
Optimal Need-Based Financial Aid

Mark Colas (University of Oregon)
Sebastian Findeisen (University of Konstanz)
Dominik Sachs (LMU Munich)

Discussion Paper No. 225

December 18, 2019

Optimal Need-Based Financial Aid*

Mark Colas
University of Oregon

Sebastian Findeisen
University of Konstanz

Dominik Sachs
LMU Munich

December 15, 2019

Abstract

We study the optimal design of student financial aid as a function of parental income. We derive optimal financial aid formulas in a general model. For a simple model version, we derive mild conditions on primitives under which poorer students receive more aid even without distributional concerns. We quantitatively extend this result to an empirical model of selection into college for the United States that comprises multidimensional heterogeneity, endogenous parental transfers, dropout, labor supply in college, and uncertain returns. Optimal financial aid is strongly declining in parental income even without distributional concerns. Equity and efficiency go hand in hand.

JEL-classification: H21, H23, I22, I24, I28

Keywords: Financial Aid, College Subsidies, Optimal Taxation, Inequality

*Contact: mcolas@uoregon.edu, sebastian.findeisen@uni-konstanz.de, dominik.sachs@econ.lmu.de. An earlier version of this paper was circulated under the title “Designing Efficient College and Tax Policies”. We thank James J. Heckman, the editor, as well as four anonymous referees. We also thank Rüdiger Bachmann, Felix Bierbrauer, Richard Blundell, Christian Bredemeier, Friedrich Breyer, David Card, Pedro Carneiro, Jonathan M.V. Davis, Juan Dolado, Alexander Gelber, Marcel Gerard, Emanuel Hansen, Jonathan Heathcote, Nathaniel Hendren, Bas Jacobs, Leo Kaas, Marek Kapicka, Kory Kroft, Paul Klein, Tim Lee, Lance Lochner, Normann Lorenz, Thorsten Louis, Alex Ludwig, Marti Mestieri, Nicola Pavoni, Michael Peters, Emmanuel Saez, Aleh Tsyvinski, Gianluca Violante, Matthew Weinzierl, Nicolas Werquin, and seminar participants at Berkeley, Bocconi, Bonn (MPI & MEE), CEMFI, Dortmund, EIEF, EUI, Frankfurt, IFS/UCL, Lausanne, Louvain (CORE), Minneapolis FED, Notre Dame, NYU, Queen’s, Salzburg, St. Gallen, Tinbergen, Toronto, Toulouse, Uppsala, Warwick, and Yale, as well as conference participants at CEPR Public Meeting 16, CEPR Conference on Financing Human Capital 18, CESifo Public Sector Economics, IIPF, NBER SI Macro Public Finance, NORMAC, SEEK, and the Taxation Theory Conference. We thank Mehmet Ayaz for careful proofreading. Dominik Sachs’ research was partly funded by a post-doc fellowship of the Fritz-Thyssen Foundation and the Cologne Graduate School in Management, Economics and Social Sciences. Financial support by Deutsche Forschungsgemeinschaft through CRC TRR 190 (project number 280092119) is gratefully acknowledged

1 Introduction

In all OECD countries, college students benefit from financial support (OECD, 2014). Moreover, with the goal of guaranteeing equality of opportunity, financial aid is typically need-based and targeted specifically to students with low parental income. In the United States, the largest need-based program is the Pell Grant. Federal spending on this program exceeded \$30 billion in 2015 and has grown by over 80% during the last 10 years (College Board, 2015). One justification for student financial aid in the policy debate is that the social returns to college exceed the private returns because the government receives a share of the financial returns through higher tax revenue (Carroll and Erkut, 2009; Baum et al., 2013). This lowers the effective fiscal costs (i.e., net of tax revenue increases) of student financial aid.¹

In this paper, we study the optimal design of financial aid and show that considering dynamic scoring aspects is crucial to assessing the desirability of need-based programs such as the Pell Grant. The reduction in the effective fiscal costs of student financial aid due to dynamic fiscal effects varies along the parental income distribution. We show that effective fiscal costs are increasing in parental income and are therefore lowest for those children that are targeted by the Pell Grant. The policy implication is that need-based financial aid is desirable not only because it promotes intergenerational mobility and equality of opportunity. Need-based financial aid is also desirable from an efficiency point of view because subsidizing the college education of children from weak parental backgrounds is cheaper for society than subsidizing students from "average" parental backgrounds. The usual equity-efficiency trade-off does not apply for need-based financial aid.

To arrive there, we start with a general model without imposing restrictions on the underlying heterogeneity in the population. Further, besides enrollment, labor supply and savings decisions, we consider dropout, labor supply during college and endogenous parental transfers. We derive a simple optimality condition for financial aid that transparently highlights the key trade-offs. At a given level of parental income, optimal financial aid decreases in the share of inframarginal students, which captures the marginal costs. These costs are scaled down by the marginal social welfare weights attached to these students. Optimal financial aid increases in the share of marginal students² and the fiscal externality per marginal student, which jointly capture the marginal benefits of the subsidy. The fiscal externality is the change in lifetime fiscal contributions causal to college attendance.³ For the optimality condition, the

¹The Congressional Budget Office (CBO), following a request by the Senate Committee on the Budget, recently documented the growth in the fiscal costs of Pell Grant spending (Alsalam, 2013). Dynamic scoring aspects are neglected in this report: the positive fiscal effects through higher tax revenue in the future are not taken into account. Generally, the CBO does consider issues of dynamic scoring: <https://www.cbo.gov/publication/50919>.

²Those students that are at the margin of attending college with respect to financial aid.

³On top of that, financial aid is also increasing in the completion elasticity with respect to financial aid and the fiscal externality due to completing college instead of dropping out. This channel, however, turns out to be quantitatively of minor importance.

specific reason why marginal students change their behavior due to a change in subsidies (e.g., borrowing constraints or preferences) is not important.

Elasticities linking changes in enrollment behavior to changes in financial aid have been estimated in the literature (e.g., by Dynarski (2003) and Castleman and Long (2016)). These papers provide guidance about the average value of this policy elasticity or about its value at a particular parental income level. However, knowledge about how this elasticity varies along the parental income distribution is missing. Knowledge of those parameters for students from different parental income groups, however, is necessary to analyze the welfare effects of need-based financial aid. Further, these elasticities are not deep parameters but do change as policy changes. The main approach of this paper is therefore a structural model of selection into college that allows us to compute this policy elasticity along the parental income distribution and for alternative policies.

As a first step, however, before studying this empirical model, we consider a simple theoretical setting. We reduce the complexity of the problem by focusing on two dimensions of heterogeneity: (i) parental transfers and (ii) returns to college. Further, we simplify the model by making the problem static, shutting down risk, labor supply during college and dropout. We first show that financial aid is decreasing in parental income even in the absence of distributional concerns if the distribution of returns is log concave (which implies a decreasing hazard rate)⁴ and if returns and parental income are independently distributed. We then show that these analytical results extend to the empirically more plausible case of a positive association between parental income and child ability.⁵

We then move to our structural life-cycle model, where we account for earnings risk, dropout, labor supply during college and, importantly, we account for crowd-out of parental transfers by explicitly modeling parental decisions to save, consume and provide transfers to their children. Another additional crucial ingredient of the model is heterogeneity in the psychic costs of education because monetary returns can only account for a small part of the observed college attendance patterns (Heckman et al., 2006). Using data from the National Longitudinal Survey of Youth 1979 and 1997, we estimate the parameters of our model via maximum likelihood and provide a detailed discussion of how variation in the data helps us to identify the crucial parameters.

We find that optimal financial aid policies are strongly progressive. In our preferred specification, the level of financial aid drops by 48% moving from the 25th percentile of the parental income distribution to the 75th percentile. The strong progressivity result does not rely on the Utilitarian welfare criterion. We show that a social planner that sets equal social welfare weights on all students or is only interested in maximizing tax revenues would choose

⁴The hazard rate pins down the ratio of marginal over inframarginal students which is also key in this simplified model.

⁵We obtain this clear analytical result if the ability distribution of high parental income children dominates the distribution of low parental income children in the hazard rate order. For a Pareto distribution, e.g., the property of hazard rate dominance always holds in case of first-order stochastic dominance.

an almost equally progressive financial aid schedule. Second, our estimates suggest that targeted increases in financial aid for students below the 59th percentile of the parental income distribution, are self-financing by increases in future tax revenue; this implies that targeted financial aid expansions could be free-lunch policies. Both results point out that financial aid policies for students are a rare case in which there is no equity-efficiency trade-off.

In a last step, we provide several extensions and robustness checks. We show that our progressivity result also holds if we (i) remove borrowing constraints, (ii) choose the merit-based dimension of financial aid optimally, (iii) allow the government to set an optimal Mirrleesian income tax schedule, (iv) model early educational investments and thereby endogenize ability and (v) if the relative wage for college educated labor is determined in general equilibrium.

Our paper contributes to the existing literature in several ways. Stantcheva (2017) characterizes optimal human capital policies in a very general dynamic model with continuous education choices. The main differences with our approach are twofold. First, theoretically, we study a model with discrete education choices as we find this a natural way to study financial aid policies. As we show, the optimality conditions are quite distinct from the continuous case and different elasticities are required to characterize the optimum. Second, the extensive margin education decision allows us to incorporate a large degree of heterogeneity without making the optimal policy problem intractable. This allows for a modeling approach that is close to the empirical, structural literature.

Bovenberg and Jacobs (2005) consider a static model with a continuous education choice and derive a “siamese twins” result: they find that the optimal marginal education subsidy should be as high as the optimal marginal income tax rate, thereby fully offsetting the distortions from the income tax on the human capital margin.⁶ Lawson (2017) uses an elasticity approach to characterize optimal uniform tuition subsidies for all college students.⁷ Jacobs and Thuemmel (2018) study the role of skill-biased technical change for optimal college subsidies and income taxation. We contribute to this line of research by developing a new framework to analyze how education policies should depend on parents’ resources and also trade off merit-based concerns. Our theoretical characterization of optimal financial aid (and tax policies) allows for a large amount of heterogeneity, and we tightly connect our theory directly to the data by estimating the relevant parameters ourselves. Finally, the paper is also related to many empirical papers, from which we take the evidence to gauge the performance of the estimated model. These papers are mentioned in Section 4.

⁶Bohacek and Kapicka (2008) derive a similar result as in a dynamic deterministic environment. Findeisen and Sachs (2016), focus on history-dependent policies and show how history-dependent labor wedges can be implemented with an income-contingent college loan system. Koeniger and Prat (2017) study optimal history-dependent human capital policies in a dynastic economy where education policies also depend on parental background. Stantcheva (2015) derives education and tax policies in a dynastic model with multi-dimensional heterogeneity, characterizing the relationship between education and bequest policies.

⁷Our work is also complementary to Abbott et al. (2018) and Krueger and Ludwig (2013, 2016), who study education policies computationally in very rich overlapping-generations models.

We progress as follows. In Section 2 we develop the general model and characterize the optimal policies in terms of reduced-form objects. In Section 3 we consider a simplified version of the model, which allows us to transparently study mild conditions on primitives under which financial aid is optimally decreasing in parental income. In Section 4 we specify our quantitative model as a special case of the general model presented in Section 2 and present our estimation approach. Section 5 presents optimal financial aid policies, and Section 6 decomposes the forces which lead to an optimal financial aid schedule. In Section 7 we discuss further robustness issues. Section 8 concludes.

2 Optimal Financial Aid Policies

In this section we characterize optimal (need-based) financial aid policies for college students. Our approach is to work with a general model and characterize the optimal financial aid in terms of reduced-form objects. This formula is general on the one hand and economically intuitive on the other hand. It clearly highlights the role of the fiscal externality as a reason for why education is subsidized (Bovenberg and Jacobs 2005). The fiscal externality arises through the tax-transfer system: if college increases human capital and therefore earnings, college education implies a fiscal externality since the individual will pay more taxes. Hence, if the government imposed lump sum taxes that were independent of earnings, there would be no fiscal externality. In Section 4, we explore the quantitative implications of this optimality condition in a fully specified structural empirical model, which is a special case of the model analyzed in Section 2. As an intermediate step, we theoretically explore a simplified framework in Section 3, for which we can derive conditions on primitives that imply that optimal financial aid is indeed need-based, i.e., that financial aid is decreasing in parental income, even in the absence of distributional concerns.

2.1 Individual Problem

Individuals start life in year $t = 0$ as high school graduates and are characterized by a vector of characteristics $X \in \chi$ and (permanent) parental income $I \in \mathbb{R}_+$. Life lasts T periods and individuals face the following decisions. At the beginning of the model, they face a binary choice: enrolling in college or not. If individuals decide against enrollment, they directly enter the labor market and make labor-leisure decisions every period. If individuals decide to enroll in college, they also make a labor-leisure decision during college and, at the beginning of the year, decide to drop out or continue. After graduating or dropping out, individuals enter the labor market.

We start by considering labor market decisions of individuals that either are out of college or have chosen to forgo college altogether. This is a standard labor-leisure-savings problem

with incomplete markets. Let $V_t^W(\cdot)$ denote the value function of an individual in the labor market in year t . Then the recursive problem is given by

$$V_t^W(X, I, e, a_t, w_t) = \max_{c_t, \ell_t} U(c_t, \ell_t) + \beta \mathbb{E} [V_{t+1}^W(X, I, e, a_{t+1}, w_{t+1}) | w_t]$$

subject to the budget constraint

$$c_t + a_{t+1} = \ell_t w_t - \mathcal{T}(\ell_t w_t) + a_t(1 + r) + tr_t(X, I, e, w_t).$$

The state variables are the initial characteristics (X, I) , the education level $e \in \{H, D, G\}$ (high school graduate, college dropout, college graduate), assets a_t , and the current wage w_t . The variables (X, I, e) are state variables because they may affect parental transfers $tr_t(X, I, e, w_t)$ and because they may affect the evolution of future wages. The dependence on the education decision then captures the returns to education. The function \mathcal{T} captures the tax-transfer system. Finally, we assume that the utility function is such that there are no income effects on labor supply. Given those value functions, we now turn to the value functions of the different education decisions. The value of not enrolling in college (i.e., choosing education level H) is simply given by

$$V^H(X, I) = \mathbb{E} [V_1^W(X, I, e = H, a_1 = 0, w_1)].$$

Regarding the realization of uncertainty, the timing is such that individuals directly enter the labor market in period one and draw their first wage w_1 , which is hence only known after the education decision has been made. Next, we turn to the decisions during college. Besides the question of how much to work and consume while in college, individuals also make the binary decision of dropping out or staying enrolled.

The value function of a college student at age t is given by

$$V_t^E(X, I, a_t, \varepsilon_t) = \max[V_t^{ND}(X, I, a_t, \varepsilon_t), V_t^D(X, I, a_t, \varepsilon_t)]$$

where $V^D(\cdot)$ is the value function associated with dropping out, $V_t^{ND}(\cdot)$ denotes the value function of staying enrolled (not dropping out), and ε_t is a vector of preference shocks. Agents who drop out of college enter the labor force and may also pay a psychic cost associated with dropping out. The value of dropping out is therefore given by:

$$V_t^D(X, I, a_t) = \mathbb{E} [V_t^W(X, I, e = D, a_t, w_t)] - d(\varepsilon_t)$$

where $d(\varepsilon_t)$ represents the psychic cost of dropping out.

