



**University of
Zurich**^{UZH}

**Zurich Open Repository and
Archive**

University of Zurich
Main Library
Strickhofstrasse 39
CH-8057 Zurich
www.zora.uzh.ch

Year: 2014

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

Winter, Christoph

Abstract: In this paper, I analyze the determinants of college enrolment and the changes in these determinants over time. I propose a quantitative life-cycle model with college enrolment. Altruistic parents provide financial support to their children. Using counterfactual experiments, I find that 24 percent of all households are financially constrained in their college decision. Constraints become more severe over time. I show that my model is consistent with a narrow college enrolment gap between students from rich and poor families, as previously reported in the empirical literature. The estimation of enrolment gaps is a popular reduced-form approach for measuring the fraction of constrained households. My results suggest that these reduced-form estimates are misleading, and that a structural model of parental transfers is needed to correctly identify constrained households. Further, I show that parental transfers are an important driver behind the changing role of family income as a determinant of college entry, a fact that is well documented for the US economy.

DOI: <https://doi.org/10.1111/sjoe.12076>

Posted at the Zurich Open Repository and Archive, University of Zurich

ZORA URL: <https://doi.org/10.5167/uzh-99452>

Journal Article

Accepted Version

Originally published at:

Winter, Christoph (2014). Accounting for the Changing Role of Family Income in Determining College Entry. *Scandinavian Journal of Economics*, 116(4):909-963.

DOI: <https://doi.org/10.1111/sjoe.12076>

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

Christoph Winter[†]

April 2013

Abstract

This paper analyzes the determinants of college enrolment and changes in these determinants over time. I propose a quantitative life cycle model with college enrolment. Altruistic parents provide financial support to their children. Using counterfactual experiments, I find that 24 percent of all households are financially constrained in their college decision. Constraints become more severe over time.

I show that my model is consistent with a narrow college enrolment gap between students from rich and poor families previously reported in the empirical literature. Estimating enrolment gaps is a popular reduced-form approach for measuring the fraction of constrained households. My results suggest that these reduced-form estimates are misleading, and that a structural model of parental transfers is needed to correctly identify constrained households. Further, I show that parental transfers are an important driver behind the changing role of family income as a determinant of college entry, a fact that is well-documented for the U.S. economy.

Keywords: Dynamic General Equilibrium Models with Overlapping Generations, Parental Transfers, Parental Altruism, 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 No. 402. For comments on earlier drafts I would like to thank my advisor Morten Ravn, and my thesis committee: Alex Michaelides, Claudio Michelacci and Salvador Ortigueira. I have received many useful comments also from Ken Judd, Dirk Krueger, John Knowles, Ramon Marimon, Costas Meghir, Franck Portier, Víctor Ríos-Rull, David Scherrer, David Schönholzer, Kjetil Storesletten, Mariya Teteryatnikova, Gianluca Violante, Ken Wolpin and Fabrizio Zilibotti and from seminar participants at various institutions. I am grateful to the Economics Department at the University of Pennsylvania, where I wrote part of this paper, for its hospitality. All errors are mine.

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

I Introduction

This paper examines the determinants of college enrolment in the U.S. economy. I present a rich quantitative life cycle model in which altruistic parents invest in their children's college education. I use this model to shed new light on two long-standing questions. How large is the fraction of households that invest inefficiently in college education because of borrowing constraints? Has the share of constrained households changed over time?

These questions are of first-order importance for public policy. In this paper, I argue that both questions can be answered only within a model where the determinants of parental transfers are modeled explicitly.

A calibrated version of my model illustrates this point. According to the model, 24 percent of households in the U.S. economy were constrained in their college decision at the beginning of the 1980s. This result contrasts sharply with the findings presented in the previous literature. In an influential paper, Carneiro and Heckman (2002) conclude that at most eight percent of the population were constrained in their college decision. Carneiro and Heckman (2002) identify borrowing-constrained households according to the college enrolment gap between children from rich and poor families.

I show that my model is perfectly consistent with the finding of Carneiro and Heckman (2002). Applying the method of Carneiro and Heckman (2002) to the data generated by my model produces an estimate for the enrolment gap between students from rich and poor families that is in line with theirs. The discrepancy between the average enrolment gap and the findings from the counterfactual experiment suggests that the latter method does not correctly identify households that are financially constrained in their college decision.

I show that we can identify constrained households more accurately through a model in which parental transfers are generated endogenously. In the U.S., financial assistance from family (i.e., parents) is an integral part of a student's college financing. Gale and Scholz (1994) document that parental payments amounted to 35 billion U.S. dollars (hereafter USD) in 1986. Keane and Wolpin (2001) report that even students from poor families cover a substantial share of their total college expenditures with the help of parental transfers.

When modeling intergenerational linkages, I follow the classical papers by Becker (1974), Barro (1974), and Becker and Tomes (1979) and assume the existence of intergenerational borrowing limits that prevent parents from borrowing against their children's future earnings. In addition, I assume that parental transfers are generated by one-sided altruism. I follow Laitner (2001) and allow for imperfect altruism.¹ The presence of an intergenerational borrowing constraint, and the fact that altruism is only one-sided, imply that parents face a trade-off between saving for their own future consumption and providing transfers to their children. As a result, parents may find it optimal to underinvest in their children's college education. If children face a life-cycle borrowing constraint that prevents them from borrowing against their future earnings (Cunha and Heckman, 2007), the overall enrolment in college may be inefficiently low.

In the model, children are borrowing-constrained if the difference between the costs of tuition and the amount of parental support exceeds the borrowing limit. The difference depends on the

¹This means that the parent weights the utility of the child less than its own utility.

decision of children to attend college and on the parental transfers. Children decide to enrol in college based on tuition fees and on their ability to graduate. Tuition fees in turn are a function of ability and parental resources. The willingness (and the ability) of parents to provide financial aid is a function of their wealth, their income and their future earnings prospects. In essence, binding borrowing constraints are the outcome of a high-dimensional decision problem involving both children and parents.

Reduced-form approaches have trouble capturing the multidimensionality of the problem, given that many factors influencing the decision of parents and children are not observable to the econometrician. Moreover, since reduced-form estimates are based only on actual enrolment data, identifying borrowing constraints requires strong assumptions about the counterfactual enrolment rates. Using the average enrolment gap employed by Carneiro and Heckman (2002) and the counterfactual enrolment rates generated by my model, I show that the required assumptions are not satisfied in an environment where altruism is one-sided and intergenerational borrowing is prohibited.

The importance of modeling altruism becomes even more apparent when I analyze how the determinants of college enrolment changed over time. In the data, family income gained importance as a determinant of college entry between 1980 and 2000; see Ellwood and Kane (2000) and more recently Belley and Lochner (2007). At the same time, tuition fees doubled in real terms, while borrowing limits remained unchanged (Lochner and Monge-Naranjo, 2010). These facts together suggest that borrowing constraints became binding for a larger fraction of the population, resulting in wider enrolment gaps between children from rich and poor families. Between 1980 and 2000, the U.S. economy also experienced an increase in the college premium and a rise in (within-group) earnings; see e.g. Krueger and Perri (2006) and Heathcote et al. (2010).

Using the changes in the economic environment as inputs, I find that the model generates wider enrolment gaps. Moreover, the fraction of households that are financially constrained in their college decision increases from 24 to 28 percent. That is, the model suggests that bigger enrolment gaps indeed are associated with a larger fraction of constrained households.

However, my results imply that larger enrolment gaps do not necessarily indicate tighter borrowing constraints. In fact, I show that the widening of the enrolment gaps is the result of the rise in earnings inequality, and not a consequence of tighter borrowing constraints. This finding highlights the response of parental transfers as key for understanding the changing role of family income in determining college entry.

Despite the importance of family transfers, intergenerational linkages are not modeled in most quantitative studies of borrowing constraints. Notable exceptions are Keane and Wolpin (2001) as well as Johnson (2011), who estimate a parental transfer function in reduced-form. However, this approach implicitly assumes that parental transfers are invariant to policy changes. The present paper is closely related to the important recent work of Brown et al. (2012), who also assume one-sided altruism and intergenerational borrowing limits. They point out that almost all financial aid packages expect financial assistance from family. This expected family contribution "is, however, neither legally guaranteed nor universally offered" (Brown et al., 2012, p. 1). The authors document that about half of the children in the population are potentially constrained in the U.S., because their parents are relatively poor or selfish. This finding suggests that the results presented here should be seen as a lower bound. Compared to Brown et al. (2012), the quantitative model

I propose in this paper allows me to calculate the size of the bias that results if inference on the extent of borrowing constraints is made without explicitly modeling parental transfer decisions. Moreover, with the help of my model, I am able to study how changes in the economic environment affect the behavior of parental transfers. This question is also particularly interesting, given the widening of college enrolment gaps observed over time.

