρ^{col} for college graduates. The variance of the innovations are σ^{hs} and σ^{col} , respectively. These parameters are estimated by Hubbard et al. (1995)(HSZ in the following) from the 1982 to 1986 Panel Study of Income Dynamics (PSID).

sufficient time to affect the behavior of households. This assumption is likely to be fulfilled because the NLSY97 sample used by Belley and Lochner (2007) includes individuals if they had reached age 21 by 2004. This implies that their parents were between 25 and 30 at the beginning of the 1980s, assuming an age gap of 30 years. According to the time trends of earnings inequality presented by Krueger and Perri (2006) and Heathcote et al. (2010), the rise in earnings inequality (both within and between education groups) began in the early years of the 1980s. Moreover, a substantial share of the total increase between 1980 and 2000 materialized during the 1980s. It thus seems reasonable to assume that the rise in earnings inequality influenced the behavior of the parents of those households who were surveyed in the NLSY97 almost throughout their entire working life. Hintermaier and Koeniger (2011) provide support for this assumption. The authors analyze the development of the distribution of net worth between 1983 and 2004. They estimate the structural parameters of their life cycle economy such that the model is consistent with the distribution of net worth in 1983. They then show that if they increase the variance of residual earnings inequality and lower the interest rate by a similar magnitude to what I consider here, their model matches the distribution of net worth observable in the 2004 SCF very well.

³⁴The implicit underlying assumption is that efficiency units supplied by high school and college graduates are perfect substitutes in the production process. GMV (2010) allow for imperfect substitutability.

³⁵The actual survival probability before 53 is close to 1. See the National Vital Statistics Report.

They find that high school graduates have a lower earnings persistence and a higher variance ($\rho^{hs} = 0.946$, $\sigma^{hs} = 0.025$) compared to college graduates ($\rho^{col} = 0.955$, $\sigma^{col} = 0.016$). Storesletten et al. (2004b) confirm these findings.³⁶

These estimates are based on data from the beginning of the 1980s. However, parental transfers that occurred in the 1980s originate from savings that were accumulated in the 1970s or even earlier. Since the US economy experienced an increase in residual earnings inequality between the 1970s and the 1980s, the earnings risk that those parents faced was lower than indicated by the estimates from Hubbard et al. (1995) for the 1980s.

Gottschalk et al. (1994) show that the permanent and the transitory variances changed on average by about 40 percent between the 1970s and the 1980s, with a somehow larger change for high school graduates and a somehow lower change for college graduates. See Gottschalk et al. (1994), Table 1.

Based on these findings, I reduce the variance of the earnings innovation for college graduates by 21 percent and by 48 percent for high school graduates. This implies a $\sigma^{hs} = 0.0169$ for high school graduates and a $\sigma^{col} = 0.013225$ for college graduates. The respective standard deviations are 0.13 and 0.115. The changes are in the range estimated by Meghir and Pistaferri (2004).

I also take the average age-efficiency profile ε_j^e from HSZ, which gives me an estimate of the college premium for different age groups. The authors find that earnings are more peaked for college families, which is in line with findings from other empirical studies. By estimating a fixed-effect model, GMV (2010) propose an alternative way of calibrating the earnings process.³⁷

College Subsidies. In the model, the government subsidizes college education. Subsidies are decreasing in parental income and asset holdings, such that children from richer families receive less support. This reflects the financial aid system at US institutions. Colleges and universities assign financial aid based on the difference between the cost of attending and the so-called "expected family contribution (EFC)", see Feldstein (1995). The EFC is computed according to the discretionary income and the available assets of the applicant's family. Discretionary income consists of capital and labor income. Available assets are calculated as the difference between current wealth holdings and a wealth level that is deemed to maintain the current standard of living, which I denote by \bar{a} .

For simplicity and because this specification is common in the literature, I assume that the fraction of college expenses that is covered by the subsidy is linearly decreasing in the level of parental resources:

$$\nu = \max(\{\nu_0 - \nu_{asset}(\max(0, a_{p,w}^{31} - \bar{a}) + (1 - \tau_k) r a_{p,w}^{31}) - \nu_{labor} (1 - \tau_w) \varepsilon_j^e \eta^{j,e} w\}, 0) \quad (8)$$

³⁶It should be noted that the estimates are rather conservative as HSZ use the combined labor income of the husband and wife (if married) plus unemployment insurance for their estimates. I approximate the earnings process with a four-state Markov process using the procedure proposed by Tauchen and Hussey (1991).

³⁷I do not adjust the college premium because it was - at least on average - constant during the 1970s. The male college premium fell slightly at the beginning of the 1970s and has been increasing since the end of the 1980s. See Heathcote et al. (2010), Figure 1.

where $\nu_0, \nu_{asset}, \nu_{labor} \geq 0$. I assume that \bar{a} is two-thirds of the average wealth level in the economy. This choice ensures that a substantial fraction of the population has asset holdings below the threshold. There are three other parameters to be chosen: ν_0 , ν_{asset} and ν_{labor} . ν_0 determines the maximum fraction of college expenses that can be covered by the subsidy. The parameters ν_{asset} and ν_{labor} determine how fast this fraction decreases with asset income and labor income, respectively.

I set ν_0 to 0.75. This choice for ν_0 is based on the observation that in 1979-1980, even students from families in the bottom income quintile finance a significant amount of their college expenses with the help of parental transfers. See the evidence cited in Keane and Wolpin (2001).³⁸

Feldstein (1995) reports that the implied marginal tax rates on (adjustable) income lies within 22 and 47 percent, according to the so-called uniform methodology, which was used to calculate the EFC at the beginning of the 1980s. I choose a value for ν_{asset} of 0.2, which is at the lower end of the spectrum reported by Feldstein (1995). Feldstein (1995) finds that the implied capital levy by the EFC can be quite high. In a sense the saving distortions implied by my calibration are expected to be lower on average than the distortions that result from the US system.

In the data, the share of college expenses that is financed with parental transfers increases rapidly in level of total parental income Keane and Wolpin (2001).³⁹ This suggests that college subsidies are actually falling sharply in parental income.⁴⁰ In order to account for this feature of the data, without at the same time overly distorting the saving decisions, I choose different values for ν_{labor} and for ν_{asset} . More specifically, I assume that ν_{labor} takes a value of 0.7.

5.3 Parameters Calibrated Internally

Discount Factor β . In order to calibrate the discount factor β , I target the ratio of aggregate net worth to aggregate income for the lower 99 % wealth quantile in the US, which is 3.1 (see Storesletten et al. (2004a)). The result is a β of 0.96. The implied interest rate is 3.5 percent per annum, which is in the range commonly reported in the literature.

Children's Weight in Parents' Utility ς . The parameter ς governs transfer behavior in the model. It is the intergenerational discount factor and determines the relative weight

³⁸In 1979-80, for youths with families in the bottom income quintile, about 19 percent of college expenses came from parental transfers. The rest was financed with other internal sources (such as youth's income, from which I abstract) or with scholarships, grants or loans. See Keane and Wolpin (2001). Notice that in my specification, college loans do not count as college subsidies ν . Hence, some students may be able to finance their total college expenses with the help of external funds.

³⁹Keane and Wolpin (2001) show that parental transfers account for 19 percent of total college expenses if parents belong to the poorest 25 percent of the population, and 60 percent if parents belong to the richest 75 percent of the population.

⁴⁰In this context, it is important to notice that the financial aid that a student actually receives in practice does not necessarily cover the difference between the EFC and the cost of college, as it is implicit in the EFC procedure. As pointed out by Dick and Edlin (1997), federal programs do not provide enough subsidized aid to meet the need of all students, and most colleges are not committed to cover the entire residual.

that parents assign to their children’s utility in their own utility function. In order to calibrate ς , I target the ratio of intended transfers to aggregate net worth, i.e. transfers that arise because parents are altruistic in my model. Gale and Scholz (1994) provide information about intra-family transfers in form of inter-vivos transfers, support for college expenses and bequest, using the 1983 and the 1986 wave of the SCF. Inter-vivos transfers and support for college expenses are classified as intended transfers. The flow of both transfer categories amounts to 0.82 percent of total net worth. Bequests, however, amount to 0.88 percent of aggregate net worth and are thus quantitatively more important than both other categories together. The total transfer flow is thus equal to 1.7 percent of total wealth. As noted by Gale and Scholz (1994), it is not clear whether bequests are intended or unintended, because there are no markets to insure against uncertainty about lifetime.

I use my model to compute the total amount of bequests that arise because of missing annuity markets. Given a discount factor $\beta = 0.96$, accidental bequests amount to about 0.3 percent of total wealth. This implies that, in order to be consistent with the data, the model needs to generate intended transfers of 1.4 percent of net worth.⁴¹ The resulting value for ς is 0.7. A ς of 0.7 implies that a parent household discounts the utility of its child household by 30 percent more than it discounts its own utility. This is in line with results obtained from Nishiyama (2002) who uses an altruistic framework to explain the observable degree of wealth inequality in US economy.

Upper Limit Loans $\bar{\chi}$. The maximum amount of loans is calibrated to match the fraction of households with negative or zero financial assets. Because I do not model collateralized debt, such as mortgages, I use net financial assets, instead of net worth. In 1983, the fraction of households with negative or zero financial assets was 25 percent (see Table 1, in Wolff (2000)). The resulting borrowing limit corresponds to about 30 percent of average college expenses for prospective college students. Keane and Wolpin (2001) report that college students can expect to finance 25 percent of their college expenses with the help of loans.

Cost of College Education and Dropout Probabilities. Given the choices for the discount factor β , the intergenerational discount factor ς , the borrowing limit for loans as well as the college subsidies I am now in the position to calibrate the parameters that govern the cost of college education as well as the dropout probabilities.

The cost of college education/observable ability \mathcal{K} are normally distributed, with density function $f(\kappa; \mu_{\mathbf{f}, e^p}, \sigma_{\mathbf{f}, e^p})$. Recall that in my model, the cost of college education and observable ability are directly linked. I assume that $\sigma_{\mathbf{f}, e^p} = 1$. Innate ability \mathbf{f} and parental education e^p thus influence only the mean of the distribution of \mathcal{K} , but not the variance.

I further assume that $\mu_{\mathbf{f}, e^p}$ consists of two parts, i.e. $\mu_{\mathbf{f}, e^p} = \bar{\kappa} + \Delta_{\mathbf{f}, edu^p}$. $\bar{\kappa}$ is the same for all prospective students and determines the average level of college enrolment, while $\Delta_{\mathbf{f}, edu^p}$ captures the differences in family background and innate ability.

⁴¹Recall that the timing of intended transfers not related to college is indeterminate in my model, unless borrowing constraints are binding. This implies that one could interpret part of the inter-vivos transfers as bequest.

\mathbf{f} takes two realizations, $\mathbf{f} \in \{low, high\}$. This results in four combinations for $\Delta_{\mathbf{f},edu^p}$ and for $\lambda_{\mathbf{f},ep}$. Together with $\bar{\kappa}$, there are now nine values which are jointly calibrated.