The value function for staying enrolled is a bit more complex and given by:

$$V_t^{ND}(X, I, a_t, \varepsilon_t) = \max_{c_t, \ell_t^E} [U^E(c_t, \ell_t^E; X, \varepsilon_t) + \beta \{ (1 - Pr_t^{Grad}(X)) \times \mathbb{E} [V_{t+1}^E(X, I, a_{t+1}, \varepsilon_{t+1})] + Pr_t^{Grad}(X) \times \mathbb{E} [V_{t+1}^W(X, I, e = G, a_{t+1}, w_{t+1})] \}]$$

subject to

$$c_t = \ell_t^E \underline{\omega} + a_t (1 + r(a_t, I)) - a_{t+1} - \mathcal{F}(X) + \mathcal{G}(X, I) + tr_t^E(X, I, \mathcal{G}(X, I))$$

and

$$a_{t+1} \geq \bar{a}_{t+1}^E.$$

$\underline{\omega}$ is the wage that students earn if they work during college and $\mathcal{F}(X)$ is tuition. Tuition might vary by X because of regional differences in college tuition, for example. We denote work in college by ℓ_t^E .⁸ The term $\mathcal{G}(X, I)$ is the amount of financial aid a student with characteristics X and parental income I receives, and $tr_t^E(X, I, \mathcal{G}(X, I))$ captures parental transfers in year t for children that are enrolled in college. They are endogenous with respect to the level of financial aid to account for the potential crowding out of parental transfers through financial aid. $Pr_t^{Grad}(X)$ is a stochastic graduation probability which can depend on the vector X . We allow the interest rate for college enrollees to vary by the agent's asset position (positive or negative) and by the agent's parental income. We denote flow utility while enrolled in college by $U^E(c_t, \ell_t^E; X, \varepsilon_t)$. Importantly, this flow utility may include the psychic costs and nonpecuniary benefits of college attendance, in addition to flow utility from consumption and labor supply. These psychic costs have been found to be important in explaining college enrollment patterns.⁹ The flow utility in college can depend directly on personal characteristics, X , allowing these psychic costs of college to vary with the individual's characteristics. Note that the vector of personal characteristics, X , may also include idiosyncratic preferences for enrolling in college.

Finally we denote the value of enrolling into college in the first place as

$$V^E(X, I) = \mathbb{E} [V_1^E(X, I, a_1 = 0, \varepsilon_1)] + v(X),$$

where $v(X)$ is a function that gives any additional nonpecuniary benefits of enrolling in college for agents with characteristics X . An individual enrolls in college if $V^E(X, I) \geq V^H(X, I)$.

Denote by $P_t^D(X, I, \mathcal{G}(X, I))$ the share of individuals of type (X, I) that drop out in period t . Importantly the model captures the idea that the dropout decision is endogenous with

⁸We assume these earnings are not taxed. In the data, the average earnings of students who work in college are so low that they do not have to pay positive income taxes; in addition, the vast majority of college students does not qualify for welfare/transfer programs.

⁹See Cunha et al. (2005), Heckman et al. (2006) or Heckman and Navarro (2007).

respect to financial aid. Further, denote by $P_t^E(X, I, \mathcal{G}(X, I)) = \prod_{s=1}^t (1 - P_s^D(X, I, \mathcal{G}(X, I))) \times \prod_{s=1}^{t-1} (1 - Pr_s^{Grad}(X))$ the proportion of all initially enrolled students that are enrolled in period t . Finally, we denote the proportion of initially enrolled students that successfully complete college by $P^C(X, I, \mathcal{G}(X, I)) = \sum_{t=1}^{t_g^{max}} P_t^E(X, I, \mathcal{G}(X, I)) Pr_t^{Grad}(X)$. We move to the policy analysis and for the remainder of the section make three simplifying assumption for the purpose of simpler notation. We assume that individuals can only drop out after two years in college such that $P_t^D(X, I, \mathcal{G}(X, I)) = 0$ if $t \neq 3$ and cannot graduate before year $t = 3$, i.e. $Pr_t^{Grad}(X) = 0$ for $t = 1, 2$.¹⁰ Finally, we assume that financial aid only depends only on parental income, and not on other characteristics, X . We therefore write financial aid as $\mathcal{G}(I)$ for the remainder of this section.¹¹

2.2 Fiscal Contributions

We now define the expected net fiscal contributions for different types (X, I) and different education levels as these will be key ingredients for the policy analysis. We start with the net present value (NPV) in net tax revenues of high school graduates of type (X, I) :

$$\mathcal{NT}_{NPV}^H(X, I) = \sum_{t=1}^T \left(\frac{1}{1+r} \right)^{t-1} \mathbb{E}(\mathcal{T}(y_t) | X, I, H),$$

where $y_t = w_t \ell_t$ is total earnings in year t .

The fiscal contribution of a dropout is given by their net present value of tax payments minus grants received:

$$\mathcal{NT}_{NPV}^D(X, I) = \sum_{t=3}^T \left(\frac{1}{1+r} \right)^{t-1} \mathbb{E}(\mathcal{T}(y_t) | X, I, D) - \mathcal{G}(I) \sum_{t=1}^2 \left(\frac{1}{1+r} \right)^{t-1}.$$

Finally, we turn to students that do not dropout but graduate. The average fiscal contribution of graduates of type (X, I) is given by:

$$\mathcal{NT}_{NPV}^G(X, I) = \frac{1}{\sum_{g=3}^{t_g^{max}} P_g^E(X, I, \mathcal{G}(I)) Pr_g^{Grad}(X)} \sum_{g=3}^{t_g^{max}} P_g^E(X, I, \mathcal{G}(I)) Pr_g^{Grad}(X) \left[\sum_{t=g+1}^T \left(\frac{1}{1+r} \right)^{t-1} \mathbb{E}(\mathcal{T}(y_t) | X, I, G) - \mathcal{G}(I) \sum_{t=1}^g \left(\frac{1}{1+r} \right)^{t-1} \right]$$

¹⁰We provide the optimal policy formulas without these simplifying assumptions in Appendix A.2. The intuition of these formulas are the same but the notation is considerably more cumbersome.

¹¹We allow for other characteristics to enter the financial aid formula in the quantitative version of the model in Section 4. We show that our main result also extends to the case in which the merit-based elements are chosen optimally in Appendix C.10.

where t_g^{max} is the latest possible graduation date. Finally, we define the expected fiscal contribution of an individual that decides to enroll:

$$\mathcal{NT}_{NPV}^E(X, I) = P^C(X, I, \mathcal{G}(I)) \times \mathcal{NT}_{NPV}^G(X, I) + (1 - P^C(X, I, \mathcal{G}(I))) \times \mathcal{NT}_{NPV}^D(X, I).$$

Before we derive optimal education subsidies, we ease the upcoming notation a little bit. Let a type (X, I) be labeled by j and define the enrollment share for income level I :

$$E(I) = \int_{\mathcal{X}} \mathbb{1}_{V_j^E \geq V_j^H} h(X|I) dX,$$

where $\mathbb{1}_{V_j^E \geq V_j^H}$ is an indicator function capturing the education choice for each type $j = (X, I)$. Next, we define the completion rate by

$$C(I) = \frac{\int_{\mathcal{X}} \mathbb{1}_{V_j^E \geq V_j^H} P^C(X, I, \mathcal{G}(I)) h(X|I) dX}{E(I)},$$

which captures the share of enrolled students of parental income level I that actually graduate. We assume that these shares, as well as the probabilities of dropping out, $P_t^D(X, I, \mathcal{G}(I))$, are differentiable in the level of financial aid.

2.3 Government Problem and Optimal Policies

We now characterize the optimal financial aid schedule $\mathcal{G}(I)$. We denote by $F(I)$ the unconditional parental income CDF, by $K(X, I)$ the joint CDF and by $H(X|I)$ the conditional one; the densities are $f(I)$, $k(X, I)$, and $h(X|I)$, respectively. The government assigns Pareto weights $\tilde{k}(X, I) = \tilde{f}(I)\tilde{h}(X|I)$, which are normalized to integrate up to one.

Importantly, we assume that the government takes the tax-transfer system $\mathcal{T}(\cdot)$ as given and consider the optimal budget-neutral reform of $G(I)$. Whereas the tax-transfer system is not changed if financial aid is reformed, a change in the financial aid schedule changes the size and the composition of the set of individuals that go to college. This implies a change in tax revenue and transfer spending that directly feeds back into the available resource for financial aid.¹² Taking the tax-transfer system as given, the problem of the government is

$$\max_{\mathcal{G}(I)} \int_{\mathbb{R}_+} \int_{\mathcal{X}} \max\{V^E(X, I), V^H(X, I)\} \tilde{k}(X, I) dX dI \quad (1)$$

subject to the net present value government budget constraint:

¹²We consider this as the more policy-relevant exercise than considering the joint optimal choice of $\mathcal{T}(\cdot)$ and $\mathcal{G}(I)$. Nevertheless, to complete the picture, in Appendix C.9, we consider the joint optimal design of financial aid $\mathcal{G}(I)$ and the tax-transfer system $\mathcal{T}(\cdot)$. Further, we also explore jointly optimal merit and need-based financial aid in Appendix C.10.

$$\begin{aligned}
& \int_{\mathbb{R}_+} \int_{\mathcal{X}} \mathcal{N}\mathcal{T}_{NPV}^H(X, I) \mathbb{1}_{V_j^E < V_j^H} k(X, I) dX dI \\
& + \int_{\mathbb{R}_+} \int_{\mathcal{X}} \mathcal{N}\mathcal{T}_{NPV}^G(X, I) \mathbb{1}_{V_j^E \geq V_j^H} P^C(X, I, \mathcal{G}(I)) k(X, I) dX dI \\
& + \int_{\mathbb{R}_+} \int_{\mathcal{X}} \mathcal{N}\mathcal{T}_{NPV}^D(X, I) \mathbb{1}_{V_j^E \geq V_j^H} (1 - P^C(X, I, \mathcal{G}(I))) k(X, I) dX dI \geq \bar{F}, \tag{2}
\end{aligned}$$

The term \bar{F} captures exogenous revenue requirements (e.g. spending on public goods) and exogenous revenue sources (e.g. tax revenue from older cohorts). Hence, $\bar{F} < 0$ could capture that the cohort for which we are reforming the financial aid schedule is effectively subsidized from other cohorts. Now we consider a marginal increase in $\mathcal{G}(I)$. As we show in Appendix A.1, it has the following impact on welfare:

$$\underbrace{\frac{\partial E(I)}{\partial \mathcal{G}(I)} \times \Delta \mathcal{T}^E(I)}_{\text{Enrollment Effect}} + \underbrace{\frac{\partial C(I)}{\partial \mathcal{G}(I)} \Big|_{E(I)} \times E(I) \times \Delta \mathcal{T}^C(I)}_{\text{Completion Effect}} - \underbrace{\tilde{E}(I) (1 - W^E(I))}_{\text{Mechanical Effect}} = 0. \tag{3}$$

The first two terms of (3) capture behavioral effects (i.e., changes in welfare that are due to individuals changing their behavior). The third term captures the mechanical welfare effect (i.e. the welfare effect that would occur for fixed behavior). We start with the latter.

The mechanical effect captures the direct welfare impact of the grant increase to inframarginal students. The more students are inframarginal in their decision to go to college and the more of them do not drop out, the higher are the immediate costs of the grant increase. The term $\tilde{E}(I)$ is the total discounted years of college attendance of income group I and is defined as

$$\tilde{E}(I) = \int_{\mathcal{X}} \mathbb{1}_{V_j^E \geq V_j^H} \left(\sum_{t=1}^{t_g^{max}} \left(\frac{1}{1+r} \right)^{t-1} P_t^E(X, I, \mathcal{G}(I)) \right) h(X|I) dX.$$

This captures the direct marginal fiscal costs of the grant increase. Since the utility of these students is valued by the government, the costs have to be scaled down by a social marginal welfare weight (Saez and Stantcheva, 2016). We denote average social marginal welfare weight of inframarginal students with parental income I by $W^E(I)$. Formally it is given by

$$W^E(I) = \frac{\int_{\mathcal{X}} \mathbb{1}_{V_j^E \geq V_j^H} \mathbb{E} \left[\sum_{t=1}^{t_g^{max}} \beta^{t-1} U_c^E(\cdot) \left(1 + \frac{\partial tr_t^E(\cdot)}{\partial \mathcal{G}(I)} \right) \prod_{s=1}^t (\mathbb{1}_{V_s^{ND} \geq V_s^D}) \prod_{s=1}^{t-1} (1 - Pr_s^{Grad}(X)) \right] \tilde{h}(X|I) dX}{\rho \frac{f(I)}{\bar{f}(I)} \tilde{E}(I)},$$

where ρ is the marginal value of public funds, U_c^E is the marginal utility of consumption, and $\mathbf{1}_{V_s^{ND} \geq V_s^D}$ is an indicator for an individual choosing not to drop out of college in year s . Thus, $W^E(I)$ is a money-metric (appropriately weighted) average marginal social welfare weight. One difference from the standard concept applies here, however. One has to correct for the implied reduction in parental transfers that accompanies an increase in resources for college students. For each marginal dollar of additional grants, students only have a change in consumption that is given by $\left(1 + \frac{\partial tr_t^E(X, I, \mathcal{G}(I))}{\partial \mathcal{G}(I)}\right)$. Ceteris paribus, the stronger the crowding out of transfers, the lower are these welfare weights since fewer of the additional grants effectively reach students.¹³

We now turn to the behavioral welfare effects in the first line of (3). The first term captures the change in tax revenues due to an increase in enrollment and $\frac{\partial E(I)}{\partial \mathcal{G}(I)}$ captures the additional enrollees. Since these individuals are marginal in their enrollment decision, this change in their decision has no first-order effect on their utility. Therefore, we only have to track the effect on welfare through the effect on public funds. The term $\Delta \mathcal{T}^E(I)$ captures the the average increase in the NPV of net tax revenues for these marginal enrollees. Formally, it is given by

$$\Delta \mathcal{T}^E(I) = \frac{\int_{\mathcal{X}} \mathbf{1}_{H_j \rightarrow E_j} \Delta \mathcal{T}^E(X, I) h(X|I) dX}{\int_{\mathcal{X}} \mathbf{1}_{H_j \rightarrow E_j} h(X|I) dX}, \quad (4)$$

where $\mathbf{1}_{H_j \rightarrow E_j}$ takes the value one if an individual of type j is marginal in her college enrollment decision with respect to a small increase in financial aid. By definition we have $\int_{\mathcal{X}} \mathbf{1}_{H_j \rightarrow E_j} h(X|I) dX = \frac{\partial E(I)}{\partial \mathcal{G}(I)}$. $\Delta \mathcal{T}^E(X, I)$ is the (expected) fiscal externality of an individual of type (X, I) : $\Delta \mathcal{T}^E(X, I) = \mathcal{N} \mathcal{T}_{NPV}^E(X, I) - \mathcal{N} \mathcal{T}_{NPV}^H(X, I)$.