Building on these insights, a growing body of the literature incorporates structural models of parental transfers into quantitative life cycle models in order to study education policies. A recent example is the work by Abbott, Gallipoli, Meghir and Violante (Abbott et al., 2013, henceforth AGMV, 2013). AGMV (2013) also introduce an aggregate production function where different types of human capital are not (necessarily) perfectly substitutable.² With the help of their model, the authors compute the equilibrium effects of different policy interventions on optimal education decisions, inequality, and output.³ AGMV (2013) find that the response of parental transfers is quantitatively significant for understanding the impact of policy interventions.

Other examples of papers that incorporate altruism in the study of education policies are Restuccia and Urrutia (2004), Cunha (2007), Bohacek and Kapicka (2010), Holter (2013) and Caucutt and Lochner (2012). Some of these contributions focus on skill formation (e.g., Cunha, 2007) or on the relative importance of early versus late credit constraints, as e.g., Caucutt and Lochner (2012). Bohacek and Kapicka (2010) study the welfare effects of educational reforms. Both Restuccia and Urrutia (2004) and Holter (2013) are interested in the determinants of earnings persistence. Cunha (2007) as well as Bohacek and Kapicka (2010) assume two-sided altruism: families care about both 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 also children provide transfers to their parents. However, there is little evidence of this in the data, as argued by Gale and Scholz (1994) and Brown et al. (2012). Caucutt and Lochner (2012), Restuccia and Urrutia (2004) and Holter (2013) also assume one-sided altruism. Relative 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) also 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. 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.

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.

²Earlier papers analyzing the general equilibrium implications of education policies include Heckman et al. (1998) and Ábrám (2008), who examines wage inequality and education policy in a general equilibrium OLG model with skill-biased technological change.

³Garriga 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.

II 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 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.

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 ($j = 46$), and 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.

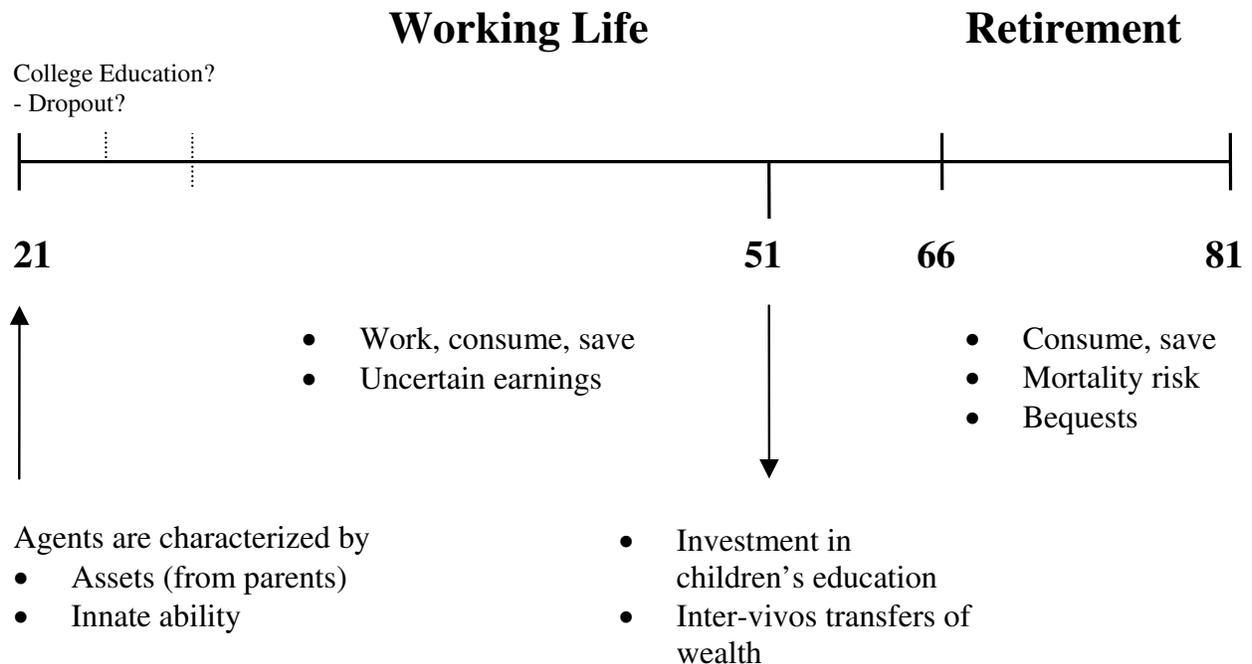


Figure 1: Life cycle and Generation Structure

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.⁴ 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 productivity of an agent with education level e :⁵

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

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,

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

where $N^e = \{\eta_1^e, \eta_2^e, \dots, \eta_n^e\}$ 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 U.S., as well as in other OECD countries, there is a significant fraction of students who enter college, but actually leave school without having obtained a degree (see e.g., Restuccia and Urrutia, 2004; Akyol and Athreya, 2005).

Upon entering college, students are required to possess 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 (2008). 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 (2012) 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 (2008), 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 do high school graduates. As a consequence, college investment is indivisible.

In what follows I discuss the formation of academic ability, the underlying force of behind college success, in greater detail.

⁴College graduates work for fewer periods; see below.

⁵I 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.

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 on earnings, other than through education.⁶ 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.⁷

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 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.⁸ Moreover, Carneiro et al. (2013) 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.⁹

Parental education e^p and innate ability 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) suggests 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 calibrate the values of λ that are associated with different combinations of 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.¹⁰ Since a single test score is unlikely to capture the full set of skills that is necessary for college success (Carneiro and Heckman, 2002), I distinguish between "academic" ability and "observable" ability in the following.¹¹

Let \mathcal{K} denote a random variable with realization κ_a . κ_a denotes the performance of a young

⁶Also see Heathcote et al. (2010)

⁷Cawley 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 that 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.

⁸See Black and Devereux (2010) for a comprehensive and recent review of the empirical literature.

⁹With 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).

¹⁰SAT scores and AFQT scores are highly correlated, see footnote 7 in Light and Strayer (2002).

¹¹Moreover, 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.

household in standardized test scores. Both κ_a and λ are known to children (and their 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 is perfectly negatively correlated with observable ability. That is, the cost of college is 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.¹² 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 U.S. 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 pay less in order to enrol in college.¹³

Individual aid packages can be based explicitly 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 note that most financial aid packages are implicitly based on student's merit, even if they are labeled differently.¹⁴

The following figures make this point clear. A 100 point difference in the SAT increases (need-based) grant aid between 500 and 2,300 USD (in 1996 prices) over the course of a college career (McPherson and Schapiro, 2006). Private institutions typically have more resources at their disposal and therefore offer more generous aid packages (Long, 2004; Collegeboard, 2010). These findings imply that a student from the lower end of the ability distribution has to pay between 3,000

¹²As 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.

¹³More 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).

¹⁴McPherson 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).

and 14,000 USD more for his college education than does a student at the upper end.¹⁵

The fact that the direct cost of college education is 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 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 are discussed in the next section. The (expected) return of attending college differs because more able students are more likely to complete their education.¹⁶

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 can a student whose parents are well off. I denote the fraction of direct college expenses (κ) that is covered by financial aid by v . The net direct costs of college attendance, after financial help is taken into account, are given by $(1 - v)\kappa(I(\kappa) > 0)$, where $I(\kappa)$ is an indicator function taking the value 1 if college expenses are positive.¹⁷ v is a function of parental income and parental assets, $v(y^p, a^p)$. 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. (2012), the expected family contribution that is calculated from the family's ability to pay is not legally guaranteed. Put differently, children cannot 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 the cost of college. Federal programs do not provide enough subsidized aid to meet the needs of all students, and most colleges are not committed to covering the entire residual (Dick and Edlin, 1997).

¹⁵These figures are based on the assumption that students at the lower end of the distribution have a SAT score below 700, while students at the upper end have a score above 1300. According to Chatterjee and Ionescu (2012), 8 percent of the students who took the SAT score below 700, while about 15 percent score higher than 1300.

¹⁶There are other ways to generate heterogeneity. In important contributions, Gallipoli et al. (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 extent 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 1,000 USD (in 1998 prices) 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).