In order to pin down these parameters, I choose the following nine targets:

- the share of college graduates in the NLSY79, which is 28 percent according to Keane and Wolpin (2001).
- another four targets are chosen such that the model is consistent with the fact that enrolment rates increase in AFQT scores, as shown in Belley and Lochner (2007) using the NLSY79. More specifically, for each of the two bottom ability quartiles, I target the average of the enrolment rates of the four family income groups reported by Belley and Lochner (2007). Since differences in enrolment rates between income groups are small for the bottom AFQT quartiles, averaging provides an accurate description of the enrolment behavior of all income groups. For the two top AFQT quartiles, differences in college enrolment rates between rich and poor families are substantial, as shown by Belley and Lochner (2007). In these cases, I target the enrolment rates of high income families only. This group is key for measuring the impact of borrowing constraints, as it will become clear in the results section.
- the last four targets are the college dropout probabilities for different quartiles of observable scholastic ability, as reported by Chatterjee and Ionescu (2010).⁴²

The choice of these targets ensures that the model is consistent with two patterns in the data, where college enrolment rates are increasing and dropout rates are falling in measured ability.

With the parameters presented in Table 2, the model is consistent with these patterns. As Table 2 shows, the model fits the data closely.

Additional assumptions are needed in order to identify λ and Δ . The reason is that a priori it is not clear how the combinations of innate ability and parental education map into the different observable ability quartiles. I thus assume that $\Delta_{low,hs} \geq \Delta_{low,col} \geq \Delta_{high,hs} \geq \Delta_{high,col}$ and $\lambda_{low,hs} \geq \lambda_{low,col} \geq \lambda_{high,hs} \geq \lambda_{high,col}$. This is equivalent to saying that innate ability has more influence on the costs and benefits of college, compared to parental education. I will also make use of this assumption in the next section to calibrate Γ .

Intergenerational Transmission of Innate Abilities. The intergenerational transmission of innate ability is governed by a 2x2 transition matrix Γ .

⁴²Chatterjee and Ionescu (2010) use the students' SAT score as a measure for academic ability. Depending on the test score, they divide their sample (which is taken from the Beginning Postsecondary Student Longitudinal Survey (BPS 1995/96)) into four groups, where each group contains between 15 and 35 percent of the total number of students in their sample (see their Table 2 on page 22). They find that degree completion rates are increasing with the level of the test scores. They report that the degree completion rates for the four groups are 0.659, 0.7627, 0.8462 and 0.8825. The dropout rates are the complements of these numbers. When computing the degree completion rates, they control for students who do not put effort but simply enroll and dropout shortly after. Since I target the dropout rates for four different ability quartiles, I interpolate the data from Chatterjee and Ionescu (2010) to compute in-between values.

For a given level of parental education, Γ determines how strongly educational achievements are correlated between generations. If the matrix Γ is identical to the identity matrix, children will inherit the same level as their parents. If Γ is instead given by a matrix where all elements on the antidiagonal are equal to one, children and parents are of the opposite type.

Using the NLSY79, Keane and Wolpin (2001) report that parents who have a college degree are 40 percent more likely to have children that are college educated as well, compared to high school educated parents. This suggests that the transmission of \mathbf{f} is somewhere in-between the two extremes. I thus assume that Γ is given by a convex combination of an identity matrix and a matrix where all elements on the antidiagonal are equal to one. A new parameter ϖ is introduced, which determines the weight of identity matrix in this linear combination. Intuitively, the role of this parameter is to shift probability mass to the main diagonal and to make the transmission of ability more persistent. See Castaneda et al. (2003) for a similar approach.

Additional Moments. In the appendix, I show that the model is consistent with other moments that were not used as targets in the calibration procedure. I present empirical facts that are informative about the trade-offs that are relevant for the college enrolment decision. In particular, I document that (i) the distribution of measured ability is skewed to the right, as in the data (see e.g. GMV 2010), (ii) the implied average college expenses are in line with the data, (iii) the model is consistent with differences in tuition expenses with respect to ability, as documented McPherson and Schapiro (2002), and that (iv) the model is line with empirical findings regarding the sensitivity of college enrolment to changes in tuition, as reported by Dynarski (2003). Given the complexity of the college enrolment decision, it is reassuring that my baseline calibration is able to reproduce these moments as well. Further empirical validation is provided in the next section, where I compute the average enrolment gap in the model, following Carneiro and Heckman (2002). Before doing so, I outline the calibration of the economy 2000.

5.4 Economy 2000

Between 1980 and 2000, the US economy was characterized by an increase in the skill premium, higher residual earnings inequality as well as a rise in tuition fees. In this section, I discuss how I adjust the parameter values such that the model is consistent with these changes.

Increase in College Premium. Heathcote et al. (2010), Figure 1, report that the college wage premium for men rose by about 40 percentage points between the beginning of the 1980s and 2000. This is in line with other estimates in the literature, see for example Katz and Autor (1999). The increase in the college premium is implemented by shifting the life cycle earnings profile of college graduates up, while the mean average earnings of high school graduates is kept constant, also see Heathcote et al. (2010).

Increase in Earnings Uncertainty. I assume that the residual earnings variance increases from 0.0169 to 0.025 for high school graduates and from 0.0132 to 0.016 for college graduates. These are the original estimates from HSZ, provided for the 1980s. This was the time when parents accumulated their saving for college investment undertaken at the end of the 1990s. Clearly, the variance of residual earnings continued to increase during the 1990s as well (see e.g. Krueger and Perri (2006) and Heathcote et al. (2010)), so my choice is rather conservative.

Increase in Tuition Fees. Between 1980 and 2000, tuition fees doubled in real-terms, see Collegeboard (2005). I adjust $\bar{\kappa}$ from 2.6 to 5.2, which raises the average tuition fees by 90 percent, relative to average income.

Borrowing Limit for College Loans $\bar{\chi}$. Lochner and Monge-Naranjo (2010) report that cumulative Stafford loan limits remained almost identical in real-terms between the beginning of the 1980 and 2000 (see their Figure 1). To keep borrowing limits constant, I adjust $\bar{\chi}$ to 1.25.

The parameter choices for the economy 2000 are summarized in Table 3.

6 Results

I present the results in the following order: First, I show that the enrolment pattern in the economy 1980 are broadly in line with the data. In particular, the results are consistent with an average enrolment gap of 8 percent, as computed by Carneiro and Heckman (2002). I then compute the fraction of borrowing constrained households with the help of a counterfactual experiment, in which I relax the borrowing limit. I find that 24 percent of all households are borrowing constrained in their college education. I then proceed with the economy 2000. I show that my model is consistent with the widening of the enrolment gaps that are visible in the data. Interestingly, I find that the fraction of borrowing constrained households declined over the course of time, from 24 percent to 18 percent.

6.1 Economy 1980

College Enrolment Rates in the Data and in the Model. I start by comparing the college enrolment rates produced by the benchmark calibration to their empirical counterparts. Figure 2 is taken from Belley and Lochner (2007) (Figure 2a in their paper) and depicts the college enrolment rate for different levels of observable ability and family income. Figure 4 shows the corresponding rates that are generated by the model.

Model and data coincide with respect to the main trends. Both figures show that (i) college enrolment rates are increasing in observable ability, (ii) college enrolment rates are increasing in family income, and (iii) the gap in college enrolment rates between students from rich and from poor families are increasing in observable ability. The first two pattern were used as targets in the calibration procedure, but not the third.

It should be noted that, compared to the data, the model overstates the size of the college enrolment gap between children from rich and poor families for more able children and it understates the enrolment gap for the less able. What is important in the context of this paper is that *on average*, the enrolment gap is in line with the empirical evidence.⁴³ In the next section, I follow the methodology of Carneiro and Heckman (2002) to compute the average enrolment gap.

Reproducing Carneiro and Heckman (2002). Carneiro and Heckman (2002) present a method to compute the fraction of the population that is financially constrained in their college decision. Their method has become the standard tool to address the question of binding borrowing constraints, see for example Restuccia and Urrutia (2004) and more recently Bohacek and Kapicka (2010).

Carneiro and Heckman (2002) use the NLSY79 and compute the fraction of borrowing constrained households with the help of the following steps. First, they divide their sample by parents' income quartiles and children's ability terciles, using AFQT test scores as a proxy for ability. They assume that youth with parents in the highest income quartile are not borrowing constrained. Second, the college enrolment gaps for each ability and income group with respect to the unconstrained income quartile are computed. The fraction of the population that is borrowing constrained is equal to the average enrollment gap across all groups, which can be computed using population weights.

As pointed out by Carneiro and Heckman (2002), even after controlling for ability, family income still seems to play an important role for college enrolment. This is also the message of Figure 4. However, as argued by the authors, family resources are likely to produce many skills which are not fully captured by a single test score. Moreover, family income at the time when students take their college decision is strongly correlated with family income throughout the life cycle.

Carneiro and Heckman thus introduce additional measures for early family background factors, such as parental education, family structure and place of residence. They find that enrolment gaps become considerably smaller, after controlling for long-run effects properly. They conclude that at most 8 percent of the population is financially constraint in their college decision.

I replicate their method, using the distribution of households which is generated by my benchmark calibration.⁴⁴ Without controlling for long-run factors other than measured ability, the average enrolment gap 0.17. Clearly, the magnitude of the average enrolment gap reflects the sizable enrolment gaps that are visible in Figure 4. I then additionally control for parental education, which also determines scholastic ability, but which is only partially incorporated in the result of test scores. I find that after doing so, the average enrolment gap is much smaller, about 7 percent. Put differently, according to the methodology of Carneiro and Heckman (2002), 7 percent of the population in the benchmark economy is

⁴³This means that the discrepancies between model and data cannot be responsible for the difference between the actual fraction of borrowing constrained households and the average enrolment gap.

⁴⁴To be consistent with the rest of my analysis, I group households according to quartiles of measured ability, not terciles, as Carneiro and Heckman (2002)

financially constrained in their college decision. This result is perfectly in line with Carneiro and Heckman’s result from the NLSY79.

Borrowing Constraints. In this section, I compute the fraction of households that would have enrolled in college, if access to credit had been unlimited, given the interest rate in the initial steady- state. I thus increase the upper limit on loans, $\bar{\chi}$ from 1 to 5.2. 5.2 is the natural borrowing limit, given the market clearing interest rate that results from the benchmark calibration in 1980.⁴⁵

The rise in college enrolment following the removal of borrowing limits is substantial: the fraction of college student increases from 36 percent to about 60 percent. This implies that 24 percent of the population were borrowing constrained in their college decision in the benchmark calibration. The fraction of college graduates also increases significantly, from 28 percent to 44 percent.

Strikingly, the fraction of borrowing constrained households resulting from the counterfactual experiment is more than three times as large as the average enrolment gap computed in the previous section. To understand this discrepancy, notice that the average enrolment gap identifies the fraction of borrowing constrained households if the following two assumptions are satisfied:

1. Youth with parents in the top income quartile are not financially constrained in their college enrolment decision.
2. In the absence of financial constraints, college enrolment rates are equal for all income groups (after controlling for observable ability).