There is a second behavioral effect due to endogenous college dropout. This second term in (3) captures the increase in tax revenue due to an increase in the completion rate of the inframarginal enrollees. The term $\left. \frac{\partial C(I)}{\partial \mathcal{G}(I)} \right|_{E(I)}$ is the partial derivative of completion w.r.t. financial aid, holding $E(I)$ constant. Therefore, the term $\left. \frac{\partial C(I)}{\partial \mathcal{G}(I)} \right|_{E(I)} \times E(I)$ captures the amount of inframarginal enrollees who did not graduate in the absence of the grant increase but graduate now. Again, the envelope theorem applies and the change in their behavior has no first-order effect on their utility. However, there is a welfare effect through the change in public funds. $\Delta \mathcal{T}^C(I)$ captures the implied change in net fiscal contributions through the increased completion rate:

$$\Delta \mathcal{T}^C(I) = \frac{\int_{\mathcal{X}} \Delta \mathcal{T}^C(X, I) \frac{\partial P^C(X, I, \mathcal{G}(I))}{\partial \mathcal{G}(I)} h(X|I) dX}{\int_{\mathcal{X}} \frac{\partial P^C(X, I, \mathcal{G}(I))}{\partial \mathcal{G}(I)} h(X|I) dX},$$

¹³Note that we are not accounting for parents' utilities here. Doing so would basically imply an increase in the social welfare weights as not only the children but also the altruistic parents are benefiting from the grants. The change in parental transfers would have no impact on parent's utility due to the envelope theorem.

where $\Delta\mathcal{T}^C(X, I) = \mathcal{N}\mathcal{T}_{NPV}^G(X, I) - \mathcal{N}\mathcal{T}_{NPV}^D(X, I)$. Finally, note that formula (3) is independent of the adjustment in labor supply during college as a response to the grant increase. This is an implication of the envelope theorem.

Formula (3) expresses the optimal policy as a function of reduced-form elasticities and provides intuition for the main trade-offs underlying the design of financial aid.¹⁴ It is valid without taking a stand on the functioning of credit markets for students, the riskiness of education decisions, or the exact modeling of how parental transfers are influenced by parental income and how they respond to changes in financial aid. Those factors, of course, influence the values of the reduced-form elasticities. For example, a tightening of borrowing constraints should increase the sensitivity of enrollment especially for low-income students.

However, note that all terms in the optimal financial aid formula are endogenous with respect to policies. Even if we know the empirical values for current policies, this is not enough to calculate optimal policies. For this purpose, a fully specified model is necessary. In the next Section 3, we consider a simplified model, for which we can derive closed-form solutions.¹⁵

3 Is Optimal Financial Aid Progressive? A Simple Model

Simplified Environment. We assume that preferences are linear in consumption and that labor incomes are taxed linearly at rate τ , which is larger than 0 and smaller than one. We consider a static problem. If individuals do not go to college, they earn income y_H . If they go to college, they pay tuition \mathcal{F} and earn $y_H(1 + \theta)$. Individuals are heterogeneous in ability/returns to college, θ , and each $\theta > 0$. There is no uncertainty. Further, individuals are heterogeneous in parental income I . If individuals go to college, they receive a parental transfer $tr(I)$ with $tr'(I) > 0$ and financial aid $\mathcal{G}(I)$.

Individual Problem. If an individual decides against college, utility is given by $U^H = (1 - \tau)y_H$. If an individual goes to college, utility is given by $U^C(\theta, I) = (1 - \tau)y_H(1 + \theta) - (\mathcal{F} - \mathcal{G}(I) - tr(I))$. For each income level I , we can define the ability of the marginal college graduate $\tilde{\theta}(I)$, implicitly given by $U^H = U^C(\tilde{\theta}(I), I)$. All types (θ, I) with $\theta \geq (<)\tilde{\theta}(I)$ (do not) attend college. Note that higher parental income here simply has the role of lowering the costs of college. This implies that high-parental-income children are more likely to select into college. This channel is reinforced if there is a positive association between I and θ .

Government Problem and Optimal Financial Aid for a Given I . The government uses non-negative Pareto weights over the types as in the general model from the last section.

¹⁴Sometimes such formulas are labeled as sufficient statistics formulas. See Kleven (2018) for a discussion on the terminology in the literature.

¹⁵We are very grateful to one of our referees for many detailed suggestions how to clarify the intuition behind the results in Section 3.

Consistent with the notation from last section, $F(I)$ is the parental income distribution and $H(\theta|I)$ the conditional distribution of ability. Appendix A.3 shows that the following version of equation (3) holds:

$$\underbrace{\frac{h(\tilde{\theta}(I)|I)}{y_H(1-\tau)}}_{\frac{\partial E(I)}{\partial \mathcal{G}(I)}} \times \underbrace{\left(\tau y_H \tilde{\theta}(I) - \mathcal{G}(I)\right)}_{\Delta \mathcal{T}^E(I)} - \underbrace{\left(1 - H(\tilde{\theta}(I)|I)\right)}_{\tilde{E}(I)} (1 - W^E(I)) = 0.$$

First note that there is no completion effect since we abstract from dropout. Second, the fiscal externality takes a simple form. Third the ratio of marginal over inframarginal students is determined by the hazard rate of the conditional skill distribution. Rewriting leads to a rather tractable expression for optimal financial aid $\mathcal{G}(I)$.

Proposition 1. *The optimal financial aid schedule in the simplified environment is given by*

$$\mathcal{G}(I) = \tau (\mathcal{F} - tr(I)) - y_H(1-\tau)^2 \frac{\left(1 - H(\tilde{\theta}(I)|I)\right)}{h(\tilde{\theta}(I)|I)} \times (1 - W^E(I)), \quad (5)$$

where $\tilde{\theta}(I) = \frac{\mathcal{F} - tr(I) - \mathcal{G}(I)}{(1-\tau)y_H}$. and $\tau (\mathcal{F} - tr(I)) = \tau y_H \tilde{\theta}(I)$.

Proof. See Appendix A.3. □

The first term in (5), $\tau (\mathcal{F} - tr(I))$, can be interpreted as a Pigouvian correction. Without any distortions, i.e. $\mathcal{G}(I) = \tau = 0$, the marginal college enrollee would be characterized by

$$\theta^*(I)y_H = \mathcal{F} - tr(I). \quad (6)$$

Here the private returns and costs are equalized to the social ones. Such a condition is typically called “first best”. When τ or $\mathcal{G}(I) \neq 0$, the marginal enrollee still equates private returns to private costs, but there is a wedge between the *social* returns and costs now. Equating private returns and costs yields:

$$\tilde{\theta}(I)(1-\tau)y_H = \mathcal{F} - tr(I) - \mathcal{G}(I). \quad (7)$$

Comparing (7) with (6) shows that the fiscal externality $\Delta \mathcal{T}^E(I) = \tau y_H \tilde{\theta}(I) - \mathcal{G}(I)$ can be seen as a wedge. This is the classical “siamese twins” result of Bovenberg and Jacobs (2005): the sole presence of taxes gives a rationale for subsidizing education and the size of the subsidy is increasing in the size of the tax. Setting $\mathcal{G}(I) = \tau (\mathcal{F} - tr(I)) = \tau y_H \theta^*(I)$ would imply $\tilde{\theta}(I) = \theta^*(I)$ and hence yield the first-best education level. When choosing the optimal education subsidy $\mathcal{G}(I)$, the social planner, however, cannot target the marginal students but has to account for the fact that an increase in $\mathcal{G}(I)$ also has to be paid to those students that are inframarginal in their decision.¹⁶ This is accounted for in the second part of (5).

¹⁶If the planner can choose $\mathcal{G}(I, \theta)$ in this simple model, she effectively has lump-sum taxes/transfers available (for all college students). She only needs to correct the fiscal externality in this case (the other

Since the decision of inframarginal students is not altered, this is a pure transfer which is valued by $W^E(I) - 1$ multiplied with the share of inframarginal students. This implies that if $W^E(I) > (<)1$, the planner would subsidize students of parental income up to a point where education is above (below) the first best level as defined above.¹⁷ Further, this second term inversely proportional to share of marginal students. Intuitively, the more marginal students can be incentivized, the higher is the relative weight on the first term.

In the following we want to explore whether financial aid optimally decreases with parental income. For this purpose, we shut down any redistributive case for financial aid and assume that $\frac{\partial W^E(I)}{\partial I} = 0$. Two useful benchmark cases generate this: (i) a government that solely wants to maximize tax revenue (implying $W^E(I) = 0$ for all I) and (ii) unweighted Utilitarianism (implying $W^E(I) = \text{constant} < 1$ for all I as we elaborate in Appendix A.3). If redistribution within college students is desired, i.e. with declining weights $W^{E'}(I) < 0$, this would strengthen the case for progressivity and need-based financial aid.

Is Optimal Financial Aid Decreasing in Parental Income? We proceed in two steps and first state a result on the progressivity if parental income and child's ability are independently distributed.

Corollary 1. *Assume that ability θ and parental income I are independent, that is, $H(\theta|I) = H(\theta) \forall \theta, I$. Further assume $\frac{\partial W^E(I)}{\partial I} = 0$, i.e. there is no desire to redistribute from high to low parental income students. Then the optimal financial aid schedule is progressive (i.e., $\mathcal{G}'(I) < 0 \forall I$) if the distribution $H(\theta)$ is log concave.*

Proof. See Appendix A.4. □

The first term in (5) is decreasing in I . The higher parental income, the lower are the costs of college $\mathcal{F} - tr(I)$ and hence, for a given rate of subsidization τ , the lower is the overall level of the subsidy. Since $\tilde{\theta}'(I) < 0$,¹⁸ the second term is decreasing in I if the inverse of the hazard rate of $H(\theta)$ is decreasing. As Bagnoli and Bergstrom (2005) point out, log-concavity of a density function is sufficient for an increasing hazard rate.¹⁹ Hence, in the illustrative case in which parental income and child's ability are independent, we have an important benchmark, where the selection mechanism through parental income in itself calls for progressive financial

considerations like the ratio of marginal to inframarginals and redistribution within students can be perfectly dealt with by choosing $\mathcal{G}(I, \theta)$ for each type. This is not the case in the more general model presented in Section 2. We analyze the case of jointly optimizing merit-based and need-based financial aid quantitatively in C.10.

¹⁷This resembles the results of the optimal income tax literature with extensive margin labor supply responses that negative participation taxes are optimal if the social welfare weight of low income workers is above one, see e.g. Saez (2002).

¹⁸Note that for this we need $tr'(I) + \mathcal{G}'(I) > 0$, i.e. that financial aid is not too progressive. As our proof in Appendix A.4 shows, this is the case.

¹⁹Log-concavity of a probability distribution is a frequent condition used in many mechanism design or contract theory applications, as this is "just enough special structure to yield a workable theory" (Bagnoli and Bergstrom, 2005).

aid. Next we turn to the empirically more appealing case in which parental income and ability are positively associated.²⁰

Corollary 2. *Assume that ability θ and parental income I are positively associated in the sense that for $I' > I$, the distribution $H(\theta|I')$ dominates $H(\theta|I)$ in the hazard rate order, that is,*

$$\forall \theta, I, I' \text{ with } I' > I : \frac{h(\theta|I)}{1 - H(\theta|I)} \geq \frac{h(\theta|I')}{1 - H(\theta|I')}. \quad (8)$$

Further assume $\frac{\partial W^E(I)}{\partial I} = 0$, i.e. there is no desire to redistribute from high to low parental income students. Then the optimal financial aid schedule is progressive (i.e. $\mathcal{G}'(I) < 0 \forall I$) if the conditional skill distributions $H(\theta|I)$ are log concave.

Proof. See Appendix A.5. □

This condition (8) is stronger than first-order stochastic dominance (FOSD) but does imply that the skill distribution of higher parental income levels first-order stochastically dominates the skill distribution of lower parental income levels. FOSD of the skill distribution, however, does not automatically imply (8).²¹ For the empirically plausible Pareto distribution, FOSD does imply dominance in the hazard rate order. Consider, for example, the specification $h(\theta|I) = \alpha(I) \frac{\theta^{\alpha(I)}}{\theta^{\alpha(I)+1}}$, where $\alpha(I)$ is the thickness parameter. Here we have $\frac{1-H(\theta|I)}{h(\theta|I)} = \frac{\theta}{\alpha(I)}$ and hence if $\alpha'(I) < 0$, then the tail of the skill distribution of high-parental-income children is thicker and the FOSD property is fulfilled. Therefore, (8) is fulfilled.

The goal of this section was to show that under some rather weak assumptions, optimal financial aid is indeed decreasing in income. Whereas the simple model provides an interesting and intuitive benchmark, a richer empirical model is needed to give more concrete policy implications. In the next section we set up such a model and quantify it for the United States.

4 Quantitative Model and Estimation

We now present the fully specified model version, which is a specific case of the model presented in Section 2.

²⁰As Carneiro and Heckman (2003, p.27) write: "Family income and child ability are positively correlated, so one would expect higher returns to schooling for children of high income families for this reason alone." In a famous paper, Altonji and Dunn (1996) find higher returns to schooling for children with more-educated parents than for children with less-educated parents.

²¹See, e.g., Shaked and Shanthikumar (2007, p.18).

4.1 Quantitative Model

4.1.1 Basics

We first specify the underlying heterogeneity. Besides parental income I , individuals differ in $X = (\theta, s, \text{ParEdu}, \text{Region}, \varepsilon^E)$, which captures ability, gender, their parents' education levels, the region in which they live, and an idiosyncratic taste for college. Workers' flow utility in the labor force is parameterized as

$$U^W(c_t, \ell_t) = \frac{\left(c_t - \frac{\ell_t^{1+\epsilon_s}}{1+\epsilon_s}\right)^{1-\gamma}}{1-\gamma},$$

where the labor supply elasticity $\frac{1}{\epsilon_s}$ is allowed to vary by gender. Individuals work until 65 and start at age 18 in case they decide to not enroll in college. Each year, individuals make a labor-supply decision and a savings decision. Life-cycle wage paths depend on ability θ , gender s , education e , and on a permanent skill shock that individuals draw upon finishing education and entering the labor market. We present the details of the wage parameterization in Appendix B.3.

4.1.2 College Problem

We now consider decisions of individuals that are enrolled in college. We assume that students can choose to work part-time, full-time, or not at all. Formally, $\ell_t^E \in \{0, PT, FT\}$. For flow utility in college we assume the following functional form:

$$U^E(c_t, \ell_t^E; X, \varepsilon_t^{\ell^E}) = \frac{c_t^{1-\gamma}}{1-\gamma} - \kappa_X - \zeta^{\ell_t^E} + \varepsilon_t^{\ell_t^E}.$$

The term κ_X is the deterministic component of the psychic cost of attending college. Workers of higher ability may find college easier and more enjoyable and therefore may have lower psychic costs of college. Furthermore, children with parents who attended college may find college easier, as they can learn from their parents' experiences. Finally, we allow the psychic cost of college to vary by an agent's gender, to reflect differences in college-going rates across genders. We therefore parameterize the psychic cost term as

$$\kappa_X = \kappa_0 + \kappa_\theta \log(\theta) + \kappa_{\text{fem}} \mathbb{I}(s = \text{female}) + \kappa_{\text{ParEdu}} \text{ParEdu}.$$

The term $\zeta^{\ell_t^E}$ is the cost of working ℓ_t^E hours in college,²² and $\varepsilon_t^{\ell_t^E}$ is a shock associated with continuing college and working ℓ_t^E hours. This represents any idiosyncratic factors associated with staying in college and working that are not captured elsewhere in the model. We assume that the idiosyncratic preference shocks for students, $\varepsilon_t^{\ell_t^E}$, are distributed with a nested logit structure, with a separate nest for the three options involving continuing in college and a

²²We normalize $\zeta^0 = 0$ w.l.o.g.

separate nest for dropping out of college. We denote the nesting parameter by λ and the scale parameter by σ^{ℓ^E} . Given these assumptions, one can define the choice-specific Bellman equations of an agent, depending on their labor supply choices $V_t^{E,\ell^E}(X, I, a_t, \varepsilon_t^{\ell^E})$. For brevity, we do so in Appendix B.4, since it is just a specific case of the problem from Section 2.1. We now turn to the value of staying in college, dropping out of college and enrolling in college initially. Note that these are all just special cases of the value functions presented in Section 2.1. The value of staying enrolled is the maximum of the three labor supply options:

$$V_t^{ND}(X, I, a_t, \varepsilon_t) = \max \left\{ V_t^{E,0}(X, I, a_t, \varepsilon_t^0), V_t^{E,PT}(X, I, a_t, \varepsilon_t^{PT}), V_t^{E,FT}(X, I, a_t, \varepsilon_t^{FT}) \right\}$$

where ε_t is the vector of choice-specific preference shocks. At the beginning of each period, the agent must either choose to drop out of college or continue in college. We parameterize the psychic cost of dropping out as $d(\varepsilon_t) = \delta - \varepsilon_t^D$, where δ is the deterministic part of the dropout cost and ε_t^D is the idiosyncratic part. Therefore, we can write $V_t^D(X, I, a_t, \varepsilon_t^D) = \mathbb{E} [V_t^W(X, I, e = D, a_t, w_t)] - \delta + \varepsilon_t^D$. As in Section 2.1, an agent's problem at the beginning of the period is to choose whether or not to drop out: $V_t^E(X, I, a_t, \varepsilon_t) = \max \{ V_t^D(X, I, a_t, \varepsilon_t^D), V_t^{ND}(X, I, a_t, \varepsilon_t) \}$.