¹⁷Recall that negative direct costs from college may occur if the measured ability of the student is very high and the student receives a scholarship.

Assets and Loans. Households accumulate savings in the form of assets a_j . They borrow by taking out loans χ . Consequently, $a_j \geq 0$ and $\chi \geq 0$. Borrowing can take place only at the beginning of the life cycle, either to finance college education, or 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. Gallipoli et al. (2010). Moreover, installment loans are the most sizable component of unsecured consumer debt.¹⁸ 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 when credit is most needed.¹⁹

I assume that the terms and conditions for χ follow those for U.S. 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.²⁰ All loan programs have in common that participants must repay their loans in full, regardless of whether they complete college successfully (Chatterjee and Ionescu, 2012). 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).²¹ Under the standard 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.²² As an approximation, I assume that all borrowers repay their loans in 25 years, starting from the year in which college students graduate.²³

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

¹⁸Using 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 the most important form of unsecured consumer debt. Installment loans appeared not only to be more frequent, but also to be bigger. The median amount of installment debt carried by households was 1,600 USD, whereas it was just 600 USD in form of credit card debt (in 1989 dollars). More recent waves of the SCF confirm this pattern.

¹⁹Kaplan 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 after age 22 is only 4 percent.

²⁰At 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.

²¹Unsubsidized loans accrue interest over the deferment period. Accrued interest payments are added to the loan principal. For subsidized loans, the government covers the interest on loans while students are enrolled (Lochner and Monge-Naranjo, 2010).

²²<http://studentaid.ed.gov/repay-loans/understand/plans> (retrieved on February 9, 2013)

²³Since death does not occur before the age of 53 ($j=33$), no agent dies in negative net worth.

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 household's 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 is 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.²⁴

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) and Brown et al. (2012).²⁵ 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

²⁴Keane 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 5,000 USD. At age 25, i.e. shortly after college, 16 percent of the group with negative net worth held debt of more than 10,000 USD and 20 percent less than 1,000 USD, 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 USD, according to Keane and Wolpin (2001)

²⁵I allow for the fact that altruism may be imperfect.

is the same as if parents decided themselves on the optimal split between education investment and pure financial transfers.²⁶

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 the 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 however children 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, 2013). 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 assume that all intended transfers resulting from parental altruism are made at the beginning of the child's life cycle, when 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 v , and subsidies on the interest rate r . The government adjusts τ_w in order to balance its budget in every period.

III Recursive Problem

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

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 $j \geq 33$.

²⁶In the following, I will use the terms "financial transfers", "inter vivos transfers" and "financial inter vivos transfers" interchangeably.

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

$$a_{j+1} = (1 + r(1 - \tau_k))a_j + pen - c_j,$$

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

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:

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

where $Q^e(\eta^{j,e}, \eta^{j+1,e})$ 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

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

where ζ is the intergenerational discount factor. I allow for imperfect altruism, that is, $0 \leq \zeta \leq 1$. If $\zeta = 0$, parents care only about their own utility. The model thus nests a pure life cycle economy ($\zeta = 0$) and a dynastic model ($\zeta = 1$) 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 its economic life. It depends on $s_0 = (tra, \kappa, f, e^P)$. $V_0(s_0)$ is described in greater detail in the next section.

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

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

Notice that the innate ability level of the child f and 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 and the likelihood of dropping out of college.

The parental budget constraint is given by

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

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

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

Young households

Young households are of age $1 \leq j \leq 30$. The problem of a young household depends on the status of its parents. Children with parents who are living expect to receive bequests in the future, and thus need to keep track of their parents' wealth holdings. In the following, I distinguish between young households with deceased parents and those with living parents. Since parent household do not die before age $j = 33$, the problem of a young household with deceased parents starts at the age of $j = 3$.²⁷

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

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

²⁷This 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 $j = 2$.

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

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

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

College students ($e=\text{col}$) do not receive income while enrolled, 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 ($l_j = 0$), independent of the education level.

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

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

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

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

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 it obtains if the parent continues to live for another period. Since the age gap between parent and child households 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:

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

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 I discussed above, I assume that dropping out

occurs after two years of college. Hence, at $j = 2$, the recursive problem of college students reads as follows:

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

where λ_{f,e^p} is the probability of dropping out from college without a degree. λ depends the innate ability 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 in the current period. The continuation value is as stated above for the young households with deceased parents.

The budget constraint depends on whether the parent household 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 enrolled and $\varepsilon_j^{e=col} = 0$ for $1 \leq j \leq 4$. Moreover, $l_j = 0$ if $1 \leq j \leq 4$, independent of the education level.

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

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

Young Households, College Investment Decision. At the beginning of their economic life, young households receive transfers tra from their parents, and decide whether to attend college and how much money to borrow. I assume that these events take place before the realization earnings shock and before consumption and saving decisions are made in the first period.

The college investment decision is based on the cost of college κ as well as on the innate ability f and parental education e^p . Because the productivity shocks $\eta^{j,e}$ depend on the education level, the decision whether to attend 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 have taken other economic actions. The expectation is taken on with respect to the set of possible productivity shocks in the first period.²⁸

The constraints of this problem are as follows:

²⁸I assume that productivity shocks in the first period are equally likely.

$$a_1 = tra + \bar{\chi} - \kappa|_{e=col} v(a_{p,1}, \varepsilon_j^e \eta^{31,e} w)$$

$$a_1 \geq 0$$

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

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

Young households take parental transfers tra as given. All young households borrow up to the maximum debt limit, which is $\bar{\chi}$.²⁹ Based on the cost of college (net of subsidies) κv , households decide whether to attend college. The optimal education choice for an individual who faces the state vector s_0 is described by $edu(s_0)$. The residual financial resources a_1 are kept for future consumption.

IV 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:

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

1. *Given prices and policies, the value functions $V(s)$ are the solutions to problems (1)-(7). The functions $\{edu(s), c(s), a(s), tra(s)\}$ 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:*

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

²⁹As outlined above, the assumption that all households borrow up to the limit is inconsequential, since the interest rates on borrowing and lending are the same.

where the total amount of college subsidies Ξ is given by

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

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:

$$\begin{aligned} K &= \int_S a(s) d\Phi \\ L &= \int_{S^{j \leq j^{Work}}} \epsilon_j^{hs} \eta^{j,hs} d\Phi + \int_{S^{j \leq j^{Work}}} \epsilon_j^{col} \eta^{j,col} d\Phi \\ C + \delta K + I &= F(K, L) \end{aligned}$$

where

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

6. The Aggregate Law of Motion is stationary:

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

The function H is generated by the policy functions $a(s)$, $c(s)$, $\text{tra}(s)$, $\text{edu}(s)$, the transition matrix of productivity shocks, $Q^e(\eta^{j,e}, \eta^{j+1,e})$, 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 households' distributions 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 parent households actually leave. I present more details about the computational procedure in the appendix.

V Parameterization and Calibration

Experiment Design. Before turning to a detailed description of my calibration strategy, I outline the design of my quantitative experiments. I conduct two main experiments. The aim of the first experiment is to calculate the fraction of households that would have enrolled in college at the beginning of the 1980s, all other things equal, if borrowing limits on student loans had been absent. In order to implement this experiment, I first compute the equilibrium allocation generated by a version of the model which is consistent with key statistics of the U.S. economy at the beginning of the 1980s. More specifically, I target the size of parental transfers, the college enrolment rate, the dropout rates and correlation of college education across generations. The resulting allocation is labeled as benchmark calibration or "economy 1980". In the appendix, I show that the economy 1980 is also consistent with a series of other empirical regularities related to college enrolment behavior in the U.S., which are not used as targets in the calibration procedure.

In a second step, I measure the fraction of households that enroll in college if I remove the borrowing limit on student loans. The fraction of constrained households is given by the difference in the enrolment rate of two economies. The first is the benchmark economy without the borrowing limit on student loans. The second is the benchmark economy where the borrowing limit on student loans is in place. I assume that households do not anticipate the regime change. As a consequence, I keep the wealth distribution of parents at age 51 constant across the two experiments. Parents cannot adjust their saving decisions prior to the relaxation of the borrowing limit. Note that they can adjust their transfer behavior after the regime change. Moreover, I hold aggregate prices fixed, in order to make sure that households make their college decision in the same economic environment.