Even small violations of the first assumption may lead to a large discrepancy between actual fraction of borrowing constrained households and the average enrolment gap, as the following examples shows.⁴⁶ Suppose that the second assumption is satisfied. Further suppose that relaxing borrowing limits increases the college enrolment rate of youth from rich families by 1 percentage point, independent of the ability level. This increase adds a fraction of 1 percentage point of constrained households to all other income groups. This is because when the average enrolment gap is computed, the difference with respect to the top income quartile is taken. In total, the average enrolment gap underestimates the fraction of borrowing constrained households by 4 percentage points in this example.⁴⁷

A comparison of Figures 4 and 5 reveals that relaxing borrowing constraints increases enrolment rates for the top income quartile by up to 11 percentage points. This shows

⁴⁵The natural borrowing limit denotes the maximum amount of debt that a household can borrow, such that debt is serviced even in the worst income state resulting from the income process of high school graduates. The income of high school graduates is relevant even for the decision about college enrolment because of the possibility of dropping out, in which case households receive the income process of high school graduates.

⁴⁶I am grateful to an anonymous referee for pointing this out.

⁴⁷Notice that when computing the Figures, I do not control for parental education because I want to be consistent with the Belley and Lochner (2007). This will be important in the next section, where I compare the changes in enrolment behavior over time. The enrolment rates after controlling for parental education are reported in Table 6. As can be seen from there, the qualitative results in this section remain unchanged.

that even rich parents fail to provide enough resources to their offspring, due to the fact that intergenerational borrowing is not permitted and altruism is only one-sided. One-sided altruism implies that even rich parents face a trade-off between saving and transferring resources to their children. In particular, parents save for their retirement and also for precautionary reasons, since they face another 15 years of earnings uncertainty.

With respect to the second assumption, it is interesting to notice that enrolment rates are not equal for all income groups after borrowing limits have been relaxed (see Figure 5). In particular, enrolment rates for the poor are higher for some ability groups than the respective rate in the top income quartile.

This finding can be explained as follows. The decision to attend college depends on the tuition costs as well as on the dropout risk. All other things equal, the cost of attending college are lower for youth from poor families, because they receive financial aid. The dropout risk instead depends on innate ability as well as the level of parental education. Children with college educated parents are more likely to enrol in college, all other things equal. Hence, the share of college educated parents in the top income quartile relative to the other quartiles determines the income gradient of college enrolment. The higher the concentration of college educated parents in the top income group, the steeper the income gradient.

I report the fraction of college graduates in different subgroups of the population in Table 7. The smallest fraction of college graduated parents in the top income quartile, absolute and also relative to other income groups, can be found for children in the lowest ability quartile. This explains why the enrolment gradient is relatively flat for children in this group (see Figure 4).

It also explains why enrolment rates increase more for low ability youth with poor parents compared to their counterparts who have rich parents, after borrowing constraints have been removed (compare Figure 4 and Figure 5). Children in the bottom ability quartile face high costs of attending college. Within this group, children from poor families are more willing to attend college than children with rich parents, since they receive need-based aid. Because intergenerational borrowing is prohibited, poor parents are not willing or not able to provide sufficient resources to finance their offspring's education, even if attending college is cheap.

In recent work, Brown et al. (2011) provide empirical support for my finding that many parents are not willing or not able to invest efficiently in their offspring's education. The authors estimate that about half of the children in their sample are potentially constrained in their college outcome. They identify parents who are unwilling (or unable) to meet their expected financial contribution as those who do not provide post-schooling transfers to their children. Brown et al. (2011) find that financial aid affects the educational attainment of those children whose parents do not make post-schooling transfers, while it does not have an impact on the schooling attainment of those who receive post-schooling transfers.

To summarize, because intergenerational borrowing is prevented, some of the rich families and poor parents with children in the bottom ability quartile fail to provide sufficient resources to their offspring. As a result, the fraction of households that would like to attend college if they could borrow against their future earnings exceeds the average enrolment gap by a factor of three. The fact that I model the intergenerational transmission of wealth and

ability explicitly allows me to identify the subgroups in the population for which borrowing constraints are binding.⁴⁸ This approach is an important contribution with respect to the previous literature, in which the income gradient of college enrolment rates is taken as a measure for borrowing constraints, as in, for example, Ellwood and Kane (2000), Cameron and Heckman (1998) and Carneiro and Heckman (2002), or in which the parental transfer function was specified in reduced form, as in Keane and Wolpin (2001).

6.2 Economy 2000

The previous section has shown that the model is broadly consistent with the enrolled patterns that were observable in the US economy at the beginning of the 1980s. Ellwood and Kane (2000) and Belley and Lochner (2007) document that the enrolment gaps between rich and poor families became larger over time. The enrolment rates that Belley and Lochner (2007) computed from the NLSY79 and the NLSY97 are shown in Figure 2 and 3. In this section, I analyze whether my model is consistent with the changing role of family income in determining college entry. This is done by comparing the enrolment rates of the economy 1980 and of the economy 2000. In a second step, I then compute the fraction of households that is borrowing constrained in their college decision in the economy 2000. I find that a share of 18 percent of the population is borrowing constrained. This means that the fraction of constrained households decreased over time.

Enrolment Gaps. In Figures 4 and 6, I plot the enrolment rates for the economy 1980 and the economy 2000, respectively. A comparison reveals that family income is much more important in determining enrolment rates in the economy 2000 than it is in the economy 1980, conditional on observable ability. This is in line with the empirical findings by Belley and Lochner (2007). Interestingly, the income gradient increases more for youth in the lowest ability quartile, both in the data and in the model. For all other ability groups, the model generates enrolment gap that are too big compared to the data.

Share of College Graduates. It is worth emphasizing that the fraction of college graduates rises from 28 percent to 36 percent if we move from the economy 1980 to the economy 2000. This is consistent with the reported increase in the data (Census Bureau 2004).⁴⁹

Borrowing Constraints. In order to measure the fraction of borrowing constrained households, I conduct the following experiment. I set $\bar{\chi}$ to the highest possible value, which

⁴⁸It should be noted that for some subgroups of the population, borrowing constraints are 'less binding' than predicted by the average enrolment gap. Otherwise, the fraction of constrained households would be even higher.

⁴⁹Moreover, the model generates a fall in the real interest rate of almost 1 percentage point, from 3.5 percent in 1980 to 2.7 percent in 2000. See Table 5. Hintermaier and Koeniger (2011) and the references therein who also find that the real interest rate fell between 1 and 2 percentage points in the period under consideration.

in this case is $\bar{\chi} = 4.9$.⁵⁰ The interest rate is kept fix at 2.7 percent per annum. As result, the share of households that enrol in college increases by 18 percentage points (see Table 5).

This result implies that the fraction of households that is borrowing constrained in their college decision actually decreased over time, from 24 percent in 1980 to 18 percent in 2000. This result is striking, given the fact that family income has become more important over time as a determinant of college entry. The latter finding has given rise to concerns that today, borrowing constraints are binding for a larger fraction of the population than two decades ago, see Ellwood and Kane (2000), Belley and Lochner (2007) and Lochner and Monge-Naranjo (2010). These concerns are intuitive, in particular against the background that in the US, tuition became more expensive and earnings inequality increased during the same period. The remainder of this section will be devoted to clarify why enrolment gaps and the share of the population affected by financial constraints developed in the opposite direction.

To understand why the fraction of constrained households declined over time, compare the enrolment rates for with and without borrowing limits for the two economies (Figures 4 and 5 for the economy 1980 and Figures 6 and 7 for the economy 2000). The striking difference between the economy 1980 and 2000 is that the enrolment rates of the rich do not rise if borrowing conditions are relaxed. Borrowing constraints are not binding anymore for the income-rich in the the economy 2000. This is in sharp contrast to the economy in 1980, where even youth from affluent families turned out to be financially constrained in their college decision.

There are three channels that explain these enrolment patterns: the increase in income inequality, a compositional effect and the fact that enrolling in college is more risky in the economy 2000 compared to the economy 1980. They are outlined in greater detail in the following.

Increase in inequality. The rise in earnings inequality, both within and between education groups, benefits parents in the upper income quartile in the sense that their income rises relative to the average income. Because parents are altruistic, this translates into an increase in financial support to their children. Therefore, the intergenerational borrowing limit becomes less binding. As a consequence, the share of constrained children coming from rich families falls, and the enrolment rate among youth from this group goes up, relative to the average.

Compositional effect. The increase in the college premium also changes the educational composition of high-income families. The fraction of college education parents in the top income quartile rises. Consider Table 7, where I report the fraction of college graduates in the different ability/income quartiles for the economies 1980 and 2000. The top income quartile in the economy 2000 consists almost exclusively of college graduated parents, independently of the ability of children. The change with respect to the economy 1980 is most drastic

⁵⁰Notice that the natural borrowing limit for the economy 2000 is stricter than for the 1980 economy because of the increase in earnings risk, which results in a lower level of the worst income state. Consequently, the maximum amount of debt a household is able to repay in the worst state is lower as well.

for the lowest ability quartile, where only 46 percent of parents have a college degree in the economy 1980.⁵¹ This compositional effect is important because children of college graduates are more likely to be successful in college, which tightens the observable link between family income and college enrolment. This is particularly true for low ability children, as the compositional effect is most pronounced for this group. Ellwood and Kane (2000) provide empirical evidence for the fact that the educational composition of income-rich households indeed changed over time. They show that the likelihood that income-rich households are at the same time better educated increased over the last decades. Their analysis also reveals that the fraction of college graduates in the top quartile of the income distribution rose by more than the average.

Increase in financial risk associated with college attendance. College education became more expensive over time. There is some probability that students drop out of college without graduating. If they dropout, the return from college is negative, because students bear the cost, but do not receive a benefit in exchange. Therefore, the decision to enrol in college can be seen as the decision to accept a gamble. An increase in tuition lowers the certainty equivalence associated with this gamble.⁵²

The risk of losing money cannot be avoided by borrowing, since college loans need to be repaid in full, independently of whether college was completed successfully or not (Chatterjee and Ionescu (2010)). This means that in the economy 2000, young households who face a high risk of dropping out are less willing to borrow in order to finance college compared to the economy 1980. Since children of high-school educated parents are on average less likely to complete college successfully, this explains why college enrolment rates do not change much for children from low income households in the economy 2000 once I relax borrowing constraints. Interestingly, the increase in financial risk implied by the rise in tuition fees reinforces the compositional channel outlined above.

To summarize, this section has shown how important it is to model the transmission of wealth and ability between parents and children in order to account for the changing role of family income in determining college entry. In the next section, I shed more light on the role that parental education plays for determining college enrolment.

6.3 The Role of Parental Education

In this section, I shed more light on the role of parental education. Parental education is commonly seen as a key determinant of children's college enrolment behavior (see e.g. Carneiro and Heckman (2002)). On the one hand, parental education serves as a good proxy for children's pre-college ability.⁵³ On the other hand, education also determines the average

⁵¹The fraction of college graduates is, all other things equal, lower for lower ability groups because of the positive correlation between parental education and children's ability.