At the beginning of the model, children must decide whether to enter college or to enter the labor market directly. Let $v(X) = \varepsilon^E$ represent idiosyncratic taste for college that is unreflected elsewhere in the model and is observed by the agent before their enrollment choice. We consider ε^E to be a random, idiosyncratic component of the nonpecuniary benefits of college enrollment, in addition to the deterministic psychic cost κ_X . We assume that ε^E is distributed as type I extreme value with scale parameter σ^E . Given this, the value of enrolling in college is

$$V^E(X, I) = \mathbb{E} [V_1^E(X, I, a_1 = 0, \varepsilon_1)] + \varepsilon^E$$

As before, an agent enrolls if $V^E(X, I) > V^H(X, I)$. For the remainder of the paper, it will be useful to separate the elements of the vector X that are observable to the econometrician from the idiosyncratic enrollment draw ε^E . We therefore let $\tilde{X} = (\theta, s, \text{ParEdu}, \text{Region})$.

4.1.3 Parent's Problem

In Section 2 we modeled parental transfers in a general reduced form fashion. Now we provide an explicit microfoundation where we model the parental life-cycle decision problem. Each year the parent makes a consumption/saving decision. The parent also chooses how much to transfer to the child dependent on the child's education choice.²³ Therefore, the parent has to trade off the utility of helping their child through parental transfers with their own consumption. Parents make transfers to their child in the year in which a child graduates from high

²³Note that this also implies that high school transfers may also be endogenous with respect to financial aid. We account for this in the calculation of optimal policy but find it to be economically unimportant quantitatively.

school. We assume that parents commit to a transfer schedule before the child's idiosyncratic enrollment benefit, ε^E , is realized. This simplifies the model solution considerably.²⁴ For all years when the transfer is not given the parent simply chooses how much to consume and save.²⁵ The parent's Bellman equation and details on the calibration of life-cycle parental earnings are given in Appendix B.5.²⁶ In the main body, we only elaborate on the portion of the utility function that arises due to transfers.

In the year of the transfer, the parent receives utility from transfers. Let $F\left(tr^H, tr^E, \tilde{X}, I\right)$ represent the expected utility the parent receives from the transfer schedule tr^H, tr^E , conditional on a child with observable characteristics and parental income (\tilde{X}, I) .

$$F\left(tr^H, tr^E, \tilde{X}, I\right) = \underbrace{\omega \mathbb{E}\left[V\left(X, I, tr^H, tr^E\right)\right]}_{\text{Altruism}} + \mathbb{E}\left[\underbrace{\left(\xi_0 + \xi_{ParEdu}\right) \mathbb{1}_E}_{\text{Paternalism}} + \underbrace{\phi \frac{\left(c_b + tr^e\right)^{1-\gamma}}{1-\gamma}}_{\text{Warm Glow}}\right]$$

where $\mathbb{1}_E$ is a dummy indicating that the child enrolls in college. There are three components, which help to match key features of the relationship between parental transfers, parental income, and the child's problem. First, parents are altruistic, which allows for the possibility that changes in the financial aid schedule crowd out parental transfers. With some abuse of notation, let a child's expected lifetime utility as a function of parental transfers be written as

$$\mathbb{E}\left[V\left(X, I, tr^H, tr^E\right)\right] = \mathbb{E}\left[\max\left\{V^H\left(X, I|tr^H\right), V^E\left(X, I|tr^E\right)\right\}\right],$$

where the expectation is taken over the child's idiosyncratic enrollment benefit, ε^E . The term ω measures the weight the parent places on the child's lifetime expected utility. Second, parents are paternalistic; they receive prestige utility if the child attends college. Allowing for such paternalism allows us to match the level of college transfers relative to transfers for children who forgo college and adds an additional crowding-out element. The parameter ξ_{ParEdu} allows prestige utility to vary by the parent's education level. Specifically, ξ_0 is the prestige utility all parents receive and ξ_{ParEdu} is the additional prestige utility parents receive if at least one of the parents has a college education. Third, parents receive *warm-glow* utility from transfers that is independent of how the transfer affects the child's utility or choices. Allowing for utility from warm-glow helps us to match the gradient between parental income and transfers. Here we adopt the functional form commonly used in the literature (De Nardi, 2004). The parameter ϕ measures the strength of the warm-glow incentive, and c_b measures the extent to which parental transfers are a luxury good.

²⁴If not, the child will have to take into account how parental transfers will respond to their preferences and ability shocks which they partially reveal through their college choice.

²⁵The fact that parents provide all transfers based on the initial enrollment decision can give the incentive to strategically enroll for one year and then drop out directly only to obtain the larger parental transfer. This is one reason for why we incorporated the dropout costs δ , which makes such strategic behavior less attractive. As we show in Section 4.3, our model performs well regarding the dynamics of dropout and graduation.

²⁶We assume that parents exogenously provide transfers to the agent's siblings as well.

4.1.4 The Optimality Condition in the Structural Model

Before turning to the estimation, it is worthwhile to get back to (3), the optimality condition for financial aid, and highlight which structural parameters are key for the relationship between optimal financial aid and parental income in our quantitative model. For brevity and clarity, we focus on the share of marginal and inframarginal enrollees because our numerical analysis below shows that these are the most important forces for our progressivity result.

Inframarginal Enrollees: For brevity we focus on the share of inframarginal enrollees $E(I)$ instead of $\tilde{E}(I)$.²⁷ It is given by:

$$E(I) = \int_{\tilde{X}} \frac{\exp\left(\tilde{V}^E(\tilde{X}, I)/\sigma^E\right)}{\exp\left(\tilde{V}(\tilde{X}, I)/\sigma^E\right) + \exp\left(V^H(\tilde{X}, I)/\sigma^E\right)} dH^*(\tilde{X}|I).$$

where $H^*(\tilde{X}|I)$ is the CDF for \tilde{X} conditional on I and where $\tilde{V}^E(\tilde{X}, I) = V^E(X, I) - \varepsilon^E$ is the value of enrolling in college minus the idiosyncratic taste for college ε^E .²⁸ This expression immediately follows from the fact that the idiosyncratic enrollment benefit ε^E is distributed according to a type I extreme value distribution with scale parameter σ^E . The number of enrollees conditional on (\tilde{X}, I) increases in the difference in the value functions of attending college or not. How $E(I)$ varies with parental income is largely determined by the relation of parental income with (i) psychic costs κ , (ii) parental transfers, (iii) ability. In Section 6.1 we provide a model-based decomposition which addresses the importance of the different elements (i)-(iii).

Marginal Enrollees: For a given (\tilde{X}, I) , the share of marginal enrollees is given by

$$\frac{\partial E(\tilde{X}, I)}{\partial \mathcal{G}(I)} = \frac{E(\tilde{X}, I) \left(1 - E(\tilde{X}, I)\right)}{\sigma^E} \frac{\partial V^E(\tilde{X}, I)}{\partial \mathcal{G}(I)},$$

where $E(\tilde{X}, I)$ is the enrollment share of individuals with observable characteristics \tilde{X} and income I , $\frac{E(\tilde{X}, I)(1-E(\tilde{X}, I))}{\sigma^E}$ is the density of the enrollment benefit parameter ε^E at the value where an (\tilde{X}, I) individual is indifferent between enrolling in college or not. Formally, this threshold is given by $\tilde{\varepsilon}^E(\tilde{X}, I) = \tilde{V}^E(\tilde{X}, I) - V^H(\tilde{X}, I)$. Intuitively, the higher this density, the more individuals are marginal in their decision and the stronger is the increase in enrollment due to higher financial aid. A property of the extreme value distribution is that the density is maximized if enrollment is at 50%, as is the case also for a normal distribution. Further, the lower the scale parameter σ^E , the higher the share of marginal students *ceteris paribus*.

²⁷The insights would be identical if we were looking at $\tilde{E}(I)$ here but notation would be unnecessarily cumbersome.

²⁸Note that $V^H(\tilde{X}, I) = V^H(X, I)$, with some abuse of notation, because the idiosyncratic preference term ε^E does not affect $V^H(X, I)$

The share of marginal enrollees also depends on how much this threshold $\tilde{\varepsilon}^E(\tilde{X}, I)$ changes due to an increase in financial aid, which is captured by:

$$\frac{\partial V^E(\tilde{X}, I)}{\partial \mathcal{G}(I)} = \mathbb{E} \left[\sum_{t=1}^{t_g^{max}} \beta^{t-1} c_t(\cdot)^{-\gamma} \left(1 + \frac{\partial tr_t^E(\tilde{X}, I, \mathcal{G}(I))}{\partial \mathcal{G}(I)} \right) \prod_{s=1}^t (\mathbb{1}_{V_s^{ND} \geq V_s^D}) \prod_{s=1}^{t-1} (1 - Pr_s^{Grad}(X)) \right].$$

Intuitively, agents with low marginal utility $c_t(\cdot)^{-\gamma}$ during college react more strongly financial aid changes. According to this logic, children with low parental income should be more responsive to increases in financial aid. How much this effect varies with parental income is governed by γ , which we estimate with maximum likelihood. In addition, the stronger the crowding out of the parental transfer ($-\frac{\partial tr_t^E(\tilde{X}, I, \mathcal{G}(I))}{\partial \mathcal{G}(I)}$), the less responsive are individuals *ceteris paribus* since less of the financial aid increase reaches them.

All the key parameters are estimated with maximum likelihood and as we document in Section 4.3 the model performs very well not only in terms of enrollment patterns (targeted moments) but also in terms of replicating quasi-experimental evidence about the impact of grant increases on enrollment which was not targeted. In Section 6.1 we provide a model-based decomposition for how the share of marginal students varies with parental income and show that the correlation between parental income and parental transfers is a key driver in our model for why the share of marginal enrollees is decreasing along the parental income distribution.

4.2 Estimation and Data

To bring our model to the data, we make use of the National Longitudinal Survey of Youth 97 (henceforth NLSY97). A big advantage of this data set is that it contains information on parental income and the Armed Forces Qualification Test score (AFQT-score) for most individuals. The latter is a cognitive ability score for high school students that is conducted by the US army. The test score is a good signal of ability. Cunha et al. (2011), for example, show that it is the most precise signal of innate ability among comparable scores in other data sets. We use the NLSY97 for data on college-going, working in college, dropout, parental transfers, and grant receipts.²⁹ Since individuals in the NLSY97 are born between 1980 and 1984, not enough information about their later-life earnings is available. We therefore also use the NLSY79 to better understand how earnings evolve throughout an agent's life. Combining both data sets has proven to be a fruitful way in the literature to overcome the limitations of each individual data set; see Johnson (2013) and Abbott et al. (2018). The underlying assumption is that the relation between the AFQT score and wages has not changed over

²⁹We calculate parental transfers using the same method as Johnson (2013) which involves summing the amount of money parents give to the child, the amount of money received from family for college related expenditures and the monetary value of living at home if the individual lives with his parents. If a child is living at home in the data, we assume the child additionally receives a transfer equal to the monetary value of living at home. We use estimates of the monetary value of living at home directly from Johnson (2013).

that time period. We use the method of Altonji et al. (2012) to make the AFQT scores comparable between the two samples and different age groups. We define an individual as a college graduate if she has completed at least a bachelor’s degree. An individual is considered enrolled in college in a given year if they report being enrolled in college for at least six months in a given academic year. Individuals who report enrolling for at least one year in a four-year college but do not report a bachelor’s degree are considered dropouts. Agents who never enroll in college are considered as high school graduates. Since individuals in the NLSY97 turn 18 years old between 1998 and 2002, we express all US dollar amounts in year 2000 dollars. We drop individuals with missing values for key variables. We also drop individuals who take off one year or more of college before re-enrolling. These agents constitute 11% of the sample. We allow college tuition to vary by the agent’s region. For the variable Region, we consider the four regions for which we have information in the NLSY: Northeast, North Central, South, and West. An overview of our calibration and estimation procedure is given in Table 1. First of all, to quantify the joint distribution of parental income and ability, we take the cross-sectional joint distribution in the NLSY97. We then proceed in four steps. First, we calibrate and preset a few parameters in Section 4.2.1. Second, we calibrate current US tax and college policies, which we document in Appendices B.1 and B.2, respectively. Third, we estimate the parameters of the wage function, which we document in Appendix B.3. Fourth, we estimate the parameters of the child’s and parent’s utility via maximum likelihood in Section 4.2.2.

4.2.1 Calibrated Parameters

We set the risk-free interest rate to 3% (i.e., $r = 0.03$) and assume that individuals’ discount factor is $\beta = \frac{1}{1+r}$. For the labor supply elasticity, we choose $\epsilon = 5$ for men and $\epsilon = 1.66$ for women, which imply compensated labor supply elasticities of 0.2 and 0.6, respectively.³⁰ We make the assumption that students can only borrow through the public loan system. In the year 2000, dependent undergraduates could borrow \$2,625 during the first year of college, \$3,500 during the second, and \$5,500 during following years up to a maximum of \$23,000. We set these as the loan yearly borrowing limits in our model. Students are eligible for either subsidized Stafford loans, under which the student does not pay interest on the loan while he/she is enrolled in college, or unsubsidized Stafford loans, where the student pays interest on the loan. Students are eligible for subsidized loans if their cost of college exceeds their expected family contribution, which is calculated as a function of parental assets and income, number of siblings, and student assets and income. For simplicity, we follow Johnson (2013), and assume that students with parental income below the sample median are eligible for subsidized loans and therefore do not pay interest on their loans while in college and that

³⁰See Blau and Kahn (2007) for a discussion of labor supply differences across gender. Our results are robust to assuming smaller gender differences in labor supply behavior and also larger differences. The labor supply elasticity is in general not a crucial parameter for optimal financial aid.

Table 1: Parameters and Targets

Object	Description	Procedure/Target
$F(I)$	Marginal distribution of parental income	Directly taken from NSLY97
(θ, I)	Joint and conditional distribution of innate abilities	Directly taken from NSLY97
$r = 0.03$	Interest Rate	
$\epsilon_{Men} = 5$	Inverse Labor Supply Elasticity for Men	
$\epsilon_{Women} = 1.66$	Inverse Labor Supply Elasticity for Women	
$Pr_t^{Grad}(\theta)$	Graduation Probabilities	Directly taken from NSLY97
	Wage Parameters	Estimated from regressions
	Parameters of Child and Parental Utility	Maximum Likelihood (Table 5)
Current Policies		
\bar{L}_t	Yearly Stafford Loan Maximum Values	Value in year 2000
$T(y)$	Current Tax Function	Heathcote et al. (2017)
$\mathcal{G}(\theta, I)$	Need- and Merit-Based Grants	Estimated from regressions

students with parental income above the median receive unsubsidized loans and therefore pay interest on loans while in college. Finally, we allow graduation probabilities to depend on an agent's ability and chose $Pr_t^{Grad}(\theta)$ as the fraction of continuing students with ability θ who graduate each year. In practice, we estimate separate yearly graduation probabilities for students with above median ability and below median ability. We assume that all agents in the model have to graduate after six years by setting $Pr_6^{Grad}(\theta) = 1$ for all ability levels.