I conduct a second experiment in order to analyze how the increase in the skill premium, greater residual earnings inequality, a rise in tuition fees, and a decline in the real interest rate affect the fraction of borrowing-constrained households. Again, I assume that the changes in the economic environment are not anticipated, as in Heathcote et al. (2010). This means that I must keep the wealth distribution of parents at the age of 51 constant with respect to the economy 1980. Households conduct the saving decisions that give rise to this distribution under the assumption that the environment of the economy 1980 persists forever.³⁰ I label the resulting allocation as "economy 2000". The computation of the constrained households is equivalent to that in the economy 1980. Finally, the development in the fraction of constrained households over time can be computed by comparing the fraction of constrained households in each of the economies 1980 and 2000.

In the following, I describe the calibration procedure in more detail.

Economy 1980

I distinguish between parameters that are set outside the model and others that are calibrated internally. Table 1 summarizes the first set of parameters, while table 2 gives an overview of the second set.

³⁰I also conducted experiments where I allow the distribution of parents to adjust. See Footnote 45 for a brief summary of the major changes.

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 than do high school graduates. The skill-premium is independent of the fraction of high-skilled labor.³¹ I set the capital share in income (α) equal to 0.36, as estimated by Prescott (1986). Following Imrohroglu 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 U.S. population. Only values between age 53 and age 80 are used. I assume that the survival probability is zero for agents at age 81. The survival probability for households younger than 53 years is set equal to 1.³² 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 within 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).

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.³³

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 U.S. 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 bigger change for high school graduates and a somehow smaller 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

³¹The implicit underlying assumption is that efficiency units supplied by high school and college graduates are perfect substitutes in the production process. AGMV (2013) allow for imperfect substitutability.

³²The actual survival probability before 53 is close to 1. See the National Vital Statistics Report.

³³It 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).

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, in line with findings from other empirical studies. By estimating a fixed-effect model, Gallipoli et al. (2010) propose an alternative way of calibrating the earnings process.³⁴

Pension Benefits. Pension benefits are calculated using the following formula:

$$pen = rep(1 - \tau_w)w \times \text{average lifetime efficiency hs graduates}$$

where rep denotes the replacement rate, which I set equal to 0.4, following De Nardi (2004). The average lifetime productivity of a high school graduate follows from HSZ. Pension benefits thus amount to 40 percent of the average net wage of a high school graduate. Recall that τ_w and w are determined in equilibrium.

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 U.S. 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:

$$v = \max(\{v_0 - v_{asset}(\max(0, a_{p,w}^{31} - \bar{a}) + (1 - \tau_k)ra_{p,w}^{31}) - v_{labor}(1 - \tau_w)\varepsilon_j^e\eta^{j,e}w\}, 0) \quad (8)$$

where $v_0, v_{asset}, v_{labor} \geq 0$. 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: v_0, v_{asset} and v_{labor} . v_0 determines the maximum fraction of college expenses that can be covered by the subsidy. The parameters v_{asset} and v_{labor} determine how fast this fraction decreases with asset income and labor income, respectively.

I set v_0 to 0.75. This decision is based on the observation that in 1979-1980, even students from families in the bottom income quintile financed a significant amount of their college expenses with

³⁴I 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.

the help of parental transfers. See the evidence cited in Keane and Wolpin (2001).³⁵

Feldstein (1995) reports that the implied marginal tax rates on (adjustable) income lie between 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 v_{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 U.S. system.

In the data, the share of college expenses financed with parental transfers increases rapidly in the level of total parental income (Keane and Wolpin, 2001).³⁶ This suggests that college subsidies actually are falling sharply in parental income.³⁷ In order to account for this feature of the data, without overly distorting the saving decisions at the same time, I choose different values for v_{labor} and for v_{asset} . More specifically, I assume that v_{labor} takes a value of 0.7.

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 percent wealth quantile in the U.S., which is 3.1 (see Storesletten et al., 2004a). The result is a β of 0.96. The implied interest rate is 3.6 percent per annum, which is in the range commonly reported in the literature.

Children's Weight in Parents' Utility ζ . The parameter ζ governs transfer behavior in the model. It is the intergenerational discount factor and determines the relative weight that parents assign to their children's utility in their own utility function. In order to calibrate ζ , 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 the 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 therefore quantitatively more important than both other categories combined. 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.

³⁵In 1979-80, for youths with families in the bottom income quintile, about 19 percent of college expenses were covered by parental transfers. The rest was financed with other internal sources (such as student's income, from which I abstract) or scholarships, grants or loans. See Keane and Wolpin (2001). Notice that in my specification, college loans do not count as college subsidies v . Hence, some students may be able to finance their total college expenses with the help of external funds.

³⁶Keane 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.

³⁷In this context, it is important to note that the financial aid that a student 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 needs of all students, and most colleges are not committed to covering the entire residual.

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.³⁸ The resulting value for ζ is 0.7. A ζ 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 U.S. economy.

Upper Limit on 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 ζ , the borrowing limit for loans, and the college subsidies I can now 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} , and 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.

\mathbf{f} takes two realizations, $\mathbf{f} \in \{low, high\}$. This results in four combinations for $\Delta_{\mathbf{f}, edu^p}$ and for $\lambda_{\mathbf{f}, e^p}$. 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 to make the model 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

³⁸Recall 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.

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 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 (2012).³⁹

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 than does parental education. I will make use of this assumption in the next section also to calibrate Γ .

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

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 their parents' type. 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 too, compared to high school educated parents. This suggests that the transmission of f is somewhere in between the two extremes. I therefore 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 relevant to the college enrolment decision. In particular, I document

³⁹Chatterjee and Ionescu (2012) use the students' SAT score as a measure of academic ability. Depending 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 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 forth effort but simply enroll and drop out shortly after. Since I target the dropout rates for four different ability quartiles, I interpolate the data from Chatterjee and Ionescu (2012) to compute in-between values.

that (i) the distribution of measured ability is skewed to the right, as in the data (see e.g. Gallipoli et al., 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 (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 too. More 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.

Economy 2000

Between 1980 and 2000, the U.S. economy was characterized by an increase in the skill premium, higher residual earnings inequality and a rise in tuition fees. Since I am interested in a one-generation response to these changes in the economic conditions, I hold the distribution of parents constant with respect to the economy 1980. This implies that aggregate prices r and w must be set exogenously, as they cannot be determined in general equilibrium. In this section, I discuss how I adjust the parameter values with respect to the economy 1980 such that the model is consistent with the changes in the economic environment between 1980 and 2000.

Increase in College Premium. Heathcote et al. (2010), in 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 lines up 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, see also 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 savings for college investment to be undertaken in the late 1990s. Clearly, the variance of residual earnings continued to increase during the 1990s as well (see e.g. Krueger and Perri, 2006; Heathcote et al., 2010), so my choice is rather conservative.

Decline in the Real Interest Rate r . The real interest rate in the U.S. declined during the 1990s by about 1 percentage point, see Caballero et al. (2008) and Caporale and Grier (2000). I thus set $r = 0.0257$, compared to $r = 0.0357$ in the economy 1980.

Wage Rate w . The wage rate w is computed from r , using the fact that firms are price takers and maximize profits. w then follows from the first-order conditions of the firms problem. In order to calculate w , a value for aggregate labor in efficiency units L needs to be imputed as well. L is determined by the education choice, which in turn depends on w . I start with an initial guess for L , compute w , solve for the resulting number of college graduates using the optimal decisions and the distribution of parents given from the economy 1980, and check whether the resulting L is

consistent with the initial guess. The L resulting from this procedure is 0.5, with wage rate w at 1.28, compared to 1.21 in the economy 1980. Notice that total income in the economy, given by $Y = rK^D + wL$, where K^D denotes total net worth of households, is roughly constant compared to the economy 1980.

This is important for my next steps. Here, I interpret tuition fees and the borrowing limit in real terms in order to compare their development over time. Since aggregate income is roughly constant, this means that the nominal values are equal to the real values.

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.

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). I thus keep $\bar{\chi}$ at its 1980 value of 1.

Other parameters I also need to assign values to the labor tax rate τ_w and pension benefits pen . I keep τ_w constant with respect to its value in the economy 1980 ($\tau_w = 0.053$). Given w and τ_w , pen is found by using the replacement ratio as outlined above.

All other parameters, in particular the preference parameters, are kept unchanged. The parameter choices for the economy 2000 are summarized in Table 3.

VI Results

In this section, I analyze the role of borrowing constraints in the economies 1980 and 2000. In particular, I show the following key results. First, borrowing constraints prevent a large fraction of households from attending college (24 percent). Second, this fraction increased over time (from 24 to 28 percent). Third, reduced-form estimates of the share of constrained households, based on differences in enrolment rates between children from rich and poor households, lead researchers to misinterpret the role of borrowing constraints. And last but not least important, the model outlined above offers interesting insights for the changing role of family income in determining college entry.