⁵²In addition, the benefit of attending became more risky over time due the increase in earnings risk.

⁵³In the model section above, some theories that imply a positive correlation between parental education and children's academic ability are summarized.

life time income of households. College educated parents thus have on average a higher lifetime income than high school educated parents (see e.g. Hubbard et al. (1994)). All other things equal, this may imply that children of college educated parents receive more transfers, which makes it easier for them to finance college.

The implications of these two channels for the measurement of college enrolment gaps are studied in this section. In this context, I also compare the extent to which the life cycle savings pattern generated by the model match with data.

Parental Education as a Proxy for Pre-College Ability. There is a debate in the literature about whether controlling for parental education and other variables improves the ability of the average college enrolment gap to measure the fraction of borrowing constrained households (see Carneiro and Heckman (2002) and Kane (2006)).⁵⁴

My model can be used to clarify the role that parental education plays in that context. Parental education determines the academic ability of children. Children of college educated parents are more likely to achieve better test scores, and are less likely to leave college without a degree. In the following, I label this relationship as the 'ability channel'. Ignoring this channel when calculating the enrolment gaps overstates the influence of parental resources - this is the point of Carneiro and Heckman (2002).

Parental education is also an important determinant of parental saving and transfer behaviour. This can be seen from Figure 8, where I plot the life cycle profile of net worth for college and high school educated households. College educated households have, on average, richer at the time their children go to college. They also transfer, on average, more resources to their offspring, as indicated by the drop in net worth that occurs at the age of 52. The change in net worth is much more pronounced for college educated households than for households with a high school degree. In the next section, I describe this pattern in greater detail.

Because lifetime income of college educated parents is on average higher than lifetime income of their high school educated counterparts, college graduates have a higher utility gain from transferring resources to their offspring. All other things equal, children with college educated parents are more likely to receive support from their parents than other children. This however also means that they are, *ceteris paribus*, also more likely to attend college. I call differences in parental transfer behavior the 'transfer channel'. Ignoring the transfer channel when computing the enrolment gap understates the role of parental resources - this is the point of Kane (2006). In fact, the transfer channel may contribute to explaining

⁵⁴Carneiro and Heckman (2002) write: "Family income in the adolescent years is strongly correlated with family income throughout the life cycle. In addition, long run family resources are likely to produce many skills that are not fully captured by a single test score. When we control for early family background factors (parental education, family structure and place of residence) we greatly weaken the relationship between family income and college enrolment." (p. 721). Kane (2006) writes: "However, to the extent that parents can help finance their children's education with current income or accumulated assets, the distinction between short-term constraints and long-term constraints is unclear. Indeed, we may be understating the effect of borrowing constraints by first conditioning on test performance, to the extent that these factors, too, are related to long-term family wealth" (p. 1394).

the difference between the average enrolment gap and the fraction of borrowing constrained households.

Separating the ability from the transfer channel is not straightforward because the two are intertwined: college educated parents have more able children, meaning that college investment costs less and promises a higher return, hence college educated parents transfer more resources to their children.

My model can be used to quantify the role of each of the two channels on the size of the average enrolment gap. In the following, I focus on the economy 1980. The difference between the college enrolment gap before and after controlling for parental education measures the joint impact of the ability and the transfer channel. In the baseline calibration, the average enrolment gap without controlling for parental education is 17 percent. If I control for parental education, the enrolment gap shrinks to 7 percent (see Table 4). That is, controlling for the ability channel and the transfer channel jointly decreases the enrolment gap by 10 percentage points.

I now conduct an experiment that allows me to measure the relative contribution of the transfer channel. I shut down the ability channel in the economy 1980 by setting ϖ to 0.5, which implies that innate ability is not transmitted from parents to children anymore. Moreover, I modify cost of college education Δ and the dropout probability λ such that the cost and return of attending college are independent of parental education.⁵⁵ In this experiment, only the transfer channel is operative. The average enrolment gap - after conditioning on measured ability but without conditioning on parental education - is about 9 percent (see Table 4). So it is about half of its value compared to the benchmark calibration. If I additionally control for parental college education, the average enrolment gap shrinks to about 6 percent. This corresponds to a decline of 3 percentage points, which can be seen as a measure for the strength of the transfer channel. This indicates that the transfer channel on its own is significantly weaker than the joint interaction between the transfer channel and the ability channel. The transfer channel thus makes only a small contribution to explaining the high fraction of households that is borrowing constrained.

Average Assets by Age. As the previous section made clear, differences in savings and transfers between high school graduates and college graduates are an important margin in the model. Aim of this section is to discuss the empirical relevance of the life cycle profiles of asset holdings that are plotted in Figure 8. I focus on the following three stages of the life cycle: beginning (young households enter into the economy), transfer stage (young households become parents and provide transfers to their children) and pre-retirement stage.

At the *beginning* of their life cycle, young households, on average, start their economic life with positive wealth holdings. In particular college graduates own a substantial amount of wealth at the beginning of the life cycle which is used to cover tuition fees and consumption during college education. Between age 26 and 48, the average asset holdings of college graduates are actually lower than average asset holdings of high school graduates. While

⁵⁵The new choices for Δ and λ are now given by $\Delta_{\mathbf{f}=high} = 0.4$, $\Delta_{\mathbf{f}=low} = -0.85$, and $\lambda_{\mathbf{f}=high} = 0.2$, $\lambda_{\mathbf{f}=low} = 0.52$. Notice that Δ and λ only depend on innate abilities. Parental education does not matter anymore.

this implication of the model is at odds with the data (see e.g. Hubbard et al. (1994) and Hubbard et al. (1995)), it is consistent with the differences in the shape of earnings profiles between high school and college graduates. See Hubbard et al. (1994), Figures 2b and 2c, and Hubbard et al. (1995) for a similar result.⁵⁶ The life cycle earnings profile of college graduates is more hump-shaped. Hence, college graduates save less and borrow more when young in order to smooth out their consumption profile over the life cycle.⁵⁷

At the age of 51, young households become parents and *provide transfers to their children*. At this stage of the life cycle, Figure 8 shows two facts that are interesting. First, prior to the point when transfers are made, college educated households own more assets than high school educated households. Second, after transfers have been made, the wealth of college educated parents declines by more than the wealth of high school educated parents.

The first observation can be explained by the fact that college graduates have stronger incentives to save for their offspring's education. They know that their children are more likely to attend college, as they are more able than the children of high school graduates. If I shut down the link between parental education and children's ability, as I did in the previous subsection, I find that difference in savings between college and high school educated parents disappears (see Figure 9). Put differently, if children's ability is independent of their parents education, parental saving for inter-vivos transfers are also independent of education.

The second observation, namely the drop in assets following transfers, is due to the 'transfer channel' outlined above. College educated parents have, on average, a higher life time income. Consequently, they provide on average more transfers to their children. This means that college educated households dissave more than high school educated households, in order to finance the transfers. In the next subsection, I analyze to what extent these predications of the model are in line with evidence from the Survey of Consumer Finances.

Third, *close to retirement*, college graduates hold more assets compared to high school graduates. As pointed out by Hubbard et al. (1994), the life cycle asset profile of college graduates is more hump-shaped in the data as well.⁵⁸ In my model, the fact that high school graduates save relatively less for retirement follows from the fact that pension benefits are independent of lifetime income, implying that retirement income accounts for a bigger fraction of the total income of high school graduates, as compared to the total income of college graduates. See Huggett and Ventura (2000) and Storesletten et al. (2004a) who show that a key feature of the US social security system is that annual benefits are not proportional to social security taxes paid previously.

Further Empirical Evidence on Parental Transfers and Saving . In this section,

⁵⁶Since the earnings profiles I use for my quantitative analysis are take from Hubbard et al. (1994) and Hubbard et al. (1995), this is to be expected.

⁵⁷Hubbard et al. (1994) show that introducing a consumption floor and an asset-based means tested social security system suppresses the savings of high school graduates and helps to reconcile the model with the empirical evidence. I leave the interesting question how a consumption floor and asset based means tested social security programs interact with saving for transfers for future research. I expect that introducing these features into my model will boost the fraction of households that is financially constrained in their college decision because they reduce the incentive to save.

⁵⁸Hubbard et al. (1994) abstract from borrowing and therefore focus on assets only.

I shed more light on the empirical relevance of some patterns that are produced by the model. At the stage when transfers occur in the model, college educated parents save less and provide more transfers to their children compared to high school educated parents. In the previous section, I showed that these findings are caused by differences in parental lifetime income ('transfer channel') and by differences in children's academic ability ('ability channel'). The ability channel turned out to be more important for determining college enrolment.

In order to evaluate to what extent there are similar pattern in the data, I use the 1983-86 panel of Survey of Consumer Finances (SCF). In 1986, extensive information about households' transfer behavior was collected. The 1983-86 panel has therefore become a standard reference with respect to parental inter-vivos transfers (see e.g. Gale and Scholz (1994)). I compute total transfers given by a household as the sum of all monetary transfers and college expenses, which are reported separately in the SCF.⁵⁹ I use savings in constant prices accumulated between 1983 and 1986. In order to be consistent with the model, I consider only those households that are between 45 and 55 years old and that have at least one child.

College graduated parents have on average higher savings compared to high school graduated parents, see panel 1, column 3 of Table 8. The result is not statistically significant, however. Panel 2 of Table 8 reveals that college graduates also transfer more to their children. This result is statistically significant.

The decision to save may be strongly influenced by current income realizations, as the life cycle hypothesis would predict. In column 4 of Table 8, I thus additionally control for current income. Moreover, in column 5, I control for current income and wealth, where wealth can be seen as a proxy for past income realizations. Any difference in transfer or saving behavior of parents is now solely due to the transfer channel and/or the ability channel. Interestingly, independently of whether I control for wealth and income or for income only, I find that college graduates save less and transfer more to their children compared to high school graduates. A possible interpretation of this finding is that for a given wealth and income level, college educated parents foresee on average high income in the future and are thus more willing to provide more resources to their offspring. According to Hubbard et al. (1995), the income difference between college graduates and high school graduates is particularly pronounced in the years before retirement.⁶⁰

Children of college graduates are on average more able and thus also more likely to attend college. My model predicts that in the absence of this ability link, behavioral differences between college graduated parents and high school graduated parents become smaller. The SCF does not contain information about children's ability. Given this constraint, I include in column 6 of Table 8 only households who had at least one child in college between 1983 and 1986. I expect that observable ability differences between children become smaller once I control for college attendance.

Strikingly, controlling for college attendance of children renders the estimated saving and

⁵⁹The 1986 SCF reports monetary transfers only if the transfer amount is above 3000 US-Dollar.

⁶⁰An alternative explanation could be that retirement benefits are tied to labor income, which would imply that college graduates receive higher pension benefits compared to high school graduates. I am grateful to an anonymous referee for pointing this out.

transfer gap insignificant. This hints to the fact that the ability channel is more relevant in determining parental transfers, which confirms the conclusion I reached above for the model.