4.2.2 Estimation

We estimate the remaining parameters with maximum likelihood. An agent's likelihood contribution consists of 1) the contribution of their initial college choice, 2) the contribution of their labor supply and continuation decision each year in college, and 3) the contribution of their realized parental transfers. We assume that parental transfers are measured with normally distributed measurement error. The set of parameters estimated via maximum likelihood consists of the CRRA parameter, γ , the set of parameters governing the amenity value of college and working in college, κ_X and ζ , the dropout cost, δ , the parameters governing the parent's altruism, paternalism, and warm glow, ω , ξ_0 , ξ_{ParEd} , ϕ and c_b , the parameters governing the distribution of the college enrollment and working in college preference shocks: σ^E , σ^{ℓ^E} and λ , and the standard deviation of the measurement error of parental transfers, σ^{etr} . The likelihood contribution of college enrollment and labor supply in college are given by the logit choice probabilities and the likelihood contribution of parental transfers by the PDF of the normal distribution. As these formulas are relatively standard, we present the full likelihood function in Appendix B.6.

Appendix B.7 provides a discussion of the identification. The maximum likelihood estimates are shown in Table 5 in Appendix B.8. We now discuss the estimates of several of the key parameters. This is kept brief, as the magnitude of the parameters is difficult to interpret

in a vacuum. The parameter γ governs the curvature of the utility function with respect to consumption and plays a key role in determining an agent’s risk aversion. We estimate $\gamma = 1.89$, which is in the middle of the range of estimates from the literature. As we have seen in Section 4.1.4, this parameter, along with the variance of the college-going shock, plays an important role in dictating the elasticity of college enrollment with respect to financial aid. The parameters governing the psychic cost of college are κ_0 , κ_θ , κ_{fem} , and κ_{ParEd} . Our estimates of these parameters imply that the psychic cost of college is decreasing in an agent’s ability and parental education. Furthermore, females have a lower psychic cost of college relative to men, reflecting the fact that women attend college in high numbers despite lower monetary returns than men.

4.3 Model Performance and Relation to Empirical Evidence

4.3.1 Model Fit

Enrollment, Graduation and Dropout. Figure 2 illustrates enrollment as a function of parental income and AFQT scores in percentiles. The solid lines indicate results from the model, and the dashed lines are from the data. The relationships in general are well fitted, though we slightly underestimate both gradients. The overall number of individuals who enroll in college is 38.4% in our sample and 39.4% in our model. In our model, 30.0% of agents graduate from college compared to 27.7% in the data. Data from the US Census Bureau are very similar: in 2009 the share of individuals aged 25-29 holding a bachelor’s degree is 30.6% – a number that comes very close to our data, where we look at cohorts born between 1980 and 1984. In Appendix B.9 we also show that the fit is equally good for graduation rates and when we examine enrollment rates separately by gender. Figure 2(a)

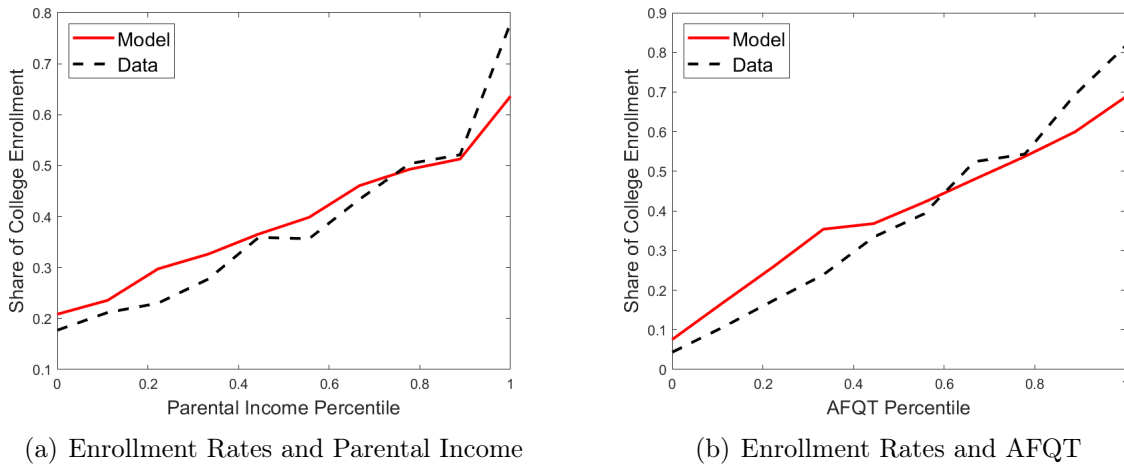




Figure 1: Enrollment Rates

Notes: The solid (red) line shows simulated enrollment shares by parental income and AFQT percentile. This is compared to the dashed (black) line which shows the shares in the data.

shows graduation and dropout fractions over time in the model and the data. The solid red line and the dashed black line show the fraction of the total population that have graduated as a function of number of years of college completed in the model and the data, respectively. In both the model and the data, graduation rates are very low for students with less than three years of college. Graduation shares peak at four years before decreasing. The dashed-dotted blue line and the dotted green line show the fraction of students that drop out in each year in the model and data, respectively. Dropout shares are slightly downward sloping as a function of years in college in both the model and the data. This slope is slightly steeper in the model compared to the data.



(a) Graduation and Dropout Over Time



(b) College Transfers and Parental Income

Figure 2: Model Fit: Graduation, Dropout and Parental Transfers

Notes: The panel on the left shows simulated graduation and dropout rates in the model versus the NLSY97. The panel on the right shows the present value of parental transfers given by parents of college enrollees and non-enrollees in data (NLSY97) versus model.

Parental Transfers. Differences in parental transfers across parental income levels can play a role in generating differential college-going rates across income groups. We analyze the fit of our model with respect to parental transfers in Figure 2(b). We can see that college transfers are strongly increasing in parental income in both the model and data, though our model slightly underestimates the average college transfers in the data. The average college transfer for enrollees with below-median parental income is \$45,000 in the model compared to \$49,000 in the data, while the average college transfer for enrollees with above-median parental income is \$57,000 in the model compared to \$60,000 in the data. The model does a good job of matching the average level of high school transfers. While in our simulations high

Age	Mean Earnings				College Premia		SD(log (y))	
	<i>High-School</i>		<i>College</i>		Model	Data	Model	Data
j	Model	Data	Model	Data	Model	Data	Model	Data
25	22,938	21,348	26,923	25,205	1.17	1.18	0.66	0.59
26	23,747	22,407	29,353	28,300	1.24	1.26	0.67	0.60
27	24,549	23,340	31,829	31,781	1.30	1.36	0.67	0.61
28	25,340	24,022	34,334	33,840	1.35	1.41	0.68	0.62
29	26,117	25,217	36,848	36,254	1.41	1.44	0.69	0.65
30	26,877	25,306	39,354	37,904	1.46	1.50	0.70	0.65
31	27,617	26,449	41,833	40,904	1.51	1.55	0.70	0.66
32	28,334	27,346	44,267	42,954	1.56	1.57	0.71	0.67
33	29,025	28,680	46,639	44,346	1.61	1.55	0.72	0.68
34	29,687	30,494	48,932	46,872	1.65	1.54	0.72	0.67

Notes: Data based on NLSY97 with cohorts born between 1980 and 1984. Mean earnings expressed in year 2000 dollars. Most recent wave from 2015. Model based moment results represent results from estimated model. Zero and small earnings below \$300 a month excluded. SD(log y) equal to standard deviation of log earnings. NLSY97 is top coded at income levels around \$155,000.

Table 2: Earnings Dynamics

school transfers are increasing globally in parental income, parental transfers for high school graduates in the data are decreasing for the highest-income children.³¹

Working During College. We match average hours worked quite well. The average college student in our simulation works 16.21 hours per week compared to 17.39 in the data.³² We observe a weak negative relationship between parental income and working during college in the model and the data.

Earnings and College Premia. Table 2 analyzes the performance of the model with respect to earnings dynamics. We can only compare the model to the NLSY97 data up to age 34 since cohorts in the NLSY97 are born between 1980 and 1984. The simulated mean earnings across ages are very close to those in the data. As described in Section 4, we account for top-coding of earnings data by appending Pareto tails to the observed earnings distribution. As such, average earnings are slightly larger in model as compared to the data. We match college earnings premia very closely until around age 32. After that, the model and data diverge slightly as more and more college students reach top-coded earnings in the NLSY97. In Figure 13 in Appendix B.10, we plot the implied earnings profiles in the model over the

³¹A reasonable suspicion is that this partly reflects measurement error because the set of high-income children who never enroll in college is relatively small. Our parameter estimates were robust ignoring this set of individuals in the estimation.

³²Note that average hours of work are calculated using data from the entire year and thus include work during summer break.

full range of ages.³³ The college-earnings premium averaged across all ages greater than 25 in our model is 85%, that is, the average income of a college graduate is nearly twice as high as the average income of a high school graduate. This is well in line with empirical evidence in Oreopoulos and Petronijevic (2013); see also Lee et al. (2017).

Untargeted Moments. The model successfully replicates quasi-experimental studies. First, it is consistent with estimated elasticities of college attendance and graduation rates with respect to financial aid expansions (Deming and Dynarski, 2009). Second, it is consistent with the causal impact of parental income changes on college graduation rates (Hilger, 2016). Further, our model yields (marginal) returns to college that are in line with the empirical literature (Card, 1999; Oreopoulos and Petronijevic, 2013; Zimmerman, 2014). More details are contained in Appendix B.11.

5 Results: Optimal Financial Aid

5.1 Optimal (Need-Based) Financial Aid

For our first policy experiment, we ask which levels of financial aid for different parental income levels maximize Utilitarian welfare. For this experiment, we consider optimal budget neutral reforms where we do not change taxes or any other policy instrument but instead only vary the targeting of financial aid.³⁴ Additionally, we work under the constraint that financial aid is nonnegative everywhere.³⁵ Figure 3(a) illustrates our main result for the benchmark case. Optimal financial aid is strongly decreasing in parental income. Compared to current policies, financial aid is higher for students with parental income below \$78,000. This change in financial aid policies is mirrored in the change in college graduation, as shown in Figure 3(b). The total graduation rate increases by 2.8 percentage points to 32.8%. This number highlights the efficient character of this reform.

³³The effect of the fatter right tails we include in the model can also be seen in the fit of standard deviation of log earnings. The simulated standard deviation of log earnings is 4-7 log points higher than that in the data from age 25 to age 34.

³⁴At this stage, we leave the merit-based element of current financial aid policies unchanged, that is, we do not change the gradient of financial aid in merit and show the financial aid level for the median ability level. In Appendix C.10, we show that our main result also extends to the case in which the merit-based elements are chosen optimally.

³⁵Relaxing this, one would get a negative subsidy at high parental income levels but nothing substantial changes in terms of results.

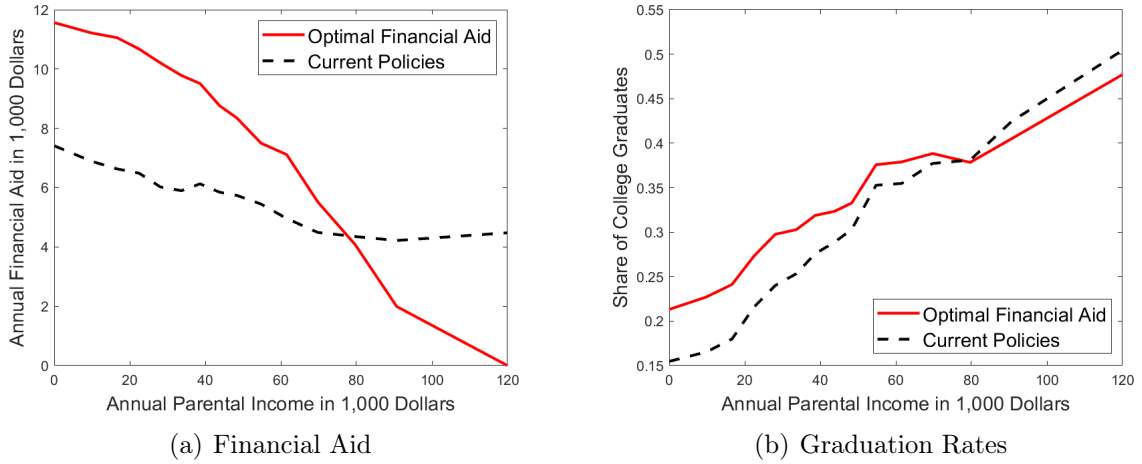


Figure 3: Optimal versus Current Financial Aid

Notes: Optimal financial aid with a Utilitarian welfare function and current financial aid in Panel (a). In Panel (b) we display the college graduation share by parental income group.

5.2 No Desire for Redistribution

One might be suspicious of whether the progressivity is driven by a desire for redistribution from rich to poor students that results in declining welfare weights.³⁶ If this were the case, the question would naturally arise whether the financial aid system is the best means of doing so. However, we now show that the result holds even in the absence of redistributive purposes. We modify the social planner’s problem such that the marginal social welfare weights are constant across parental income levels, i.e. $\frac{\partial W^E(I)}{\partial I} = 0$. In this case, the social planner values a dollar transferred to any inframarginal student equally, independent of the student’s marginal utility of consumption or level of parental crowdout. The results are in Figure 4(a). The optimal financial aid schedule is slightly less progressive than the optimal financial with a Utilitarian welfare function. The implied graduation patterns are illustrated in Figure 4(b). The results show that the social planner’s redistribution motive only plays a minor role in generating progressive optimal financial aid.

5.3 Tax-Revenue-Maximizing Financial Aid

In this section we ask the following question: how should a government that is only interested in maximizing tax revenue (net of expenditures for financial aid) set financial aid policies? Fig-

³⁶In fact, $1 - W^C(I)$, which is the relevant term for the formula, increases from around by a factor of around 2.3 between the 75th and 25th percentile of parental income. Note that this welfare weight is defined such that it accounts for crowding out of parental transfers. In fact, we find that crowding out is stronger for high parental income students. Going from the lowest to the highest parental income, the crowding out rate is monotonically increasing, from 9% at the 25th percentile to 25% at the 75th percentile of parental income. The fact that $1 - W^C(I)$ increases by a factor of 2.3 is hence not only due to the Utilitarian welfare function but also due to the fact that an increase in financial aid for the poorest students will be crowded out much less than financial aid for the richest students.