Economy 1980

Borrowing Constraints. In order to compute the fraction of households that would have enrolled in college, if access to credit had been unlimited, I increase the upper limit on loans, $\bar{\chi}$ from 1 to 5.⁴⁰ I hold all other parameters, including the parental distribution as well as aggregate prices r and w , constant.

⁴⁰This choice ensures that a large fraction of households can cover their tuition fees. Higher values of $\bar{\chi}$ may even lead to a larger enrolment rates, however, they are also associated with numerical instabilities. My choice of $\bar{\chi}$ works in all experiments and thus ensures comparability.

The rise in college enrolment following the removal of borrowing limits is substantial: the fraction of college students increases from 36 percent to about 60 percent. This implies that 24 percent of the population are borrowing-constrained in their college decision in the benchmark calibration.

Below, I compare the results to those obtained in the empirical literature by using reduced-form approaches. Here, binding borrowing constraints are commonly identified by larger effects of family income on college enrolment, see e.g. Ellwood and Kane (2000), Carneiro and Heckman (2002) or Belley and Lochner (2007).

Parental Income as a Determinant of College Enrolment in the Data and in the Model.

In assessing the impact of borrowing constraints, a common approach is to analyze family income effects. When doing so the literature carefully controls for pre-college ability, e.g. by conditioning AFQT scores and by controlling for parental education and other proxies. This control is important, since family income and ability are likely to be correlated.

Table 4 depicts the estimated family income effects on college attendance at age 21 by AFQT quartiles. The table is based on information in Belley and Lochner (2007), Table 4.

Table 5 shows the corresponding income effects, by quartiles of measured ability, that are generated by the model. Following Belley and Lochner (2007), I also condition on test scores and control for parental education.⁴¹ I compute family income effects using the stationary distribution generated by the benchmark calibration.

A visual inspection of Tables 4 and 5 reveals that family income effects in the model are qualitatively in line with the data. For most ability/family income quartiles, the deviations between model and data are smaller than one standard error. An exception are children with an AFQT test score that falls into the second quartile. Here, the model performs poorly, because it predicts a positive relationship between family income and college enrolment, whereas this relationship is negative in the data. More future work will be necessary in order to understand why college enrolment is increasing with family income for most children, but not for those in the second lowest ability group.

The fact that the model implies a stronger role of family income is potentially problematic, since it may indicate that the role of borrowing constraints is overstated. In order to assess whether this claim is true, I drop young households with test scores in the second ability quartile when computing the fraction of borrowing-constrained households. As a result, the fraction of constrained households falls from 24 to 16 percent and thus remains twice as high as suggested by the estimates in the previous literature. This is discussed below.⁴²

⁴¹Belley and Lochner (2007) control for mother's education by incorporating dummy variables for high school graduation and college attendance. Since I do not model high school dropouts and assume that college dropouts are identical to high school graduates, I only distinguish between high school and college graduates and control for college graduation.

⁴²The large fraction of potentially constrained households may also appear to be puzzling against the finding of Lochner and Monge-Naranjo (2010), who find that only a few percent of college students borrowed up to the maximum of what they could from the Stafford loan program at the beginning of the 1980s. However, Brown et al. (2012) show that a large fraction of parents are unable or unwilling to meet the expected family contribution, which implies that many children did not even attempt to enroll in college because they were lacking financial support from their parents.

In the next section, I follow the method of Carneiro and Heckman (2002) and compute the average income effect. Carneiro and Heckman (2002) interpret this statistic as the fraction of the population that is financially constrained in their college decision. I show that the average enrolment gap between children from rich and poor families computed with the help of model-generated data coincides with the empirical counterpart which Carneiro and Heckman (2002) calculated using the NSLY79.

Reproducing Carneiro and Heckman (2002). Carneiro and Heckman (2002) present a method for computing the fraction of the population that is financially constrained in their college decision. Their method has become the standard tool used 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 through 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. The authors argue that family resources are likely to produce a number of skills that a single test score cannot capture. 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 (2002) therefore 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 constrained in their college decision.

In order to apply their method to the data generated by my model, I first compute the college enrolment rates for all family income/measured ability quartiles. I control for parental education by focusing on children whose parents have a high school degree only.⁴³ In a second step I calculate, for each ability quartile, the enrolment gap as the difference in the college enrolment rate between children from the top income quartile and children from the bottom income quartile. Finally, I compute the average enrolment gap by using the population weights that are implied by the equilibrium distribution of my benchmark economy. The average enrolment gap indicates that a fraction of 6.2 percent of the population in the benchmark economy is financially constrained in their college decision. This result is perfectly in line with the estimate of Carneiro and Heckman (2002) from the NLSY79.

Strikingly, the fraction of borrowing-constrained households resulting from the counterfactual experiment is more than three times as large as suggested by the average enrolment gap.

⁴³This is done for consistency with the rest of my analysis. See Footnote 41. For the same reason, I also group households according to quartiles, and not terciles, of measured ability.

In the following subsection, I show that the source of this discrepancy is the fact that the average enrolment gap identifies borrowing-constrained households only under certain assumptions, which are not satisfied if parental transfers are generated by one-sided altruism and if intergenerational borrowing is prohibited.

Identifying Borrowing-Constrained Households. The procedure of Carneiro and Heckman (2002) identifies the fraction of borrowing-constrained households only if the following assumptions are satisfied:

1. In the absence of financial constraints, college enrolment rates are equal for all income groups (after controlling for ability by using test scores and other proxies, e.g. parental education).
2. Youth with parents in the top income quartile are not financially constrained in their college enrolment decision.

Put differently, the first assumption implies that differences between family income groups reported in Table 5 should disappear when borrowing constraints for college education are removed. Clearly, this is not the case, as Table 6 shows. According to Table 6, children from poor families are more likely to enrol in college, compared to children from rich families, in the economy 1980 without borrowing constraints. As a result, the measured impact of family income on college enrolment becomes negative.

The fact that children from poor families are more eager to enrol in college is a consequence of the college subsidy system. The costs of attending college are lower for youth from poor families, because they receive financial aid. Since I control for parental education in Table 6, differences in tuition are the only source of heterogeneity affecting college decisions of young households. In particular, other determinants of college enrolment, such as the dropout risk, should be approximately equal for children within the same ability quartile.⁴⁴

Clearly, violations of the first assumption have a significant impact on the total number of constrained households. This is the message of Table 7, which reports the fraction of constrained households per family income quartile. A large fraction of children with poor parents are borrowing-constrained. They account for almost 50 percent of the total fraction of households that is financially constrained in their college decision.

The ultimate reason behind binding borrowing constraints is the fact that parents are unwilling or unable to provide enough support. Because intergenerational borrowing is not permitted and altruism is only one-sided, parents face a trade-off between consuming for themselves and transferring resources to their children. As Table 7 makes clear, all parents face this trade-off, even those at the top of the income distribution. According to Table 7, two percent of those young households with rich parents are borrowing-constrained in their college decision. This shows that the second assumption also is violated. The next subsection sheds further light on the determinants of parental transfers and the parental trade-off.

⁴⁴Of course, there are differences in the amount of parental transfers, as I argue below.

The Role of Altruism. I now describe the results of an experiment in which I set the degree of altruism, ζ , from 0.7 to 0.75. All other parameters, including the distribution of parents, are held constant with respect to the economy 1980. A higher ζ means that parents assign greater weight to their children's utility. Put differently, parents become more willing to provide transfers to their children. This alleviates the impact of borrowing constraints on college enrolment.

There are two important points to learn from this experiment. First, it further illustrates the fact that reduced-form regressions are unsuitable for identifying borrowing constraints. I find that the fraction of households that is financially constrained in their college education decision drops from 24 percent to 22 percent. Strikingly, however, the average enrolment gap, which I computed following Carneiro and Heckman (2002), increases from 6.2 to 7.2, thus suggesting that borrowing constraints have tightened, while in fact the opposite is the case.

Second, the above experiment also improves our understanding of the determinants of parental transfers. The combination of one-sided altruism and intergenerational borrowing limits introduces a trade-off in the parental optimization problem between transferring resources to children or saving for future consumption. In other words, there is a trade-off between smoothing out consumption across generations and smoothing out consumption over the life cycle. Changing ζ provides an opportunity to illustrate some of the factors that influence this trade-off.