It is important to stress that this does not constitute a formal test, since there are potentially many alternative explanations that are consistent with the patterns documented in this section. Nevertheless, it is interesting to see that the model's prediction about the joint behavior of transfers and saving appear to be qualitatively consistent with the data, given that most of the literature has focused on difference in saving alone (see e.g. Hubbard et al. (1995) or Dynan et al. (2004)).

7 Conclusion

In this paper, I have analyzed the driving forces behind the changing role of family income in determining college entry. Moreover, I have studied to what extent the average enrolment gap is informative about the fraction of borrowing constrained households.

In order to do so, I have developed an incomplete markets model that features parental altruism and an educational choice. My model allowed me to study the behavioral adjustments of parental transfers following from changes in the economic environment that have occurred in the US economy over the last 30 years, in particular, the increase in earnings inequality (between and within education groups).

My results show that the increase in earnings inequality is key for explaining the increase in the college enrolment gap between youth from rich and from poor families. Interestingly, my model indicates that the changing role of family income is not due to a larger fraction of households that is borrowing constrained in their college choice. Actually, it turns out that the opposite is the case: the fraction of households that is borrowing constrained in their college decision decreased from 25 percent at the beginning of the 1980s to about 20 percent in 2000. In general, the share of borrowing constrained households is much larger than indicated by the average enrolment gap, which is only 7 percent in 1980.

An important extension of my model would be to model the determinants of academic ability more carefully. Mainly for simplicity, I assumed that children's ability is partly determined by parental education, and partly by innate ability, which is also partly inherited from parents. I will leave it for future research to replace this reduced form with a more structural approach, in which early education is explicitly taken into account. By extending my model in this direction one could, for example, investigate how (intergenerational) borrowing constraints at various stages in the life cycle interact. This result would be important for the design of optimal education policies.

References

- Ábráhám, A. (2004). Wage differentials and educational policy with skill-biased technological change in og setting. Mimeo, University of Rochester.
- Aiyagari, S. R. (1994). Uninsured idiosyncratic risk and aggregate saving. *The Quarterly Journal of Economics*, 109(3):659–84.
- Akyol, A. and Athreya, K. (2005). Risky higher education and subsidies. *Journal of Economic Dynamics and Control*, 29(6):979–1023.
- Attanasio, O., Banks, J., Meghir, C., and Weber, G. (1999). Humps and bumps in lifetime consumption. *Journal of Business & Economic Statistics*, 17(1):22–35.
- Barczyk, D. and Kredler, M. (2010). A dynamic model of altruistically-motivated transfers. Mimeo, Department of Economics, Universidad Carlos III de Madrid.
- Barro, R. J. (1974). Are government bonds net wealth? *Journal of Political Economy*, 82(6):1095–1117.
- Becker, G. S. (1974). A theory of social interactions. *Journal of Political Economy*, 82(6):1063–93.
- Becker, G. S. and Tomes, N. (1979). An equilibrium theory of the distribution of income and intergenerational mobility. *Journal of Political Economy*, 87(6):1153–89.
- Belley, P. and Lochner, L. (2007). The changing role of family income and ability in determining educational achievement. *Journal of Human Capital*, 1(1):37–89.
- Black, S. E. and Devereux, P. J. (2010). Recent developments in intergenerational mobility. NBER Working Papers 15889, National Bureau of Economic Research, Inc.
- Bohacek, R. and Kapicka, M. (2010). A quantitative analysis of educational reforms in a dynastic framework. Mimeo, Department of Economics, University of California, Santa Barbara.
- Brown, M., Scholz, J. K., and Seshadri, A. (2011). A new test of borrowing constraints for education. *Forthcoming, Review of Economic Studies*.
- Cameron, S. V. and Heckman, J. J. (1998). Life cycle schooling and dynamic selection bias: Models and evidence for five cohorts of american males. *Journal of Political Economy*, 106(2):262–333.
- Carneiro, P. and Heckman, J. J. (2002). The evidence on credit constraints in post-secondary schooling. *Economic Journal*, 112(482):705–734.

- Carneiro, P., Meghir, C., and Pairey, M. (2007). Maternal education, home environments and the development of children and adolescents. IZA Discussion Papers 3072, Institute for the Study of Labor (IZA).
- Carroll, C. D. (2006). The method of endogenous gridpoints for solving dynamic stochastic optimization problems. *Economics Letters*, 91(3):312–320.
- Castaneda, A., Diaz-Gimenez, J., and Rios-Rull, J.-V. (2003). Accounting for the u.s. earnings and wealth inequality. *Journal of Political Economy*, 111(4):818–857.
- Caucutt, E. and Lochner, L. (2004). Early and late human capital investments, credit constraints, and the family. 2004 Meeting Papers 129, Society for Economic Dynamics.
- Cawley, J., Heckman, J., and Vytlacil, E. (2001). Three observations on wages and measured cognitive ability. *Labour Economics*, 8(4):419–442.
- Chatterjee, S. and Ionescu, F. (2010). Insuring college failure risk. Working Papers 10-1, Federal Reserve Bank of Philadelphia.
- Clausen, A. and Strub, C. (2011). Envelope theorems for non-smooth and non-concave optimization. mimeo, University of St. Gallen.
- Collegeboard (2005). Trends in college pricing 2005.
- Collegeboard (2010). What it costs to go to college.
- Cox, D. (1990). Intergenerational transfers and liquidity constraints. *The Quarterly Journal of Economics*, 105(1):187–217.
- Cunha, F. (2007). A time to plant and a time to reap. Mimeo, University of Chicago.
- Cunha, F. and Heckman, J. (2007). The technology of skill formation. *American Economic Review*, 97(2):31–47.
- Dick, A. W. and Edlin, A. S. (1997). The implicit taxes from college financial aid. *Journal of Public Economics*, 65(3):295–322.
- Dynan, K. E., Skinner, J., and Zeldes, S. P. (2004). Do the rich save more? *Journal of Political Economy*, 112(2):397–444.
- Dynarski, S. M. (2003). Does aid matter? measuring the effect of student aid on college attendance and completion. *American Economic Review*, 93(1):279–288.
- Ellwood, D. and Kane, T. (2000). Who is getting a college education? family background and the growing gaps in enrollment. In Danziger, S. and Waldfogel, J., editors, *Securing the Future*. New York: Russell Sage Foundation,.

- Feldstein, M. (1995). College scholarship rules and private saving. *American Economic Review*, 85(3):552–66.
- Gale, W. G. and Scholz, J. K. (1994). Intergenerational transfers and the accumulation of wealth. *Journal of Economic Perspectives*, 8(4):145–60.
- Gallipoli, G., Meghir, C., and Violante, G. L. (2010). Education decisions and policy interventions: A general equilibrium evaluation. Mimeo, New York University.
- Garriga, C. and Keightley, M. P. (2007). A general equilibrium theory of college with education subsidies, in-school labor supply, and borrowing constraints. Working Papers 2007-051, Federal Reserve Bank of St. Louis.
- Gomes, J., Greenwood, J., and Rebelo, S. (2001). Equilibrium unemployment. *Journal of Monetary Economics*, 48(1):109 – 152.
- Gottschalk, P., Moffitt, R., Katz, L. F., and Dickens, W. T. (1994). The growth of earnings instability in the us labor market. *Brookings Papers on Economic Activity*, pages 217–272.
- Gourinchas, P.-O. and Parker, J. A. (2002). Consumption over the life cycle. *Econometrica*, 70(1):47–89.
- Haveman, R. and Wolfe, B. (1995). The determinants of children’s attainments: A review of methods and findings. *Journal of Economic Literature*, 33(4):1829–1878.
- Heathcote, J., Storesletten, K., and Violante, G. L. (2010). The macroeconomic implications of rising wage inequality in the united states. *Journal of Political Economy*, 118(4):681–722.
- Heckman, J. J., Lochner, L., and Taber, C. (1998). General-equilibrium treatment effects: A study of tuition policy. *American Economic Review*, 88(2):381–86.
- Heckman, J. J., Lochner, L. J., and Todd, P. E. (2005). Earnings functions, rates of return and treatment effects: The mincer equation and beyond. NBER Working Papers 11544, National Bureau of Economic Research, Inc.
- Heckman, J. J., Stixrud, J., and Urzua, S. (2006). The effects of cognitive and noncognitive abilities on labor market outcomes and social behavior. *Journal of Labor Economics*, 24(3):411–482.
- Heer, B. (2001). Wealth distribution and optimal inheritance taxation in life-cycle economies with intergenerational transfers. *Scandinavian Journal of Economics*, 103(3):445–65.
- Heer, B. and Maußner, A. (2005). *Dynamic General Equilibrium Modelling: Computational Methods and Applications*. Springer, Springer.
- Hintermaier, T. and Koeniger, W. (2011). On the evolution of the us consumer wealth distribution. *Review of Economic Dynamics*, 14(2):317–338.

- Holter, H. (2011). Accounting for cross-country differences in intergenerational earnings persistence: The impact of taxation and public education expenditure. Mimeo, University of Pennsylvania.
- Hubbard, R. G., Skinner, J., and Zeldes, S. P. (1994). The importance of precautionary motives in explaining individual and aggregate saving. *Carnegie-Rochester Conference Series on Public Policy*, 40(1):59–125.
- Hubbard, R. G., Skinner, J., and Zeldes, S. P. (1995). Precautionary saving and social insurance. *Journal of Political Economy*, 103(2):360–99.
- Huggett, M. (1996). Wealth distribution in life-cycle economies. *Journal of Monetary Economics*, 38(3):469–494.
- Huggett, M. and Ventura, G. (2000). Understanding why high income households save more than low income households. *Journal of Monetary Economics*, 45(2):361–397.
- Imrohoroglu, A., Imrohoroglu, S., and Joines, D. H. (1995). A life cycle analysis of social security. *Economic Theory*, 6(1):83–114.
- Kane, T. (2006). Public intervention in postsecondary education. In Hanushek, E. A. and Welch, F., editors, *Handbook of the Economics of Education*. Amsterdam: Elsevier.
- Kaplan, G. and Violante, G. L. (2010). How much consumption insurance beyond self-insurance? *American Economic Journal: Macroeconomics*, 2(4):53–87.
- Katz, L. F. and Autor, D. H. (1999). Changes in the wage structure and earnings inequality. In Ashenfelter, O. and Card, D., editors, *Handbook of Labor Economics*, volume 3 of *Handbook of Labor Economics*, chapter 26, pages 1463–1555. Elsevier.
- Keane, M. P. and Wolpin, K. I. (2001). The effect of parental transfers and borrowing constraints on educational attainment. *International Economic Review*, 42(4):1051–1103.
- Kennickell, A. and Shack-Marquez, J. (1992). Changes in family finances from 1983 to 1989: evidence from the survey of consumer finances. *Federal Reserve Bulletin*, (Jan):1–18.
- Krueger, D. and Perri, F. (2006). Does income inequality lead to consumption inequality? evidence and theory. *Review of Economic Studies*, 73(1):163–193.
- Laitner, J. (2001). Wealth accumulation in the u.s.: Do inheritances and bequests play a significant role? Working Papers wp019, University of Michigan, Michigan Retirement Research Center.
- Light, A. and Strayer, W. (2000). Determinants of college completion: School quality or student ability? *Journal of Human Resources*, 35(2):299–332.
- Light, A. and Strayer, W. (2002). From bakke to hopwood: Does race affect college attendance and completion? *The Review of Economics and Statistics*, 84(1):34–44.