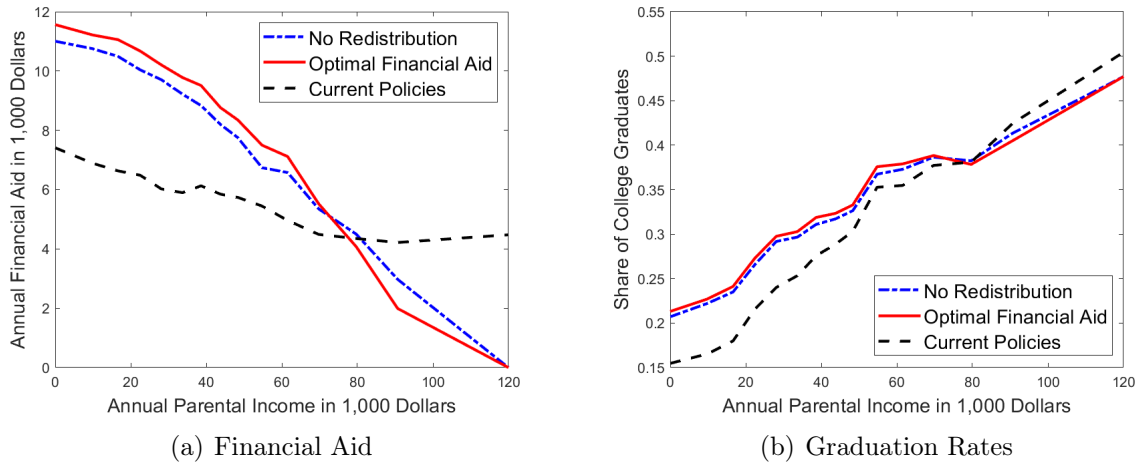


Figure 4: Financial Aid Policies with no Redistribution Motive

Notes: The dashed-dotted (blue) line shows the optimal schedule for a social planner with no redistribution motive. Optimal financial aid with a Utilitarian welfare function and current financial aid are also shown for comparison in Panel (a). In Panel (b) we display the college graduation share by parental income group for each of the three scenarios.

Figure 5(a) provides the answer: revenue-maximizing financial aid in this case is very progressive as well. Whereas the overall level of financial aid is naturally lower if the consumption utility of students is not valued, the declining pattern is basically unaffected. For lower parental income levels, revenue-maximizing aid is more generous than the current schedule, which implies that an increase must be more than self-financing. We study this in more detail in Section 5.4. The implied graduation patterns are illustrated in Figure 5(b).

5.4 Self-Financing Reforms

An increase in financial aid can be self-financing if properly targeted. The solid red line in Figure 6 illustrates the fiscal return, that is, the net effect on government revenue were financial aid for a particular income level to be increased by \$1. For example, a 40% return implies that the net present value increase in tax revenue is 40% larger than the cost of increasing financial aid. Returns are positive for parental income between \$0 and \$33,000; the latter number corresponds to the 32nd percentile of the parental income distribution. This result is striking: increasing subsidies for this group is a free lunch. An alternative would be to consider reforms where financial aid is increased for students below a certain parental income level. This case is illustrated by the dashed-dotted blue line in Figure 6. An increase in financial aid targeted to children with parental income below \$54,000 – corresponding to the 59th percentile – is slightly above the margin of being self-financing.

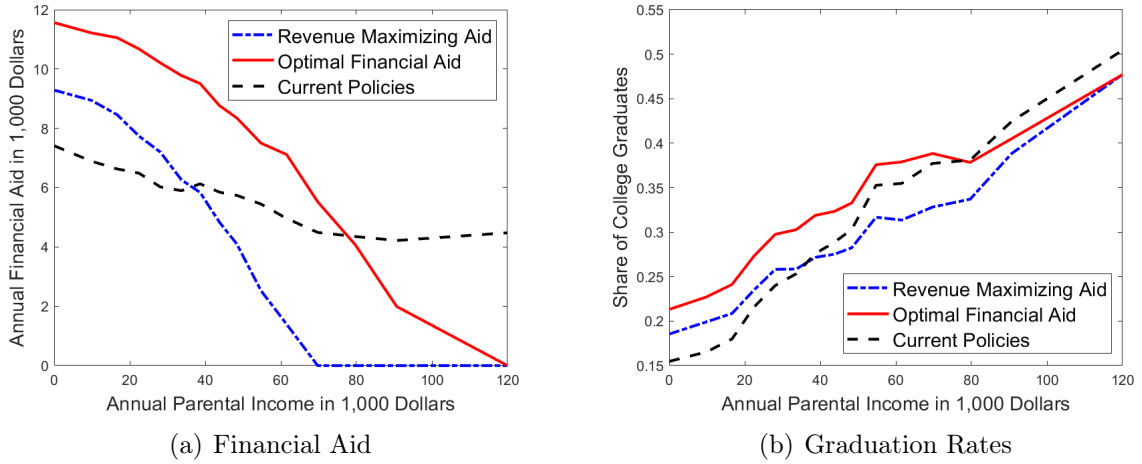


Figure 5: Tax-Revenue-Maximizing Financial Aid Policies

Notes: The dashed-dotted (blue) line shows the optimal schedule under the objective of maximizing net-tax revenue (net of expenditures for financial aid). Optimal financial aid with a Utilitarian welfare function and current financial aid are also shown for comparison in Panel (a). In Panel (b) we display the college graduation share by parental income group for each of the three scenarios.

6 Why Are Optimal Policies Progressive?

We have just shown in Section 5 that optimal financial aid is progressive and more so than the current US policies. In Sections 5.2 and 5.3, we have also shown that the results are not driven by the desire to redistribute from richer to poorer students. We now explore the key forces determining this progressivity result. Recall that the change in welfare due to a small increase of $\mathcal{G}(I)$ is given by (3)

$$\underbrace{\frac{\partial E(I)}{\partial \mathcal{G}(I)} \times \Delta \mathcal{T}^E(I)}_{\text{Enrollment Effect}} + \underbrace{\frac{\partial C(I)}{\partial \mathcal{G}(I)} \Big|_{E(I)}}_{\text{Completion Effect}} \underbrace{E(I) \times \Delta \mathcal{T}^C(I) - \tilde{E}(I) (1 - W^E(I))}_{\text{Mechanical Effect}}.$$

To explain why the optimal financial aid schedule is more progressive than the current US financial aid, we illustrate the two most important determinants of this welfare effect: the enrollment effect and the mechanical effect evaluated at the current US system financial aid. We only found a quantitatively very small contribution of the completion effect and therefore focus on the other two effects. Figure 7(a) plots the increase in enrollment for a \$1,000 increase starting from the current financial aid system against parental income. The curve is decreasing in income – children with low parental income react more strongly. This contributes to the result that optimal financial aid is more progressive than the current US benchmark. By contrast, the fiscal externality $\Delta \mathcal{T}^E(I)$ is increasing in parental income because (i) marginal

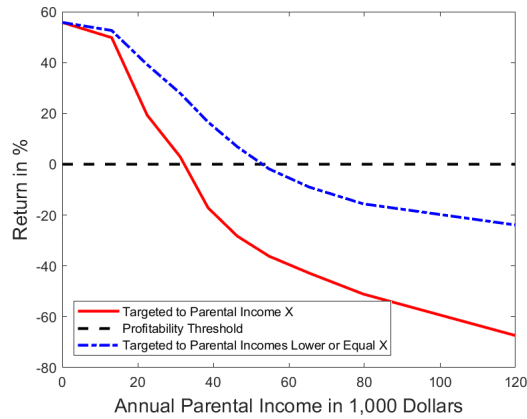


Figure 6: Fiscal Returns on Increase in Financial Aid

Notes: The dashed-dotted (blue) line shows the net fiscal return for a \$1 increase in financial aid targeted to all students with a parental income level lower than X . It crosses the profitability line at \$54,000 – corresponding to the 59th percentile. The solid (red) line shows the net fiscal return for a \$1 increase in financial aid targeted to all students with a parental income level equal to X . It crosses the profitability line at \$33,000 – corresponding to the 32nd percentile.

enrollees from higher income households have higher returns³⁷ and (ii) the fiscal externality is higher for children with high income parents because they receive less financial aid. We now turn to the mechanical effect. Since we have already shown in Sections 5.2 and 5.3 that the redistributive preferences play a minor role, we turn again to the of inframarginal enrollees as plotted in Figure 1(a). As discussed above, there is a strong parental income gradient, as the simulated share of enrollees increases from around 21% to around 63%. Note that this basically implies that the direct marginal fiscal costs of a grant increase by a factor of three with parental income.

Summing up, both the increasing share of inframarginal students and the declining share of marginal students are important drivers for why optimal financial aid is more progressive than current financial aid. An open question is what exactly drives how the share of marginal and inframarginal students vary with parental income. We now provide a model-based decomposition to shed light on the key drivers.

6.1 Relationship between Inframarginal Students, Marginal Students, and Parental Income

We have just seen that the following two features mainly explain why optimal financial aid is more progressive than current financial aid. First, students with low parental income are more likely to be on the margin of enrolling in college, therefore an increase in financial aid

³⁷The relationship between parental income and the average ability of marginal students depends on how strongly college enrollees are selected on ability. Ultimately, we find that average ability of marginal enrollees is increasing in parental income. As the college wage premium is increasing in ability, this implies that increase of tax payments of marginal enrollees is increasing in parental income.




Figure 7: Marginal Students and Tax Revenue Changes

Notes: In (a), we plot the change in enrollment rates for a simulated \$1,000 change in financial aid for each parental income level. The average (across all individuals in the sample) is 1.69 percentage points. In (b), we show the implied average fiscal externality across all students who are marginal w.r.t. the financial aid increase.

targeted at low income families will induce larger increases in college enrollment. Second, the positive relation between college enrollment and parental income is strong and therefore the direct fiscal costs of increasing financial aid is lower for children with low parental income. We now perform a model-based decomposition exercise to better understand which factors drive these two relationships. For this decomposition, all changes to the model specification are cumulative. That is, each new model specification contains the same model alterations as the previous specification.³⁸

To isolate the effects of model primitives, we perform this decomposition for a hypothetical flat financial aid schedule instead of the current US tax schedule. This allows us to isolate the influence of the model primitives, instead of mixing the effects of current policies and model primitives. Specifically, we set the aid for all parental income groups to the mean level of financial aid in the data. Results are similar if the decomposition is performed for the current financial aid system, as we document in Appendix C.1. We start by focusing on the relationship between parental income and college enrollment in Figure 8(a).³⁹ The solid line captures the baseline case. College enrollment rates are strongly increasing in parental income: 69% of students at the top of the parental income distribution enroll in college compared to only 17% at the bottom of the distribution.⁴⁰ One factor that leads to this positive relationship is the correlation between parental income and ability. To understand the contribution of this

³⁸In Appendix C.2, we consider an alternative decomposition in which we equalize parental transfers first. In Appendix C.3, we consider a decomposition where we remove the borrowing constraint before we equalize parental transfers.

³⁹Alternatively, we could have focused on graduation instead of enrollment. The implications are very similar.

⁴⁰Note that this relationship is stronger than the one in Figure 1(a) where the current financial aid schedule is used as benchmark.

correlation toward differential college-going rates by parental income, we simulate a version of the model in which we remove the correlation between parental income and ability by drawing each agent’s ability from the unconditional ability distribution. Recall that ability affects both the returns to college and the psychic costs of attending college. Figure 8(a) shows that the relation between college enrollment and parental income reduces substantially, with 27% of children at the bottom of the income distribution enrolling in college compared to 69% of children from the top of the income distribution.

Additionally, children with higher parental income are more likely to go to college because parental income is positively correlated with parental education. Since higher parental education lowers psychic costs, this implies a negative correlation between parental income and psychic costs. We remove the relation between parental education and psychic costs in college by setting $\kappa_{ParEd} = 0$, in addition to removing the correlation between parental income and ability.⁴¹ After removing these differences in psychic costs, the relationship between parental income and college enrollment becomes again flatter with 33% of children from the bottom of the income distribution and 48% from the top of the income distribution.

In our model, there are further factors that influence the parental income gradient in college education. The individual returns to college are not known at the time of the enrollment decision. As individuals are risk averse and as parents with higher income levels give higher transfers for students attending college, this riskiness of college is another mechanism which can generate a positive relationship between college and parental income. In addition to the modifications above, we remove this risk in the monetary return to college by simulating a version of the model in which each agent with certainty receives a fixed labor market ability draw. Removing the riskiness of college leads to a further flattening of the relationship between parental income and college enrollment. Yet there is still a gradient as enrollment increases from 35% to 48% due to the fact that high parental income children obtain more transfers from their parents. We finally remove this relationship by providing all children the mean parental transfer levels for enrollees and non-enrollees and assuming that no families are eligible for subsidized Stafford loans. As a consequence, the relationship between parental income and college enrollment becomes flat.⁴² We conclude that all components play an important role for the increasing share of inframarginal students with the exception of the risk channel.

We turn to the determinants of the negative relation between parental income and share of marginal enrollees (again considering a \$1,000 increase in financial aid) in Figure 8(b). The solid line shows the relationship between parental income and the share of marginal students in the baseline case. The share of marginal students is decreasing in parental income. The dotted line and the dash-dotted line show the cases in which we remove the correlation between

⁴¹Furthermore, we set κ_0 so that the average psychic cost of going to college is unchanged.

⁴²The fact that it is not totally flat is due to the fact that individuals still differ in gender and region and these variables are not distributed in exactly the same way for each parental income group.

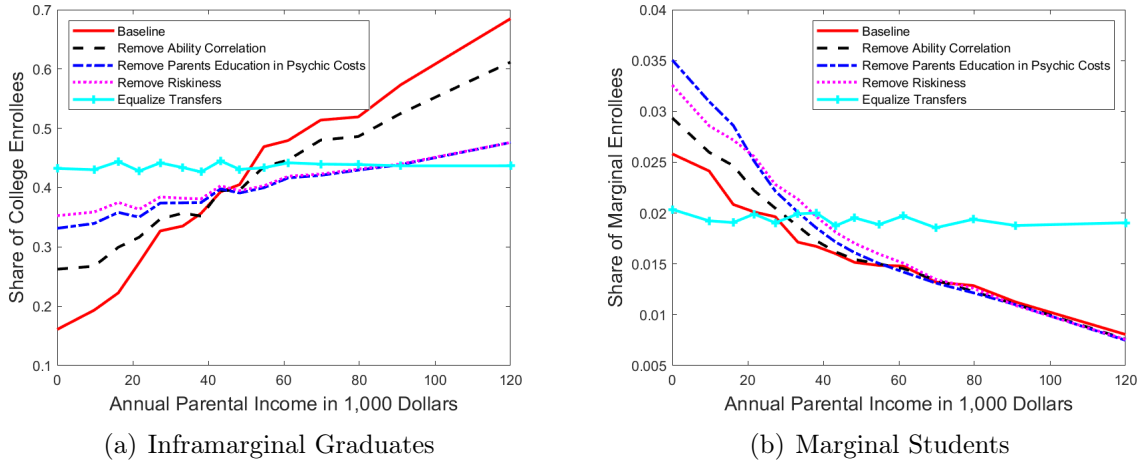


Figure 8: Model-Based Decomposition for Marginal and Inframarginal Students

Notes: We plot the share of college enrollees and marginal college enrollees given a flat financial aid schedule for different model specifications. The solid red line represent the baseline model (but also with the flat financial aid schedule). For the dashed black line we simulate a model version for which we remove the correlation between ability and parental income. For the dashed-dotted blue line we simulate a model version for which we remove additionally the correlation between the psychic costs and parental education. For the dotted pink line we simulate a model version for which we additionally removes labor market riskiness; i.e. education decisions are made with no uncertainty about future wages. For the turquoise line with crosses we simulate a model version for which we set parental transfers to the mean parental transfers in the data, conditional on education.

parental income and ability and in which we remove the relation between parental income and parental education, respectively. In both cases, the share of marginal students among low income families increases slightly. One reason why this occurs is because the enrollment share of low income families moves closer to 50%. As we have seen in Section 4.1.4, 50% corresponds with the mode of the extreme value distribution, which implies a higher marginal share, all else equal. Removing the riskiness of college returns (the dotted line) reduces the share of marginal low income students slightly, but the relation between parental income and the share of marginal students is still strongly decreasing. Finally, we set parental transfers to the mean levels for enrollees and non-enrollees and assume no families are eligible for subsidized Stafford loans. The relation between parental income and the share of marginal enrollees disappears completely when there is no relationship between parental income and the child's financial resources. We conclude that the correlation between parental transfers and parental incomes play the strongest role in differences in the share of marginal students across parental income levels.