A comparison of Tables 7 and 8 reveals that the decline in the number of constrained households is entirely due to a fall in the fraction of constrained children with poor parents. Hence, even when everyone is more altruistic, implying that parents are more willing to share resources with their children, this is not sufficient to change the behavior of the rich.

One might have expected the opposite. However, due to the redistributive nature of the pension system in the model, rich parents need to save more resources in order to smooth out their consumption over the life cycle. Thus, the trade-off between transfers and own saving is tighter for the rich, all other things equal. In the next section, I argue that this trade-off is key for explaining the changing role of family income in determining college entry.

Economy 2000

Over the last decades, the U.S. economy has experienced a sharp increase in tuition fees and in earnings inequality, both between and within education groups. At the same time, borrowing limits remained constant in real terms and the real interest rate declined by 1 percentage point.

In this subsection, I use my model to analyze the impact of these changes on the link between college enrolment and family income, as well as on the development of the fraction of borrowing-constrained households over time. I start by briefly reviewing the stylized facts.

Empirical Evidence. Ellwood and Kane (2000) and Belley and Lochner (2007) present evidence showing that the impact of family income as a determinant of college enrolment has become more important over time. This can be seen by comparing Tables 4 and 9, which are based on results from Belley and Lochner (2007).

In addition, Lochner and Monge-Naranjo (2010) report that borrowing limits for student loans remained constant while tuition fees doubled (in real terms). This means that the ratio of borrowing opportunities to tuition fees decreased over time. Borrowing limits became tighter in real terms.

Together with the changing role of family income as determinant of college entry, this suggests that borrowing constraints are now binding for a larger fraction of the population, relative to the beginning of the 1980s.

I use my model to shed more light on this point. Before doing so, I describe the development of the enrolment gap in the model.

The Changing Role of Family Income in the Model. Table 10 shows the impact of family income in the economy 2000. A comparison with the income effects in the economy 1980 (see Table 5) reveals that family income is more important in the economy 2000, relative to the economy 1980. Qualitatively, this is in accordance with the data. Quantitatively, however, the model overstates the impact of family income in the economy 2000, in particular for youth with high levels of measured ability.

In order to summarize the impact of family income on college enrolment with the help of a single number, I follow the procedure by Carneiro and Heckman (2002) and compute the average enrolment gap as outlined above. I find that the average enrolment gap is 15.7 percent in the economy 2000, relative to 6.2 percent in the economy 1980. Clearly, family income is more important in determining college enrolment in the economy 2000.

The Development of the Fraction of Constrained Households. The fraction of households that is financially constrained in their college enrolment decision is 28 percent in the economy 2000. Hence, the share of constrained households increased over time, from 24 to 28 percent. Thus, in the model, the growing role of family income as a determinant for college entry is associated with a higher fraction of constrained households.

I compute the share of constrained households by adjusting the limit on student loans, $\bar{\chi}$, from 1 to 5, and by measuring the associated change in the overall college enrolment rate. All other parameters are held constant, including the aggregate prices r and w and the distribution of parents at age 51.

Are Tighter Borrowing Constraints Responsible for the Changing Role of Family Income in Determining College Entry? I now analyze whether the fact that borrowing limits became tighter, relative to the costs of tuition, is responsible for the increase in enrolment gaps between students from rich and poor families and for the associated increase in number of borrowing constrained households.

I design an experiment in which borrowing limits on student loans rise in accordance with tuition fees. Since average tuition fees doubled between 1980 and 2000, I raise $\bar{\chi}$ from 1 to 2. All other parameters, including the distribution of parents at age 51, remain constant.

According to this experiment, tighter borrowing limits account only partly for the changing role of family income. If $\bar{\chi}$ is increased to 2, the average enrolment gap drops from 15.7 percent to 10.5 percent. Hence, the enrolment gap is still a lot larger than the 6.2 percent measured for the economy 1980.

At the same time, the relative tightening of borrowing limits seems to be responsible for an increase in the fraction of constrained households over time that is implausibly large, given the overall change in the fraction of constrained households between the economies 1980 and 2000. Instead of 28 percent, only 18 percent of households are constrained if I increase $\bar{\chi}$ from 1 to 2.

Together, these two observations suggest that tighter borrowing constraints are not the sole explanation for the increase in the fraction of constrained households and the growing importance of family income as a determinant for college enrolment.

Changes in the Distribution of Constrained Households. According to the model, not only the number of constrained households changed, but also their share in each family income quartile. This can be seen by comparing Tables 7 and 12. In 1980, 47 percent of children with poor parents are constrained. In the economy 2000, only 17 percent of children in the bottom family income quartile are constrained. At the same time, the fraction of constrained households in the top income quartile rose from 2 percent to 16 percent. Similarly, among children with parents in the second highest income quartile, the fraction of constrained households increased from 19 to 32 percent.

This shift in the distribution is difficult to reconcile with a theory in which tighter borrowing limits are the sole driver behind the increase in the fraction of borrowing-constrained households. In this case, the rise in the number of constrained households would stem from a larger fraction of constrained households at the lower end of the income distribution.

Parental Transfers and the Changing Role of Family Income. Understanding the causes behind the changing role of family income and the larger number of borrowing-constrained households has significant policy relevance.

Here, I argue that parental transfers are key for explaining these changes as well as the shift in the distribution of constrained children.

The development in the aggregate size of transfers is striking. While tuition fees doubled in real terms, the ratio of parental inter-vivos transfers relative to output falls from 3.6 percent in the economy 1980 to 3.3 percent in the economy 2000.⁴⁵ This suggests that, on the aggregate level, parental support is not sufficient to cover the increase in tuition fees. As a consequence, more students need to rely on loans to finance their studies. The fraction of households in debt increases from 22 percent in 1980 to 24 percent in 2000.

Parental transfers shrunk in real terms, because the increase in earnings inequality raises earnings uncertainty. This means that parents become more reluctant to transfer resources to their children, as they shift to accumulating more precautionary savings.

The increase in earnings inequality also worsens the position of families at the lower end of the distribution. In the economy 2000, poor families have less income in absolute terms, compared to the economy 1980. Since poor families transfer less, all other things equal, the fall in income at the lower end of the distribution implies that children from the bottom quartile of the income distribution receive less support, relative to the economy 1980.

As a consequence, children from poor families need to take out more loans to finance their studies. This is because in the model, college subsidies never cover the full amount of tuition fees,

⁴⁵Notice that I keep the distribution of parents at age 51, the age in which transfers are made, constant across all experiments with respect to the economy 1980. I also conducted experiments where the distribution of parents adjusts, such that it reproduces itself in equilibrium. In this case, the fact that parents internalize the increase in tuition fees and adjust their savings accordingly shifts the wealth distribution of parents at the age at which transfers are made. Parental transfers are higher in equilibrium, and the fraction of constrained households is thus significantly lower. The overall mechanisms at play are the same, however. Results from these experiments, which were also circulated in an earlier version of this paper, can be obtained upon request.

in line with evidence by Keane and Wolpin (2001) and Brown et al. (2012). The fact that tuition fees are higher in the economy 2000 increases the amount of borrowing, all other things equal.

Notice that the increase in the college premium raises the return of college attendance, *ceteris paribus*. Since the real interest rate declines between 1980 and 2000, borrowing becomes also cheaper. Despite this, children from poor families are discouraged from attending college, even if funding is readily available. This is because bigger leverage makes the decision to enrol in college more risky. Loans need to be paid back in full, even if the student fails to graduate, see Chatterjee and Ionescu (2012).

As for children of rich households, their parents are relatively richer compared to the economy 1980. This makes the parents more likely to provide transfers, all other things equal. However, the larger transfers do not keep pace with the rise in tuition. As a consequence, children from rich families also need to rely more on external funding to finance their studies in the economy 2000. The fraction of tuition that they need to finance through borrowing is smaller, though, than for those children from poor families. Hence, the joint decline in parental transfers and rising tuition fees modifies the risk-return trade-off between attending college for children at different parts of the family income distribution, such that children from rich families are more eager to attend college, and children with poor families are less so, compared to the economy 1980.

In sum, the joint development of parental transfers and (net) tuition across the family income distribution is able to generate (i) the increase in enrolment gaps between 1980 and 2000; (ii) an increase in the overall fraction of borrowing-constrained households; and (iii) a shift in the distribution of constrained children.

VII Conclusion

This paper analyzes the determinants of college enrolment in the U.S. and how these have changed over time. I propose a rich quantitative life cycle model with college enrolment. An important feature of the model is that altruistic parents provide financial support to their children. Using counterfactual experiments, I find that 24 percent of all households were financially constrained in their college decision at the beginning of the 1980s. The share of constrained households increased to 28 percent at the beginning of the 2000s.