- Linsenmeier, D. M., Rosen, H. S., and Rouse, C. E. (2002). Financial aid packages and college enrollment decisions: An econometric case study. NBER Working Papers 9228, National Bureau of Economic Research, Inc.
- Lochner, L. J. and Monge-Naranjo, A. (2010). The nature of credit constraints and human capital. *Forthcoming, American Economic Review*.
- Long, B. T. (2004). How do financial aid policies affect colleges?: The institutional impact of the georgia hope scholarship. *Journal of Human Resources*, 39(4).
- McPherson, M. S. and Schapiro, M. O. (2002). The blurring line between merit and need in financial aid. *Change: The Magazine of Higher Learning*, 34(2):38 – 46.
- McPherson, M. S. and Schapiro, M. O. (2006). Us higher education finance. In Hanushek, E. A. and Welch, F., editors, *Handbook of the Economics of Education*. Amsterdam: Elsevier.
- Meghir, C. and Pistaferri, L. (2004). Income variance dynamics and heterogeneity. *Econometrica*, 72(1):1–32.
- Nishiyama, S. (2002). Bequests, inter vivos transfers, and wealth distribution. *Review of Economic Dynamics*, 5(4):892–931.
- Ozdagli, A. and Trachter, N. (2009). College enrollment, dropouts and option value of education. Mimeo, Einaudi Institute.
- Prescott, E. C. (1986). Theory ahead of business cycle measurement. Staff Report 102, Federal Reserve Bank of Minneapolis.
- Restuccia, D. and Urrutia, C. (2004). Intergenerational persistence of earnings: The role of early and college education. *American Economic Review*, 94(5):1354–1378.
- Ríos-Rull, J. V. (1997). *Computation of Equilibria in Heterogeneous Agent Models*. Federal Reserve Bank of Minneapolis. Research Department.
- Rodriguez, S. B., Diaz-Gimenez, J., Quadrini, V., and Rios-Rull, J.-V. (2002). Updated facts on the u.s. distributions of earnings, income, and wealth. *Quarterly Review*, pages 2–35.
- Stinebrickner, T. R. and Stinebrickner, R. (2007). The effect of credit constraints on the college drop-out decision a direct approach using a new panel study. Working Paper 13340, National Bureau of Economic Research.
- Storesletten, K., Telmer, C. I., and Yaron, A. (2004a). Consumption and risk sharing over the life cycle. *Journal of Monetary Economics*, 51(3):609–633.
- Storesletten, K., Telmer, C. I., and Yaron, A. (2004b). Cyclical dynamics in idiosyncratic labor market risk. *Journal of Political Economy*, 112(3):695–717.

- Tauchen, G. and Hussey, R. (1991). Quadrature-based methods for obtaining approximate solutions to nonlinear asset pricing models. *Econometrica*, 59(2):371–96.
- Wolff, E. N. (2000). Recent trends in wealth ownership, 1983-1998. Working Paper wp300, Levy Economics Institute, The.

8 Appendix I: Additional Empirical Validation

The distribution of observable ability in my model implies a certain distribution of tuition. The distribution of tuition, in turn, is key for explaining college enrolment rates as well as the fraction of borrowing constrained households. In this section, I present a number of stylized facts related to college enrolment behavior. I argue that my quantitative model is consistent with these facts.

Skewness of the Observable Ability Distribution. The distribution of measured ability levels that results from these parameter choices is skewed to the right, with a median (-3.15) which is smaller than the mean (-2.9). Right-skewness of skills is often found in empirical work, see for example GMV (2010), who compute the distribution of AFQT results in the NLSY.

Average Direct Cost of College Education. The average amount students pay in the model (based on their measured ability), relative to GDP per capita, is 2.1. Notice that this is for four-years of education, before need-based financial aid is subtracted.

The respective ratio in the data is 1.7. This takes into account the average yearly fees for tuition, room and board charged by 4-year institutions (public and private), which was approximately 5000 Dollar at the beginning of the 1980s (in current Dollars of 1980) (see the Digest of Educational Statistics). GDP per capita was about 12 000 Dollars (current prices) in 1980 (see World Bank Economic Indicators).

Distribution of direct cost of college education. I now turn to the distribution of direct cost that prospective students face. McPherson and Schapiro (2006) report that 100 points more in the SAT score lower total tuition fees by between 500 to 2300 Dollars. According to Chatterjee and Ionescu (2010), the difference between the 30th percentile and the 65th percentile in the distribution of SAT scores is exactly 200 points. That means that students at the 65th percentile of the distribution of observable abilities pay between 8 and about 40 percent less than students at the 30th percentile, measured in terms of per-capita GDP. In my model, the respective difference is 50 percent. The distribution of SAT scores is highly non-linear (see Chatterjee and Ionescu (2010)). The same applies to the distribution of observable abilities in the model. A comparison of percentiles that are further apart thus becomes increasingly difficult. It is interesting to note that the fraction of the population that faces tuition costs of zero or less is small (less than 1 percent). Generous scholarships at many US schools are typically reserved to the top 1 percent of the applicant pool.

Sensitivity of college enrolment to tuition. Dynarski (2003) estimates that subsidizing college with an additional 1000 Dollars (in 1998 Dollars) increases college enrolment by about 4 percentage points. I find that giving an equivalent amount to college students in the benchmark calibration raises college enrolment by 3.5 percentage points. In line with empirical evidence, I also find that the response to changes in tuition decreases with family income. I thus conclude that the model describes the sensitivity of college enrolment with respect to changes in tuition well.

9 Appendix II: Solution Algorithm

I solve the quantitative model using a nested fixed point algorithm. The outer loop searches for a fixed point in the interest rate, while the inner loop solves the dynamic program given by (1) - (7) as described in the next section. The inner loop solves the hybrid model which nests both the pure life cycle economy and a model with infinitely lived dynasties as special cases. The hybrid nature of the model manifests itself in the fact that the parental value function $V_{p,w}(s_{p,w}^{31})$ contains the discounted future utility of the child and vice versa. I follow the Laitner (2001) when solving this problem. I start with a guess for the parental value function, $V'_{p,w}(s_{p,w}^{31})$. Given this guess, I solve the child's problem as specified in (5), (4), (6) and (7). I describe the solution technique in greater detail in the next section. I then compute an update for the parental value function, $V''_{p,w}(s_{p,w}^{31})$. I keep on repeating this process until convergence is achieved.⁶¹

9.1 Computing the Decision Rules

I compute the optimal decision rules for consumption and saving by adapting the 'endogenous grid point method' (EGM), first outlined by Carroll (2006).⁶² The EGM derives the optimal choices based on inverting the first-order conditions, for a given grid of tomorrow's asset choices. As a result, a grid of corresponding optimal asset levels for today's problem arises.⁶³

Different from 'standard' dynamic programming problems, the program specified by (1) - (7) contains a discrete choice, namely the education decision. The presence of a discrete choice may generate kinks in the parental value function $V_{p,w}(s_{p,w}^{31})$, leading to non-differentiable and non-concave parts. Only young households who receive transfers above a certain threshold are able to attend college. This in turn implies that only parents whose wealth is higher than a certain threshold may be able to provide sufficiently high transfers, such that their children are able to attend college. While it can be shown that first-order conditions are still necessary for optimality in this case (see Clausen and Strub (2011)), they are not sufficient anymore.⁶⁴ Concavity guarantees that the solution is a global maximum and is thus a desirable property of any maximization problem.

Fortunately, there are several model features which smooth out the kinks generated by the education choice. First of all, parental transfers can be used for college and non-college related expenditures. At the threshold, when young households decide to enter college, they are forced to reduce their other consumption expenditures accordingly. The resulting utility loss partly outweighs the utility gain associated with college attendance. Therefore, the kink in the (parental) value function at the threshold is less pronounced.⁶⁵ Second of all, skill

⁶¹The algorithm converges at a geometric rate, see Laitner (2001)

⁶²An exception is the parental problem (3), for reasons that I describe in the next paragraph.

⁶³Therefore its name 'endogenous grid method'.

⁶⁴Clausen and Strub (2011) derive envelope theorems for non-concave and non-smooth optimization problem. They show that optimal decisions are never at the kinks induced by discrete choices.

⁶⁵A similar argument applies to the savings of a parent household in model period $j = 31$. Parental savings increment parental wealth holdings, which are part of the child household's state space. Because parents

accumulation in form of college success and (observable) ability is stochastic. Uncertainty generates a 'smoothing effect', as demonstrated by Gomes et al. (2001).

It turns out that these two model elements are sufficient to make $V_{p,w}(s_{p,w}^{31})$ concave. $V_{p,w}(s_{p,w}^{31})$ is plotted in Figure 10. In order to make sure that this finding is not the result of some kind of numerical approximation routine, I solve the parental problem in $j = 31$ using a standard grid search procedure.

9.2 Computation of the Equilibrium

Using the policy functions which were computed previously, I can now solve for the equilibrium allocation. Computing an equilibrium involves the following steps:

1. Choose the policy parameters, that is, determine the social security replacement rate rep , the tax rate for capital income τ_k and a college subsidy rule ν .
2. Provide an initial guess for the aggregate (physical) capital stock K_0 , the aggregate human capital stock H_0 and the labor tax rate τ_w . Given the guesses for K and H , use the first-order conditions from the firm's problem to obtain the relative factor prices r and w .
3. Compute the optimal decision rules as outlined in the previous section.
4. Compute the time invariant measure Φ of agents over the state space.
5. Compute the aggregate asset holdings K_1 and the new human capital stock L_1 using asset market clearing condition. Given K_1 and L_1 , update r , w and τ_w .
6. If $m = \max\left(\frac{K_1 - K_0}{K_1}, \frac{L_1 - L_0}{L_1}\right) < 10^{-3}$ stop; otherwise return to step 2 and replace K_0 with K_1 and L_0 with L_1 .

In step 4, I find the time-invariant measure of agents Φ by iterating on the aggregate law of motion, as it is commonly done in models with an infinite time horizon. In the model, the measure of parents in their first period of adulthood depends on the transition of children (because all parents were children one period before). In turn, the measure of children in their first period of life depends on the measure of their parents (because children receive transfers and education). Stationarity requires that the probability measure is constant over time. This implies that, for a given measure of parents, the measure of children exactly reproduces the measure of their own parents.

decide simultaneously about savings and transfers, given their budget constraint in $j = 31$, an increase in transfers reduces savings, all other things equal. From the point of view of the children, the utility loss associated with a reduction in savings partially outweighs the utility gain associated with an increase in transfers.