6.2 Recalculation of Optimal Policies

This decomposition has illustrated the key factors in the positive relation between parental income and college education and the negative relation between parental income and the




Figure 9: Optimal Financial Aid for Different Model Specifications

Notes: For each model specification (see Figure 8), we illustrate the respective optimal financial aid schedule.

share of marginal students. Removing the different elements from the model also affects the other forces that determine the optimal financial aid schedule. To get the complete picture, we therefore now simulate the respective optimal financial aid schedule for all these model specifications. Figure 9 shows the implied optimal policies for each model specification. First, the positive relation of parental income with ability and parental education are not the main drivers of the financial aid results. Although graduation rates flatten (see the lines in Figure 8(a)), there are offsetting effect as low-income children are now more likely to be marginal (see the lines in Figure 8(b)). Policies are still relatively progressive in those models.

Removing riskiness in addition has a stronger impact, because there are no offsetting effects in this case. Both the inframarginal and marginal schedule are flatter now. Finally, the turquoise crossed line shows an almost zero slope, as transfers are equalized and we are in a world where parental income plays no more role. The effects for marginal students and inframarginal students again work in tandem, pushing towards flat aid.⁴³

From this decomposition, we conclude that the correlations between parental income and parental transfers, psychic costs, and ability all play important roles in the progressive optimal aid schedule with parental transfers probably playing the biggest role. The correlation of parental income with parental transfers drives the negative correlation of parental income and share of marginal students and also plays a role in the positive correlation of parental income and the share of inframarginal students. The relationships between parental income and ability and psychic costs play important roles in the correlation between share of inframarginal students and parental income, and therefore also play important roles in the progressive optimal aid schedule.⁴⁴

⁴³Again, the fact that it financial aid not totally flat is due to the fact that individuals still differ in gender and region and these variables are not distributed in exactly the same way for each parental income group.

⁴⁴In Appendix C.2 we consider a decomposition in which we first remove the correlation with parental transfers before removing the correlations with ability and psychic costs. We reach similar conclusions. In Appendix C.3, we consider a decomposition which is as in the main body but we also remove borrowing

7 Extensions

We consider five important extensions: the role of borrowing constraints, endogenous abilities of children, general equilibrium effects, endogenous optimal taxation, and merit-based aid. The latter three are found in the appendix, because of space constraints. General equilibrium effects can be found in Appendix C.8, endogenous optimal taxation can be found in Appendix C.9, and merit-based aid can be found in Appendix C.10.

7.1 The Role of Borrowing Constraints

We have shown that optimal progressivity is not primarily driven by redistributive tastes but rather by efficiency considerations in Section 5.3. Given that our analysis assumes that students cannot borrow more than the Stafford Loan limit, the question arises whether these efficiency considerations are driven by borrowing limits that should be particularly binding for low-parental-income children. To elaborate on this question, we ask how normative prescriptions for financial aid policies change if students can suddenly borrow as much as they want (up to the natural borrowing limit, which is not binding). For this thought experiment, we first remove borrowing constraints and keep the current financial aid system. This will increase college enrollment and imply a windfall fiscal gain for the government. In a second step, we choose optimal financial aid but restrict the government to not use this windfall gain. As we show in Section C.4, optimal financial aid policies become less progressive in this case. This is expected. More low-income children are close to the borrowing constraint in the baseline specification. When we remove borrowing constraints, redistributing funds towards these students becomes less attractive for the utilitarian social planner. Quantitatively, however, optimal policies are still very progressive even when borrowing constraints are removed. We also re-estimated a version of the model in which borrowing constraints varied by parental resources. We found that the optimal financial aid schedule was very similar to the baseline schedule. Details can be found in Appendix C.5. We also considered alternative unreported versions, where exogenous borrowing constraints depend differently on characteristics of the child and the parent. The policy implications were not affected much.

7.2 Endogenous Ability

Up to this point, we have assumed that a child's ability at the beginning of the model, θ , is exogenous. One might be concerned that parents may respond to changes in the financial aid schedule by adjusting their investment in their child's development, therefore changing their child's ability at the time of the college entrance decision. To better understand how the optimal financial aid schedule would differ if ability were endogenous with respect to financial

constraints as an intermediate case before removing parental transfers. In this case the role of parental transfers is much more limited.

aid, we posit a model extension in which a child’s ability is determined endogenously as a function of parental investment.

Children are endowed with an initial ability at birth θ_0 , where θ_0 is a random variable with CDF $\theta_0 \sim F_{\theta_0}(\cdot|I)$. A child’s ability at the time of college, θ , is produced as a function of the child’s initial ability and parental monetary investment, *Invest*. The parent observes θ_0 at the beginning of the child’s life and then chooses investment in the child. Additionally, the parent chooses parental transfers when the child attends college, as in the baseline model. For simplicity, we assume that grants are only a function of income when solving the model with endogenous ability. This considerably simplifies the model solution.

For the production of the child’s ability, we assume the following functional form, which is very similar to and based on the translog functional form employed in Agostinelli and Wiswall (2016)⁴⁵

$$\theta = \ln A + \gamma_1 \ln \theta_0 + \gamma_2 \ln Invest + \gamma_3 \ln \theta_0 \ln Invest + \iota,$$

where ι is a normally distributed error that is unknown by the parent at the time of choosing *Invest*, and where A , γ_1 , γ_2 , and γ_3 are parameters of the ability production function. After the parent chooses *Invest*, the ability production shock ι is realized. The parent’s problem is then the same as in the baseline case: each year, the parent continues to make consumption/saving decisions and chooses a parental transfer schedule when the agent reaches the college enrollment choice. Therefore, increases in early childhood investment increase the child’s expected ability, but come at the cost of reduced consumption for the parent and potentially lower transfers when the child reaches the enrollment decision.

We calibrate the parameters of the childhood ability production function to match the joint distribution of parental income and ability we observe in our data and selected moments from Agostinelli and Wiswall (2016). Details on the calibration are included in Appendix C.6. Dahl and Lochner (2012) use changes in the EITC to instrument for family income and find that a \$1000 increase in family income leads to an increase in ability scores by 6% of a standard deviation. We simulate an increase in yearly family income of parents by \$1,000 in our model. The increase in income leads to an average increase in AFQT scores of 2.2% of a standard deviation across all children, and an increase of 5.1% of a standard deviation for children in the lowest quintile. Therefore, the simulated responsiveness of ability with respect to parental income is slightly smaller than but in line with Dahl and Lochner (2012).

The optimal financial aid schedule, graduation rates, and ability levels with endogenous ability are shown in Figure 10. Panel 10(a) shows the new optimal financial aid schedule when ability is endogenous. Compared to the baseline case when ability is exogenous, the

⁴⁵Agostinelli and Wiswall (2016) estimate a model of early childhood developments with multiple periods in which childhood skills are latent. Additionally, they use a broader concept of parental investment; the investment we refer to here is strictly monetary.

optimal aid schedule is now much higher, reflecting that increases in financial aid are now much more profitable for the government. With endogenous ability, increases in financial aid lead to increases in child ability, which increase tax payments of *both* marginal and inframarginal children. The optimal aid schedule is still highly progressive. Panel 10(b) shows the graduation rates evaluated at the optimal aid schedule with endogenous ability. Switching to the optimal schedule leads to an increase in college graduation rates of over 10%, reflecting that 1) the optimal schedule is considerably more generous than the current schedule and 2) increases in financial aid lead to larger increases in college-going when ability is endogenous.⁴⁶

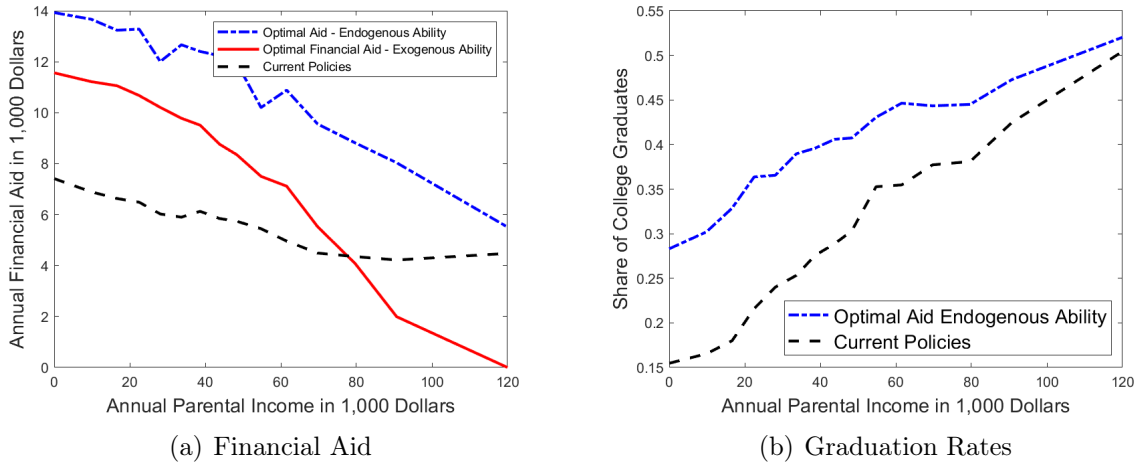


Figure 10: Financial Aid and Graduation with Endogenous Ability

Notes: The dashed-dotted (blue) line shows the optimal schedule when child's ability is endogenous. Optimal financial aid with exogenous ability and current financial aid are also shown for comparison in Panel (a). In Panel (b) we display the college graduation share by parental income group when ability is endogenous with the optimal aid schedule and with the current aid schedule.

As emphasized by Caucutt and Lochner (2017) and Lee and Seshadri (2019), low income parents may face borrowing constraints when their child is young and therefore unable to adjust their investment in their child's ability in response to changes in education policy. In Appendix C.7, we calculate the optimal aid schedule with endogenous ability under the assumption that low income parents may be borrowing constrained when their child is young. We find that the optimal progressivity of the system decreases as we increase the percentage of low-education families who are borrowing constrained. However, the optimal schedule remains more progressive than the current schedule in all cases.

⁴⁶We show the end of childhood ability as a function of parental income under the current financial aid schedule and under the optimal schedule in Section C.6. Switching to the optimal aid schedule leads to substantial increases in child ability, especially for children in the lower end of the parental income distribution.

8 Conclusion

This paper has analyzed the normative question of how to optimally design financial aid policies for students. We find that optimal financial aid policies are strongly progressive. This result holds for different social welfare functions, assumptions on credit markets for students, and assumptions on income taxation. Moreover, we find that a progressive expansion in financial aid policies could be self-financing through higher tax revenue, thus benefiting all taxpayers as well as low-income students directly. It seems to be that financial aid policies are a rare case with no classic equity-efficiency trade-off because a cost-effective targeting of financial aid goes hand in hand with goals of social mobility and redistribution. We also think that our results can be used for policy recommendations according to the criteria of Diamond and Saez (2011):⁴⁷ the economic mechanism is empirically relevant and of first order importance to the problem, it is very robust and progressive financial aid systems are clearly implementable, as they are universal across all OECD countries.

Future work could focus on adding heterogeneity in the quality of colleges, which would allow for rich interactions with financial aid policies. In such a setting, it would seem natural to let the government optimize over financial aid as a joint function of parental background and college quality. In addition, college quality could adjust endogenously. Thinking more seriously about these issues could also extend the scope of the analysis to the level of community colleges. We leave that for future research.

References

- ABBOTT, B., G. GALLIPOLI, C. MEGHIR, AND G. L. VIOLANTE (2018): "Education Policy and Intergenerational Transfers in Equilibrium," *Cowles Foundation Discussion Paper No. 1887R*.
- AGOSTINELLI, F. AND M. WISWALL (2016): "Estimating the Technology of Children's Skill Formation," *NBER Working Paper No. 22442*.
- ALSALAM, N. (2013): "The Federal Pell Grant Program: Recent Growth and Policy Options," *Washington, DC: Congressional Budget Office*.
- ALTONJI, J. G., P. BHARADWAJ, AND F. LANGE (2012): "Changes in the Characteristics of American Youth: Implications for Adult Outcomes," *Journal of Labor Economics*, 30, 783–828.
- ALTONJI, J. G. AND T. A. DUNN (1996): "The Effects of Family Characteristics on the Return to Education," *Review of Economics and Statistics*, 78, 692–704.
- BAGNOLI, M. AND T. BERGSTROM (2005): "Log-Concave Probability and Its Applications," *Economic Theory*, 26, 445–469.

⁴⁷Diamond and Saez (2011) write in their abstract: "We argue that a result from basic research is relevant for policy only if (a) it is based on economic mechanisms that are empirically relevant and first order to the problem, (b) it is reasonably robust to changes in the modeling assumptions, (c) the policy prescription is implementable (i.e., is socially acceptable and is not too complex)."

- BAUM, S., J. MA, AND K. PAYEA (2013): *Education Pays, 2013: The Benefits of Higher Education for Individuals and Society. Trends in Higher Education Series.*, New York: The College Board.
- BLAU, F. D. AND L. M. KAHN (2007): “Changes in the Labor Supply Behavior of Married Women: 1980–2000,” *Journal of Labor Economics*, 25, 393–438.
- BOHACEK, R. AND M. KAPICKA (2008): “Optimal Human Capital Policies,” *Journal of Monetary Economics*, 55, 1–16.
- BOVENBERG, A. L. AND B. JACOBS (2005): “Redistribution and Education Subsidies Are Siamese Twins,” *Journal of Public Economics*, 89, 2005–2035.
- CARD, D. (1999): “The Causal Effect of Education on Earnings,” *Handbook of Labor Economics*, 3, 1801–1863.
- CARD, D. AND T. LEMIEUX (2001): “Can Falling Supply Explain the Rising Return to College for Younger Men? A Cohort-Based Analysis,” *The Quarterly Journal of Economics*, 116, 705–746.
- CARNEIRO, P. AND J. J. HECKMAN (2003): “Human Capital Policy,” *IZA Discussion Paper No. 821*.
- CARROLL, S. J. AND E. ERKUT (2009): *The Benefits to Taxpayers from Increases in Students’ Educational Attainment*, vol. 686, Santa Monica, CA: RAND Corporation.
- CASTLEMAN, B. L. AND B. T. LONG (2016): “Looking beyond Enrollment: The Causal Effect of Need-Based Grants on College Access, Persistence, and Graduation,” *Journal of Labor Economics*, 34, 1023–1073.
- CAUCUTT, E. M. AND L. LOCHNER (2017): “Early and Late Human Capital Investments, Borrowing Constraints, and the Family,” Tech. rep.
- COLLEGE BOARD (2015): *Trends in Student Aid 2015. Trends in Higher Education Series.*, New York: The College Board.
- CUNHA, F., J. HECKMAN, AND S. NAVARRO (2005): “Separating Uncertainty from Heterogeneity in Life Cycle Earnings,” *Oxford Economic Papers*, 57, 191–261.
- CUNHA, F., F. KARAHAN, AND I. SOARES (2011): “Returns to Skills and the College Premium,” *Journal of Money, Credit and Banking*, 43, 39–86.
- DAHL, G. B. AND L. LOCHNER (2012): “The Impact of Family Income on Child Achievement: Evidence from the Earned Income Tax Credit,” *American Economic Review*, 102, 1927–1956.
- DE NARDI, M. (2004): “Wealth Inequality and Intergenerational Links,” *Review of Economic Studies*, 71, 743–768.
- DEMING, D. AND S. DYNARSKI (2009): “Into College, Out of Poverty? Policies to Increase the Postsecondary Attainment of the Poor,” *NBER Working Paper 15387*.
- DIAMOND, P. A. AND E. SAEZ (2011): “The Case for a Progressive Tax: From Basic Research to Policy Recommendations,” *Journal of Economic Perspectives*, 25, 165–190.
- DYNARSKI, S. M. (2003): “Does Aid Matter? Measuring the Effect of Student Aid on College Attendance and Completion,” *American Economic Review*, 93, 279–288.
- FINDEISEN, S. AND D. SACHS (2016): “Education and Optimal Dynamic Taxation: The Role of Income-Contingent Student Loans,” *Journal of Public Economics*, 138, 1–21.