The large fraction of borrowing-constrained households contrasts sharply with the findings in the previous literature. Using a reduced-form approach, Carneiro and Heckman (2002) concluded in a highly influential paper that at most 8 percent of households are financially constrained in their college decision. I replicate the procedure of Carneiro and Heckman (2002) using enrolment data generated by my model. I show that the reduced-form estimates obtained from the data and model coincide.

The considerable discrepancy between reduced-form estimates and counterfactual experiments suggests that reduced-form approaches do not correctly identify constrained households. Using the counterfactual enrolment data generated by my model, I show that the assumptions needed in order to identify constrained households through enrolment gaps are not fulfilled.

The interplay between parental transfers and college enrolment decisions of young households is a major factor underlying the difference. I conclude that a structural model of parental transfers

and college enrolment is necessary to identify constrained households adequately.

Two model assumptions turn out to be key for explaining the large share of borrowing-constrained households and its development over time, as well as the changing role of family income in determining college entry. First, parental altruism is only one-sided. And second, intergenerational borrowing is not permitted. These assumptions imply that parents face a trade-off between transferring resources to their children and saving for their own future consumption. As a consequence, college investment may be inefficiently low, even for children whose parental resources exceed the cost of college education.

This finding means that even controlling for parental resources, as it is typically done in reduced-form estimates, does not guarantee that constrained households are identified correctly. Moreover, the increase in earnings uncertainty observable between 1980 and 2000 increases household's need to self-insure. As a result, the amount of resources parents transfer to their children declines in the aggregate. The fraction of constrained households increases. The rise in within-group earnings inequality instead contributes to wider enrolment gaps between children from rich and poor families.

An important extension of this paper would be to model the determinants of academic ability more carefully. Mainly for simplicity, I assumed here that children's ability is determined partly 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 taken explicitly 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 extended model could be used to study the design of optimal education policies, a question I did not consider in this paper. Moreover, a structural model of the transmission of abilities could shed more light on the complex interaction between family income and children's ability in determining college entry.

References

- Abbott, B., Gallipoli, G., Meghir, C., and Violante, G. L. (2013), Education Policy and Intergenerational Transfers in Equilibrium, Mimeo, New York University.
- Ábrám, A. (2008), Earnings Inequality and Skill-biased Technological change with Endogenous Choice of Education, *Journal of the European Economic Association*, 6(2-3), 695–704.
- 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. (2013), A Dynamic Model of Altruistically-Motivated Transfers, *Forthcoming, Review of Economic Dynamics*.
- 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. (2012), A New Test of Borrowing Constraints for Education, *Review of Economic Studies*, 79(2), 511–538.
- Caballero, R. J., Farhi, E., and Gourinchas, P.-O. (2008), An Equilibrium Model of Global Imbalances and Low Interest Rates, *American Economic Review*, 98(1), 358–93.
- Caporale, T. and Grier, K. B. (2000), Political Regime Change and the Real Interest Rate, *Journal of Money, Credit and Banking*, 32(3), 320–34.
- 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 Parey, M. (2013), Maternal Education, Home Environments, And The Development Of Children And Adolescents, *Journal of the European Economic Association*, 11(s1), 123–160.
- 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. M. and Lochner, L. (2012), Early and Late Human Capital Investments, Borrowing Constraints, and the Family, Working Paper 18493, National Bureau of Economic Research.
- 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. (2012), Insuring student loans against the financial risk of failing to complete college, *Quantitative Economics*, 3(3), 393–420.
- 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.
- De Nardi, M. (2004), Wealth Inequality and Intergenerational Links, *Review of Economic Studies*, 71(3), 743–768.
- Dick, A. W. and Edlin, A. S. (1997), The implicit taxes from college financial aid, *Journal of Public Economics*, 65(3), 295–322.
- 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 S. Danziger and J. Waldfogel, eds., *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.
- Holter, H. (2013), Accounting for Cross-Country Differences in Intergenerational Earnings Persistence: The Impact of Taxation and Public Education Expenditure, Mimeo, Uppsala University.

- Hubbard, R. G., Skinner, J., and Zeldes, S. P. (1995), Precautionary Saving and Social Insurance, *Journal of Political Economy*, 103(2), 360–99.
- Imrohoroglu, A., Imrohoroglu, S., and Joines, D. H. (1995), A Life Cycle Analysis of Social Security, *Economic Theory*, 6(1), 83–114.
- Johnson, M. T. (2011), Borrowing Constraints, College Enrollment, and Delayed Entry, Working Papers 2011-006, Human Capital and Economic Opportunity Working Group.
- 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 O. Ashenfelter and D. Card, eds., *Handbook of Labor Economics, Handbook of Labor Economics*, volume 3, 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 E. A. Hanushek and F. Welch, eds., *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, R. and Stinebrickner, T. (2008), The Effect of Credit Constraints on the College Drop-Out Decision: A Direct Approach Using a New Panel Study, *American Economic Review*, 98(5), 2163–84.
- 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, The Levy Economics Institute.

VIII Appendix I: 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
v_0	max. share of tuition fees financed by subsidies	0.75
v_{asset}	impact of capital income on college subsidies	0.2
v_{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/wealth: 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 (2012) 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
$\bar{\kappa}$	average cost of college	2.16	5.2
r	interest rate	0.0357	0.0257
w	wage rate	1.21	1.28

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 economy 1980 and 2000. For detailed information regarding modifications and definitions, please refer to the main text.

Table 4: Effects of Family Income on College Attendance at Age 21 by AFQT Quartile, NLSY79.

	Ability Quartile:			
	1	2	3	4
Family Income Quartile 2	0.0273 (0.0408)	-0.0526 (0.0547)	0.0956 (0.0668)	0.0862 (0.0544)
Family Income Quartile 3	0.0258 (0.0490)	-0.0730 (0.0572)	0.1222 (0.0664)	0.0887 (0.0544)
Family Income Quartile 4	0.1171 (0.0600)	-0.0486 (0.0627)	0.1818 (0.0691)	0.1541 (0.0539)

Notes: Estimates obtained from Table 4 in Belley and Lochner (2007), Panel A. Regressions include dummy variables for mother's education and other controls. Standard errors are in parenthesis.

Table 5: Effects of Family Income on College Attendance at Age 21 by Ability Quartile, Model Economy 1980.

	Ability Quartile:			
	1	2	3	4
Family Income Quartile 2	0.0716	0.1777	0.1882	0.1391
Family Income Quartile 3	0.0491	0.1342	0.1427	0.1031
Family Income Quartile 4	0.0696	0.2160	0.2314	0.1606

Notes: Family income is measured at the time the college decision is taken (age=21). Ability refers to measurable ability, as defined in the main text. Family income effects are measured as the difference from the enrolment rate of the lowest income quartile. I control for parental education, as described in the main text.

Table 6: Effects of Family Income on College Attendance at Age 21 by Ability Quartile, Model Economy 1980, No Borrowing Limit on Student Loans

	Ability Quartile:			
	1	2	3	4
Family Income Quartile 2	-0.0995	0.0527	0.0677	-0.0026
Family Income Quartile 3	-0.2369	-0.0477	-0.0289	-0.1166
Family Income Quartile 4	-0.3696	-0.1687	-0.1477	-0.2441

Notes: Family income is measured at the time the college decision is taken (age=21). Ability refers to measurable ability, as defined in the main text. Family income effects are measured as the difference from the enrolment rate of the lowest income quartile. I control for parental education, as described in the main text.

Table 7: Fraction of households constrained in their college decision, per income quartile, Economy 1980

	% constrained
Family Income Quartile 1	47
Family Income Quartile 2	31
Family Income Quartile 3	19
Family Income Quartile 4	2

Remark: Table shows the fraction (expressed in percentage) of households that are financially constrained in their college decision for each family income quartile.

Table 8: Fraction of households constrained in their college decision, per income quartile, Economy 1980, Experiment $\zeta = 0.75$

	% constrained
Family Income Quartile 1	38
Family Income Quartile 2	27
Family Income Quartile 3	19
Family Income Quartile 4	2

Remark: Table shows the fraction (expressed in percentage) of households that are financially constrained in their college decision for each family income quartile.

Table 9: Effects of Family Income on College Attendance at Age 21 by AFQT Quartile, NLSY97.