I approximate the measure of agents by means of a probability density function.⁶⁶ The density function is computed and stored on a finite set of grid points. Following Ríos-Rull (1997), I choose a grid $D^{density}$ which is finer than the one used in the previous step for computing the decision rules, that is $D \subseteq D^{density}$. Choosing a finer grid for the density increases the precision with which the aggregate variables are computed.⁶⁷

The optimal choice will almost surely be off-grid. In order to map the optimal choices onto the grid, I introduce some kind of lottery. An individual with asset choice $a'(\cdot) \in (a_i, a_{i+1})$ is interpreted as choosing asset holdings a_i with probability λ and asset holdings a_{i+1} with probability $(1 - \lambda)$ where λ solves $a'(\cdot) = \lambda a_i + (1 - \lambda)a_{i+1}$. That is, I compute a piecewise linear approximation to the density function. No lottery is needed for agents for which the lower bounds on asset holdings is binding, which is the case for a positive fraction of the population. I thus allocate the grid points such that there are closely spaced in the neighborhood of the lower bound. This is achieved by choosing grid points which are equally spaced in logarithms. I select the upper bound of $D^{density}$ and D such that it is never found to be binding.

I find the time invariant measure of agents Φ by iterating on the aggregate law of motion. The forward recursion starts with an initial distribution of young agents in model period $j = 1$, Φ_1 . This requires an initial guess for the distribution of parents in model period $j = 31$. Following Heer (2001), a uniform distribution is taken as an initial guess. Using the decision rules, one can then derive Φ_{31} , from which an update of Φ_1 can be obtained. Φ_{31} is then updated until convergence.

As a check on the internal consistency, aggregate consumption, investment, transfers and output are computed in order to ensure that the good market clearing condition is approximately satisfied.

⁶⁶Heer and Maußner (2005) argue that approximating the time-invariant measure of agents with the help of a density function saves up to 40 percent of CPU time compared to an approximation using the distribution. Computing the distribution function requires computing the inverse of the policy function.

⁶⁷The gains in precision (as measured by aggregate excess demand) by doing so are enormous. The reason is that the aggregate good market clearing condition is just a weighted average of the individuals' budget constraints, where the weights are derived from the grid points of the density Φ . The finer the grid in Φ , the better will be the correspondence between the optimal policies and the resulting weights, leading to better aggregation results.

10 Appendix III: Graphs




Figure 2: **Enrolment Rates NLSY79.** Source: Belley and Lochner (2007), Fig. 2a.




Figure 3: **Enrolment Rates NLSY97.** Source: Belley and Lochner (2007), Fig. 2b.

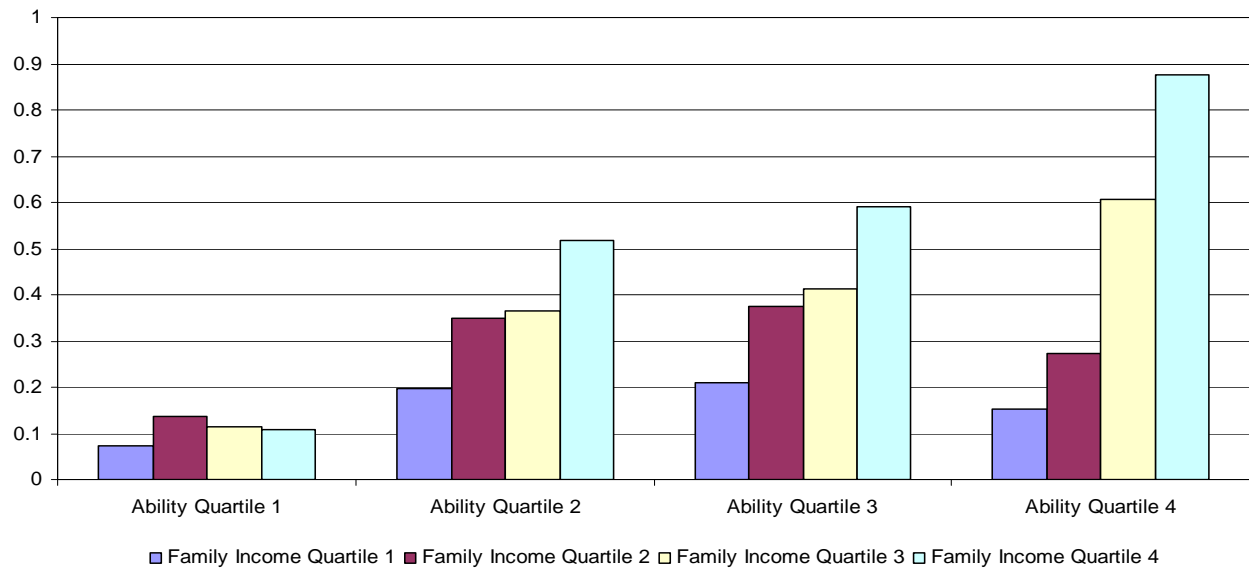


Figure 4: **Enrolment Rates Model Economy 1980.** Fraction of youth enrolled in college, conditional on observable ability/family income quartiles. Variables are defined as described in the text.




Figure 5: **Enrolment Rates Model Economy 1980, No Borrowing Constraints.** Fraction of youth enrolled in college, conditional on observable ability/family income quartiles. Variables are defined as described in the text.

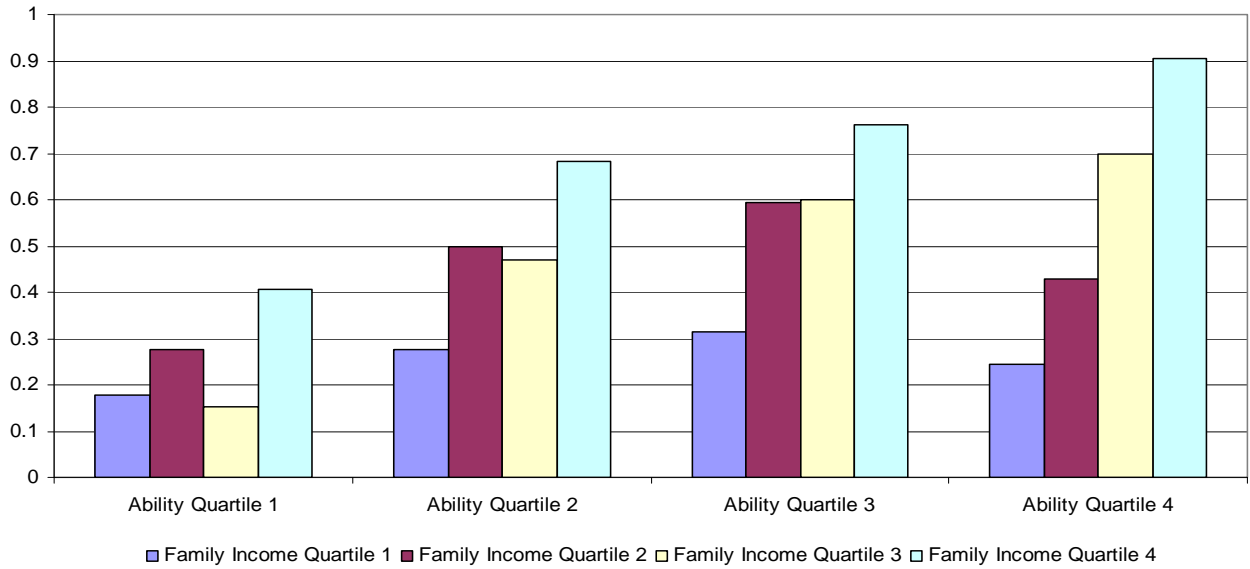


Figure 6: **Enrolment Rates Model Economy 2000.** Fraction of youth enrolled in college, conditional on observable ability/family income quartiles. Variables are defined as described in the text.




Figure 7: **Enrolment Rates Model Economy 2000, No Borrowing Constraints.** Fraction of youth enrolled in college, conditional on observable ability/family income quartiles. Variables are defined as described in the text.

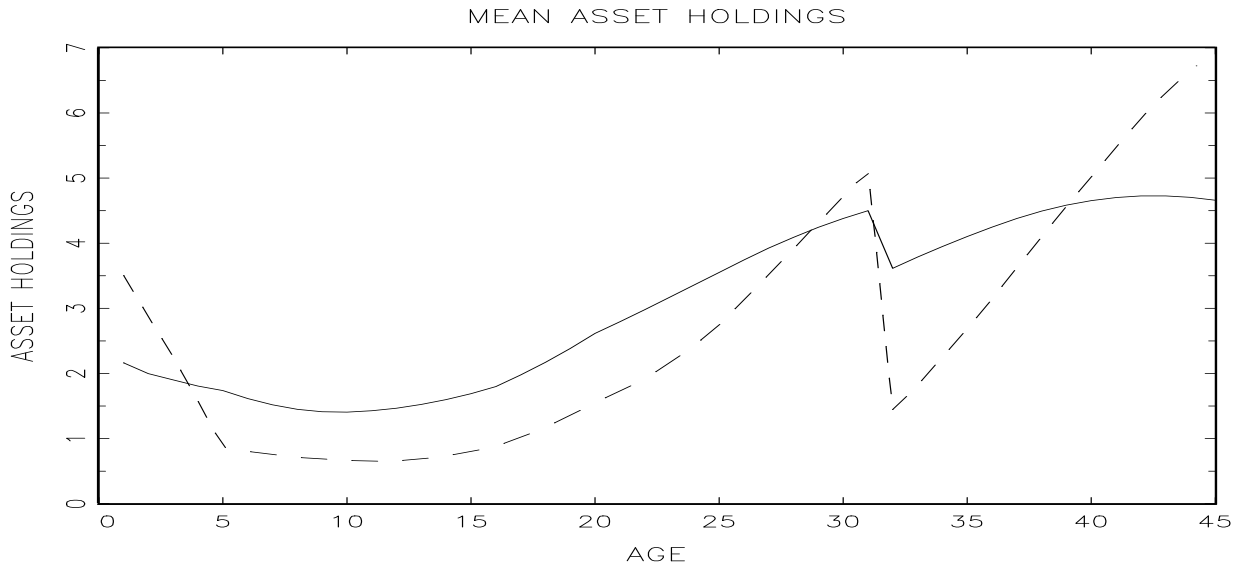


Figure 8: Mean Asset Holdings over the Life Cycle, Economy 1980, Baseline Calibration. Solid Line = High School Graduates, Dashed Line = College Graduates




Figure 9: Mean Asset Holdings over the Life Cycle, Economy 1980, No Transmission of Innate Ability Between Generations. Solid Line = High School Graduates, Dashed Line = College Graduates. Remark: Details of the calibration procedure and an interpretation are given in section 6.3

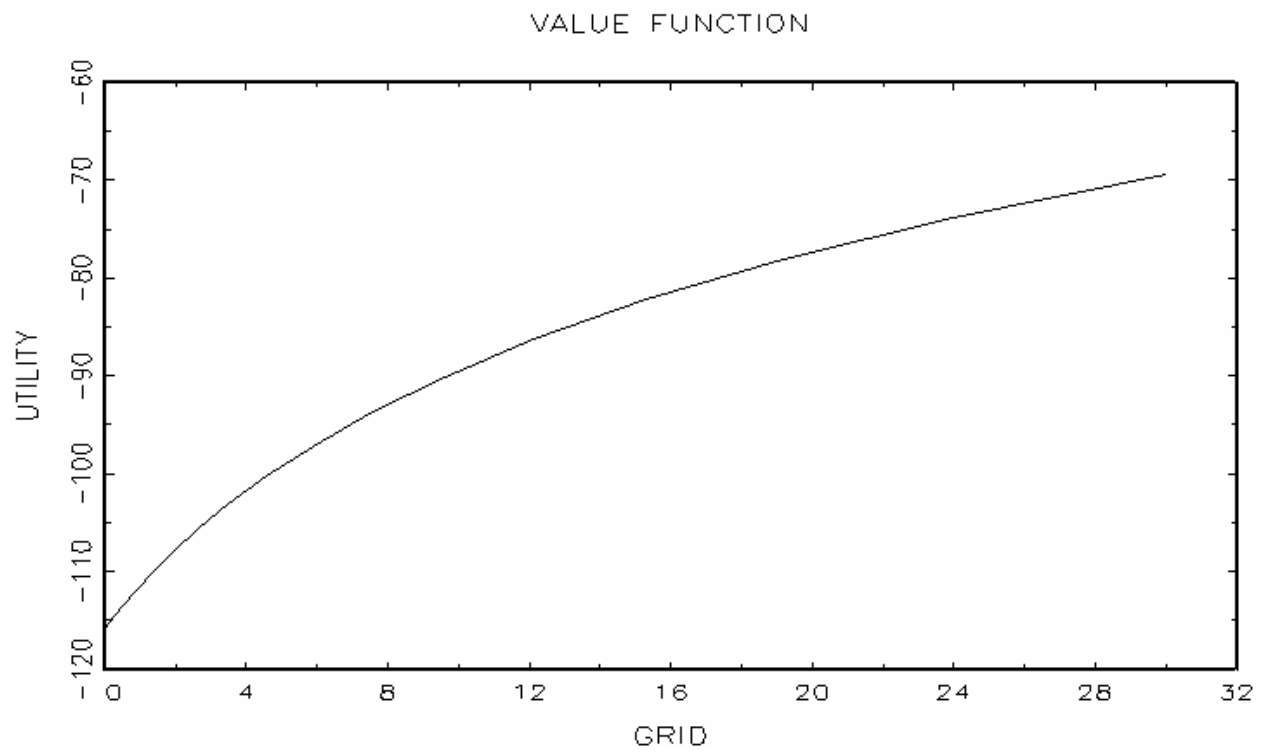


Figure 10: Value function for different levels of household wealth, after controlling for education, productivity and children's ability shock

11 Appendix IV: Tables

Table 1: Calibrated Parameters with Direct Empirical Counterpart for 'Economy 1980'

Parameter	Description	Value
ψ_j	survival probabilities	see text
α	capital share of income	0.36
δ	capital depreciation rate	0.08
γ	risk aversion	1.5
ε_j^e	life cycle earnings profile	see text
ρ^{hs}	earnings persistence high school	0.946
ρ^{hs}	variance shocks	0.015
ρ^{col}	earnings persistence college	0.955
σ^{col}	variance shocks	0.010
τ_K	capital income tax rate	0.2
rep	replacement ratio pensions	0.4
ν_0	max. share of tuition fees financed by subsidies	0.75
ν_{asset}	impact of capital income on college subsidies	0.2
ν_{asset}	impact of labor income on college subsidies	0.7

Notes: The source for all parameters describing the income process (ε_j^e , ρ^e and σ^e) is Hubbard et al. (1995). For detailed information regarding modifications and definitions, please refer to the main text.

Table 2: Parameters Calibrated Internally

Parameter	Description	Moment to Match in Data	Moment in Model	Value
β	discount factor	wealth/income: 3.1	3.1	0.96
ς	children's weight in parents' utility	intended transfers/income: 0.014	0.012	0.7
ϖ	transmission of innate abilities	correlation of education: 0.4	0.43	0.7
$\bar{\chi}$	upper limit loans	% of households with assets ≤ 0 : 25	25.5	1
$\bar{\kappa}$	mean cost of college	% of college graduates: 28	28.5	2.6
$\lambda_{1,hs}$	dropout probability, low innate ability, parents hs educated	dropout rate, 1st ability quartile: 0.34	0.27	0.52
$\lambda_{1,col}$	dropout probability, low innate ability, parents col educated	dropout rate, 2nd ability quartile: 0.28	0.23	0.52
$\lambda_{2,hs}$	dropout probability, high innate ability, parents hs educated	dropout rate, 3rd ability quartile: 0.19	0.22	0.25
$\lambda_{2,col}$	dropout probability, high innate ability, parents col educated	dropout rate, 4th ability quartile: 0.14	0.16	0.15
$\Delta_{1,hs}$	cost of college, low innate ability, parents hs educated	enrolment rate, 1st ability quartile: 0.21	0.11	1
$\Delta_{1,col}$	cost of college, low innate ability, parents col educated	enrolment rate, 2nd ability quartile: 0.35	0.35	0.7
$\Delta_{2,hs}$	cost of college, high innate ability, parents hs educated	enrolment rate, 3rd ability quartile: 0.68	0.60	0.2
$\Delta_{2,col}$	cost of college, high innate ability, parents col educated	enrolment rate, 4th ability quartile: 0.91	0.88	-1

Notes: Parameters calibrated internally, baseline calibration (economy 1980). The parameters $\bar{\kappa}$, $\lambda_{f,ep}$, $\Delta_{f,ep}$ are calibrated jointly. The data sources are as follows: Storesletten et al. (2004a) for wealth/income ratio, Gale and Scholz (1994) for intended transfers/income, Keane and Wolpin (2001) for the share of college graduates and the correlation of education across generations, Chatterjee and Ionescu (2010) for the dropout probabilities, and Belley and Lochner (2007) for the college enrolment rates. For more information regarding the construction of moments and definitions, please refer to the main text.

Table 3: Parameters characterizing the Economy 2000

Parameter	Description	Value 1980	2000
σ^{hs}	variance shocks	0.015	0.025
σ^{col}	variance earnings shocks	0.01	0.016
$\bar{\chi}$	upper limit on loans	1	1.25
$\bar{\kappa}$	average cost of college	2.16	5.2

Notes: This table compares the parameter values of the economy 1980 and the economy 2000. Important: the increase in the college premium is achieved by multiplying the age-earnings profile ε_j^e by 1.4. All other parameter values, not mentioned in this table, are identical in both the economy 1980 and 2000. For detailed information regarding modifications and definitions, please refer to the main text.

Table 4: *Average Enrolment Gap for Different Experiments*

Experiment	controlling for obs. ability	and for parental edu.
Economy 1980	17	7
Economy 1980, no transmission of ability	9	6

Notes: Table displays the average enrolment gap, that is differences in college enrolment rates between youth from families in the highest income quartile and youth from other families. The average enrolment gap is computed by weighting the gaps with the population share of the respective group. In all experiments, I control for observable ability. In the second column, I also control for parental education. In the second row, I shut down the transmission of academic ability between generations, as described in section 6.3

Table 5: *Enrolment Characteristics for Different Experiments*

Experiment	% enrolled	% graduates	% interest rate
Economy 1980	36	28	3.6
borrowing unrestricted	60	44	3.6
Economy 2000	49	36	2.7
borrowing unrestricted	67	47	2.7

Notes: The Table shows the total enrolment rate, the fraction of college graduates and the interest rate for different experiments. 'Economy 1980' refers to the baseline calibration, while the line 'Economy 2000' denotes the calibration 2000. The entries for 'borrowing unrestricted' refer to the respective experiments without borrowing limits. The share of households that is borrowing constrained in their college decision is thus given by the difference between the college enrolment rates.

Table 6: *College Enrolment Rates in %, Controlling for Parental Education, Economy 1980 % (Economy 1980, No Borrowing Constraints)*

	Ability Quartile 1	Ability Quartile 2	Ability Quartile 3	Ability Quartile 4
Income Quartile 1	7.4(49.4)	19.6(56.2)	21(64.5)	15.2(55.4)
Income Quartile 2	13.8(37.7)	35(57.3)	37.4(81.7)	27.3(55.1)
Income Quartile 3	11.3(28.2)	30.3(49.6)	32.4(76.2)	23.4(47.2)
Income Quartile 4	13.4(20)	38.8(54)	41.7(1)	29.5(50.0)

Notes: This table reports college enrolment rates in the economy 1980 for different subgroups of the population, after parental education has been controlled for. The values in parentheses refer to the experiment where borrowing limits are relaxed, see the main text for further details. The values confirm pattern that are visible by comparing Figures 4 and 5. See footnote 47 in the main text.

Table 7: *Share of College Graduates in Economy 1980 in % (Economy 2000)*

	Ability Quartile 1	Ability Quartile 2	Ability Quartile 3	Ability Quartile 4
Income Quartile 1	<1(<1)	<1(<1)	<1(<1)	<1(<1)
Income Quartile 2	<1(<1)	<1(<1)	<1(<1)	<1(<1)
Income Quartile 3	5.7(13)	25(40)	30(56)	57(70)
Income Quartile 4	46(99)	83(99)	87(99)	95(99)

Notes: Fraction of parent households with a college degree, in different subgroups of the population (income quartile in which parental income falls, quartile of the child's observable ability). <1 denotes "less than 1 percent". Values for the Economy 2000 are shown in parentheses.

Table 8: *Average Savings and Transfers of High School and College Graduates, Age 45-55 with at least one child (in 1983 US-Dollar)*

		no controls	controlling for income quartile	controlling for income & wealth quartile	controls as before & kids in college
Saving	High School (<i>Std.Error</i>)	13512 (8780)	23598** (6015)	35497** (8436)	-16817 (48000)
	College (<i>Std.Error</i>)	75014 (50457)	-51676* (27964)	-50254* (26985)	-79008 (61185)
	H_0 : Col=HS	Not Rejected	Rejected	Rejected	Not Rejected
Transfers	High School (<i>Std.Error</i>)	2547** (435)	1372** (489)	659 (707)	16279.03** (3259.39)
	College (<i>Std.Error</i>)	11678** (2192)	7211** (1925)	6559** (1857)	13409** (4204)
	H_0 : Col=HS	Rejected	Rejected	Rejected	Not Rejected

Notes: Data are from the 1983 and 1986 Survey of Consumer Finances. Variables are constructed as described in the text. Estimates are obtained from regressing transfers/saving on a set of dummies. The omitted categories are chosen as follows: high school graduates (column 3); high school graduates, 3rd income quartile (column 4); high school graduates, 3rd income quartile, 3rd wealth quartile (columns 5 and 6). SCF sample weights are used as probability weights in all specifications. I also experimented with other omitted categories, and found that results are similar.

* indicates statistical significance at 10% level

** indicates statistical significance at 5% level