	Ability Quartile:			
	1	2	3	4
Family Income Quartile 2	-0.0560 (0.0460)	0.0766 (0.0556)	0.0518 (0.0544)	0.0869 (0.0431)
Family Income Quartile 3	0.1336 (0.0559)	0.1545 (0.0590)	0.0957 (0.0564)	0.0520 (0.0424)
Family Income Quartile 4	0.2793 (0.0656)	0.1611 (0.0616)	0.1992 (0.0582)	0.0905 (0.0437)

Notes: Estimates obtained from Table 4 in Belley and Lochner (2007), Panel B. Regressions include dummy variables for mother's education and other controls. Standard errors are in parenthesis.

Table 10: Effects of Family Income on College Attendance at Age 21 by Ability Quartile, Model Economy 2000.

	Ability Quartile:			
	1	2	3	4
Family Income Quartile 2	0.0671	0.2568	0.2755	0.1879
Family Income Quartile 3	0.0288	0.2064	0.2241	0.1417
Family Income Quartile 4	0.1874	0.3952	0.4185	0.3141

Notes: Family income is measured at the time the college decision is taken (age=21). Ability refers to measurable ability, as defined in the main text. Family income effects are measured as the difference from the enrolment rate of the lowest income quartile. I control for parental education, as described in the main text.

Table 11: Effects of Family Income on College Attendance at Age 21 by Ability Quartile, Model Economy 2000, No Borrowing Limits on Student Loans

	Ability Quartile:			
	1	2	3	4
Family Income Quartile 2	0.4680	0.4624	0.4618	0.4644
Family Income Quartile 3	0.1215	0.3187	0.3383	0.2468
Family Income Quartile 4	0.1077	0.2877	0.3078	0.2174

Notes: Family income is measured at the time the college decision is taken (age=21). Ability refers to measurable ability, as defined in the main text. Family income effects are measured as the difference to the enrolment rate of the lowest income quartile. I control for parental education, as described in the main text.

Table 12: Fraction of households constrained in their college decision, per income quartile, Economy 2000.

	% constrained
Family Income Quartile 1	17
Family Income Quartile 2	45
Family Income Quartile 3	32
Family Income Quartile 4	16

Remark: Table shows the fraction (expressed in percentage) of households that are financially constrained in their college decision for each family income quartile.

IX Appendix II: 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 Gallipoli et al. (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 5,000 USD at the beginning of the 1980s (in 1980 USD) (see the Digest of Educational Statistics). GDP per capita was about 12,000 USD (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 2,300 USD. According to Chatterjee and Ionescu (2012), 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 do, 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, 2012). The same applies to the distribution of observable abilities in the model. A comparison of percentiles that are farther 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 U.S. 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 1,000 USD (in 1998 USD) 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 therefore conclude that the model well describes the sensitivity of college enrolment with respect to changes in tuition.

X Appendix III: 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 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 repeat this process until convergence is achieved.⁴⁶

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).⁴⁷ 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.⁴⁸

Notice that the program specified by (1) - (7) can be simplified by assuming that parents take the education decision on behalf of their children. This can be done because parents are altruistic, which implies that they anticipate the division of their transfers that is optimal from the children's point of view. Technically, this means that the problem (7) can be integrated into problem (3).

Different from "standard" dynamic programming problems, the program 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.⁴⁹ Concavity guarantees that the solution is a global maximum and thus is 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)

⁴⁶The algorithm converges at a geometric rate, see Laitner (2001).

⁴⁷An exception is the parental problem (3), for reasons that I explain in the next paragraph.

⁴⁸Thus its name "endogenous grid method".

⁴⁹Clausen 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.

value function at the threshold is less pronounced.⁵⁰ Second of all, skill accumulation in the 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 2. 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.

Small Constant Ψ . In a few instances, the code generates zero consumption when computing the optimal policy function. This is the case when the routine computes the optimal consumption and saving decisions for households with no wealth and with labor income that is approximately equal to the repayment of student loans. If the difference between the latter two values is very, very, close, the computer program produces a zero while in fact the difference is distinct from zero (machine zero).

Hence, household consumption is zero and marginal utility is infinity. The problem is no longer well-defined, and the software produces a missing value in these cases. Through the iterations, the missing values are distributed to the whole state space and the program crashes.

In order to avoid this, I add a very small constant Ψ to the income component in the budget constraint of the young households. This is a common device in the literature, see e.g. Heer and Maußner (2005), Ch. 7. I choose $\Psi = 10^{-11}$. Together with the maximum amount of loans I use to compute the counterfactual ($\bar{\chi} = 5$), this choice turns out to be sufficient to avoid numerical instabilities in all the experiments I conducted. I use the software package Gauss 9.0 (32-bit) on a Windows XP operating system.

Computation of the Equilibrium

Using the policy functions 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 v .
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.

⁵⁰A similar argument applies to the savings of a parent household in model period $j = 31$. Parental savings increase parental wealth holdings, which are part of the child household's state space. Because 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.

5. Compute the aggregate asset holdings K_1 and the new human capital stock L_1 using the 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 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.

I approximate the measure of agents by means of a probability density function.⁵¹ 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.⁵²

The optimal choice will almost surely be off-grid. In order to map the optimal choices onto the grid, I introduce a 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, as is the case for a positive fraction of the population. I thus allocate the grid points such that they 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.

⁵¹Heer 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.

⁵²The 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 the correspondence between the optimal policies and the resulting weights will be, leading to better aggregation results.

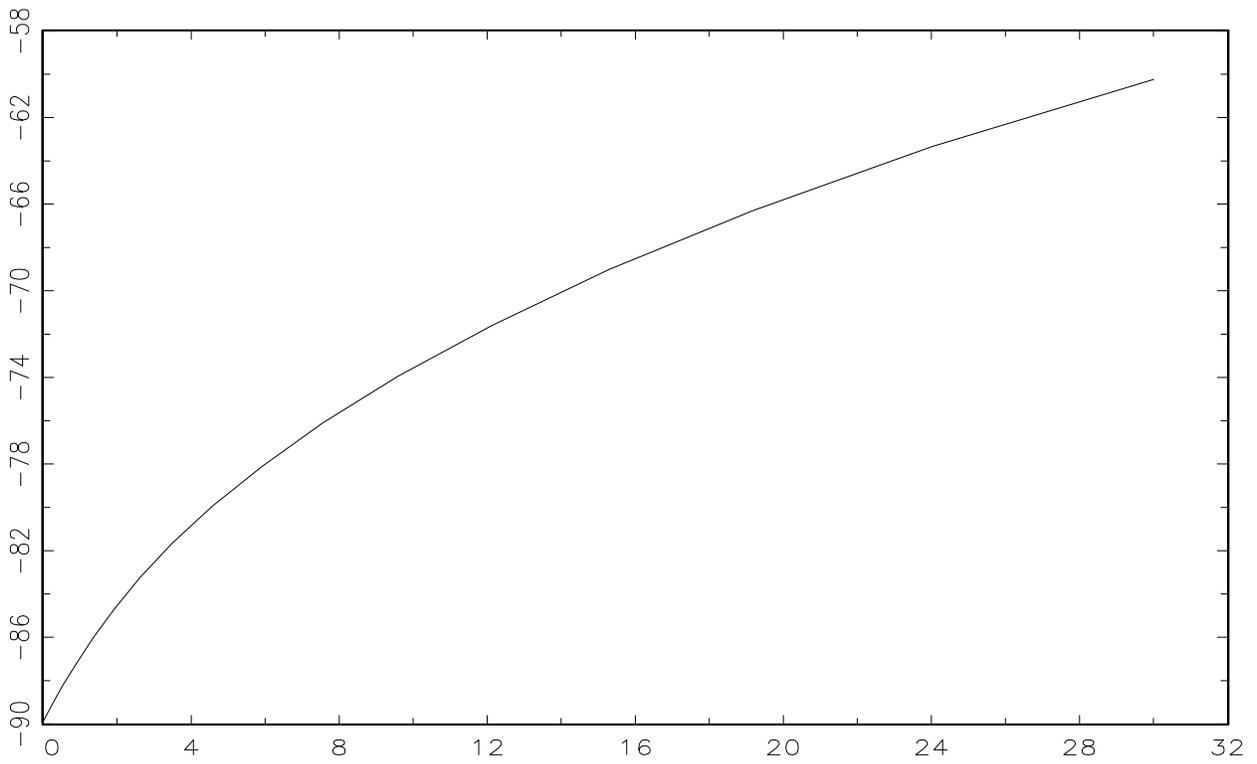


Figure 2: Value function for different levels of household wealth, after controlling for education, productivity and children's ability shock