- (2017): “Redistribution and Insurance with Simple Tax Instruments,” *Journal of Public Economics*, 146, 58–78.
- GOLOSOV, M., A. TSYVINSKI, AND N. WERQUIN (2014): “A Variational Approach to the Analysis of Tax Systems,” *Working Paper*.
- GUNER, N., R. KAYGUSUZ, AND G. VENTURA (2014): “Income Taxation of US Households: Facts and Parametric Estimates,” *Review of Economic Dynamics*, 17, 559–581.
- HEATHCOTE, J., K. STORESLETTEN, AND G. L. VIOLANTE (2017): “Optimal Tax Progressivity: An Analytical Framework,” *Quarterly Journal of Economics*, 132, 1693–1754.
- HECKMAN, J. J., L. LOCHNER, AND C. TABER (1998): “Explaining Rising Wage Inequality: Explorations with a Dynamic General Equilibrium Model of Labor Earnings with Heterogeneous Agents U,” *Review of Economic Dynamics*, 1, 158.
- HECKMAN, J. J., L. J. LOCHNER, AND P. E. TODD (2006): “Earnings Functions, Rates of Return and Treatment Effects: The Mincer Equation and Beyond,” *Handbook of the Economics of Education*, 1, 307–458.
- HECKMAN, J. J. AND S. NAVARRO (2007): “Dynamic discrete choice and dynamic treatment effects,” *Journal of Econometrics*, 136, 341–396.
- HILGER, N. G. (2016): “Parental Job Loss and Children’s Long-Term Outcomes: Evidence from 7 Million Fathers’ Layoffs,” *American Economic Journal: Applied Economics*, 8, 247–283.
- JACOBS, B. AND U. THUEMMEL (2018): “Optimal Taxation of Income and Human Capital and Skill-Biased Technical Change,” Tech. rep., Mimeo, Tinbergen Institute, Amsterdam/Rotterdam.
- JACQUET, L. AND E. LEHMANN (2016): “Optimal Taxation with Heterogeneous Skills and Elasticities: Structural and Sufficient Statistics Approaches,” *THEMA Working Paper 2016-04*.
- JACQUET, L., E. LEHMANN, AND B. VAN DER LINDEN (2013): “The Optimal Marginal Tax Rates with Both Extensive and Intensive Responses,” *Journal of Economic Theory*, 148, 1770–1805.
- JOHNSON, M. T. (2013): “Borrowing Constraints, College Enrollment, and Delayed Entry,” *Journal of Labor Economics*, 31, 669–725.
- KATZ, L. F. AND K. M. MURPHY (1992): “Changes in Relative Wages, 1963–1987: Supply and Demand Factors,” *Quarterly Journal of Economics*, 107, 35–78.
- KLEVEN, H. J. (2018): “Sufficient Statistics Revisited,” *Working Paper*.
- KOENIGER, W. AND J. PRAT (2017): “Human Capital and Optimal Redistribution,” *Working Paper*.
- KRUEGER, D. AND A. LUDWIG (2013): “Optimal Progressive Labor Income Taxation and Education Subsidies When Education Decisions and Intergenerational Transfers Are Endogenous,” *American Economic Review: Papers & Proceedings*, 103, 496–501.
- (2016): “On the Optimal Provision of Social Insurance: Progressive Taxation versus Education Subsidies in General Equilibrium,” *Journal of Monetary Economics*, 77, 72–98.
- LAUN, T. AND J. WALLENIS (2016): “Social Insurance and Retirement: a Cross-Country Perspective,” *Review of Economic Dynamics*, 22, 72–92.

- LAWSON, N. (2017): “Liquidity Constraints, Fiscal Externalities and Optimal Tuition Subsidies,” *American Economic Journal: Economic Policy*, 8, 313–343.
- LEE, D., S. Y. LEE, AND Y. SHIN (2017): “The Option Value of Human Capital: Higher Education and Wage Inequality,” *Working Paper*.
- LEE, S. Y. AND A. SESHADRI (2019): “On the intergenerational transmission of economic status,” *Journal of Political Economy*, 127, 855–921.
- LEHMANN, E., L. SIMULA, AND A. TRANNOY (2014): “Tax Me If You Can! Optimal Nonlinear Income Tax between Competing Governments,” *Quarterly Journal of Economics*, 129, 1995–2030.
- MIRRLEES, J. A. (1971): “An exploration in the theory of optimum income taxation,” *The Review of Economic Studies*, 38, 175–208.
- OECD (2014): “Education at a Glance 2014: OECD Indicators,” URL: <http://dx.doi.org/10.1787/eag-2014-en>.
- OREOPOULOS, P. AND U. PETRONIJEVIC (2013): “Making College Worth It: A Review of Research on the Returns to Higher Education,” *NBER Working Paper 19053*.
- PIKETTY, T. AND E. SAEZ (2013): “Optimal Labor Income Taxation,” *Handbook of Public Economics, Vol 5*.
- SAEZ, E. (2001): “Using Elasticities to Derive Optimal Income Tax Rates,” *Review of Economic Studies*, 68, 205–229.
- (2002): “Optimal Income Transfer Programs: Intensive versus Extensive Labor Supply Responses,” *Quarterly Journal of Economics*, 117, 1039–1073.
- SAEZ, E. AND S. STANTCHEVA (2016): “Generalized Social Marginal Welfare Weights for Optimal Tax Theory,” *American Economic Review*, 106, 24–45.
- SHAKED, M. AND J. G. SHANTHIKUMAR (2007): *Stochastic Orders*, Springer Series in Statistics, New York: Springer.
- SNYDER, T. D. AND C. M. HOFFMAN (2001): *Digest of Education Statistics 2000*, National Center for Education Statistics.
- STANTCHEVA, S. (2015): “Optimal Income, Education, and Bequest Taxes in an Intergenerational Model,” *NBER Working Paper No. 21177*.
- (2017): “Optimal Taxation and Human Capital Policies over the Life Cycle,” *Journal of Political Economy*, 125, 1931–1990.
- ZIMMERMAN, S. (2014): “The Returns to College Admission for Academically Marginal Students,” *Journal of Labor Economics*, 32, 711–754.

A Theoretical Appendix

A.1 Derivation of Equation 3

The Lagrangian for the government's problem reads as:

$$\begin{aligned} \mathcal{L} = & \int_{\mathbb{R}_+} \int_{\mathcal{X}} \max\{V^E(X, I), V^H(X, I)\} \tilde{k}(X, I) dX dI \\ & + \rho \left\{ \int_{\mathbb{R}_+} \int_{\mathcal{X}} \mathcal{N}\mathcal{T}_{NPV}^H(X, I) \mathbb{1}_{V_j^E < V_j^H} k(X, I) dX dI \right. \\ & + \int_{\mathbb{R}_+} \int_{\mathcal{X}} \mathcal{N}\mathcal{T}_{NPV}^G(X, I) \mathbb{1}_{V_j^E \geq V_j^H} P^C(X, I, \mathcal{G}(I)) k(X, I) dX dI \\ & \left. + \int_{\mathbb{R}_+} \int_{\mathcal{X}} \mathcal{N}\mathcal{T}_{NPV}^D(X, I) \mathbb{1}_{V_j^E \geq V_j^H} (1 - P^C(X, I, \mathcal{G}(I))) k(X, I) dX dI - \bar{F} \right\}. \end{aligned}$$

The derivative w.r.t. $\mathcal{G}(I)$ is given by:

$$\begin{aligned} \frac{\partial \mathcal{L}}{\partial \mathcal{G}(I)} = & \int_{\mathcal{X}} \mathbb{1}_{V_j^E \geq V_j^H} \frac{\partial V^E(X, I)}{\partial \mathcal{G}(I)} \tilde{h}(X|I) dX \tag{9} \\ & + \rho \int_{\mathcal{X}} \left\{ P^C(X, I, \mathcal{G}(I)) \frac{\partial \mathcal{N}\mathcal{T}_{NPV}^G(X, I)}{\partial \mathcal{G}(I)} + (1 - P^C(X, I, \mathcal{G}(I))) \frac{\partial \mathcal{N}\mathcal{T}_{NPV}^D(X, I)}{\partial \mathcal{G}(I)} \right\} h(X|I) dX \\ & + \rho \int_{\mathcal{X}} \mathbb{1}_{H_j \rightarrow E_j} \left\{ P^C(X, I, \mathcal{G}(I)) \mathcal{N}\mathcal{T}_{NPV}^G(X, I) + (1 - P^C(X, I, \mathcal{G}(I))) \mathcal{N}\mathcal{T}_{NPV}^D(X, I) \right. \\ & \quad \left. - \mathcal{N}\mathcal{T}_{NPV}^H(X, I) \right\} h(X|I) dX \\ & + \rho \int_{\mathcal{X}} \frac{\partial P^C(X, I, \mathcal{G}(I))}{\partial \mathcal{G}(I)} (\mathcal{N}\mathcal{T}_{NPV}^G(X, I) - \mathcal{N}\mathcal{T}_{NPV}^D(X, I)) h(X|I) dX \end{aligned}$$

Recall that $\mathbb{1}_{H_j \rightarrow E_j}$ takes the value one if an individual of type j is pushed over the college enrollment margin due to a small increase in financial aid.

The first term captures the direct utility increase of inframarginal enrollees due to receiving more financial aid. The second term captures the direct fiscal effect of paying more financial aid to inframarginal students. The third term captures the fiscal effect of additional enrollees. The fourth effect captures the fiscal effect due to the increase in the completion rate of inframarginal students. The implied change in the enrollment and dropout rate has no direct first-order effect on welfare: individuals that are marginal in their decision to enroll or not and to continue studying or drop out, were just indifferent between the two respective options, hence this change in behavior has no effect on their utility.

The definitions of $E(I)$ and $\Delta T^E(I)$ directly imply that the third term equals the enrollment effect in (3) multiplied by ρ . The definitions of $\Delta \mathcal{T}^C(I)$, $E(I)$ and $C(I)$ directly imply that the fourth term equals the completion effect in (3) multiplied by ρ .

Now it remains to be shown that the first and second term are equal to the mechanical effect in (3). The application of the envelope theorem implies that the first term reads as

$$\int_{\mathcal{X}} \mathbb{1}_{V_j^E \geq V_j^H} \mathbb{E} \left[\sum_{t=1}^{t_g^{max}} \beta^{t-1} U_c^E(\cdot) \left(1 + \frac{\partial t r_t^E(\cdot)}{\partial \mathcal{G}(I)} \right) \prod_{s=1}^t (\mathbb{1}_{V_s^{ND} \geq V_s^D}) \prod_{s=1}^{t-1} (1 - Pr_s^{Grad}(X)) \right] \tilde{h}(X|I) dX \tilde{E}(I). \quad (10)$$

The second term, using the definitions of $\mathcal{NT}_{NPV}^G(X, I)$ and $\mathcal{NT}_{NPV}^D(X, I)$, can be written as

$$-\rho \int_{\mathcal{X}} \sum_{t=1}^{t_g^{max}} \frac{1}{1+r} \prod_{s=1}^t P_s^E(X, I, \mathcal{G}(I)) h(X|I) dX. \quad (11)$$

Adding (10) and (11), using the definition of the social marginal welfare weight yields equation 3.

A.2 More General Version of Equation 3 with Annual Dropout Decisions

We now show the generalization in which individuals can drop out each period. For this case, we have to distinguish between individuals that drop out in different periods. Hence, for the education decision we have: $e \in \{H, G, D_1, D_2, \dots, D_{t_g^{max}}\}$, where D_t implies that individuals drop out at the beginning of year t . Accordingly we can define the net fiscal contribution of an individual of type (X, I) that drops out in period t by $\mathcal{NT}_{NPV}^{Dt}(X, I)$:

$$\mathcal{NT}_{NPV}^{Dt}(X, I) = \sum_{s=t}^T \left(\frac{1}{1+r} \right)^{s-1} \mathbb{E}(\mathcal{T}(y_s)|X, I, D_t) - \mathcal{G}(I) \sum_{s=1}^{t-1} \left(\frac{1}{1+r} \right)^{s-1}.$$

We also have to define the net fiscal contribution of an individual that is enrolled in year t_g^{max}

$$\mathcal{NT}_{NPV}^{Et_g^{max}}(X, I) = P_{t_g^{max}}(X, I, \mathcal{G}(I)) \mathcal{NT}_{NPV}^G(X, I) + (1 - P_{t_g^{max}}(X, I, \mathcal{G}(I))) \mathcal{NT}_{NPV}^{Dt_g^{max}}(X, I)$$

and for $t = 2, 3, \dots, t_g^{max} - 1$:

$$\mathcal{NT}_{NPV}^{Et}(X, I) = P_t(X, I, \mathcal{G}(I)) \mathcal{NT}_{NPV}^{Et+1}(X, I) + (1 - P_t(X, I, \mathcal{G}(I))) \mathcal{NT}_{NPV}^{Dt}(X, I).$$

The Lagrangian for the government's problem reads as:

$$\begin{aligned} \mathcal{L} = & \int_{\mathbb{R}_+} \int_{\mathcal{X}} \max\{V^E(X, I), V^H(X, I)\} \tilde{k}(X, I) dX dI \\ & + \rho \left\{ \int_{\mathbb{R}_+} \int_{\mathcal{X}} \mathcal{N}\mathcal{T}_{NPV}^H(X, I) \mathbb{1}_{V_j^E < V_j^H} k(X, I) dX dI \right. \\ & + \int_{\mathbb{R}_+} \int_{\mathcal{X}} \mathcal{N}\mathcal{T}_{NPV}^G(X, I) \mathbb{1}_{V_j^E \geq V_j^H} P^C(X, I, \mathcal{G}(I)) k(X, I) dX dI \\ & \left. + \sum_{\tau=1}^{t_g^{max}} \left[\int_{\mathbb{R}_+} \int_{\mathcal{X}} \mathcal{N}\mathcal{T}_{NPV}^{D\tau}(X, I) \mathbb{1}_{V_j^E \geq V_j^H} \prod_{t=1}^{\tau-1} P_t(X, I, \mathcal{G}(I)) (1 - P_\tau(X, I, \mathcal{G}(I))) k(X, I) dX dI \right] - \bar{F} \right\}. \end{aligned}$$

The FOC for $\mathcal{G}(I)$ shares the same basic structure as (9). However, here the fiscal effects due to change in dropout behavior are more involved:⁴⁸

$$\begin{aligned} & \rho \int_{\mathcal{X}} \prod_{t=1}^{t_g^{max}-1} P_t(X, I, \mathcal{G}(I)) \frac{\partial P_{t_g^{max}}(X, I, \mathcal{G}(I))}{\partial \mathcal{G}(I)} \left(\mathcal{N}\mathcal{T}_{NPV}^G(X, I) - \mathcal{N}\mathcal{T}_{NPV}^{Dt_g^{max}}(X, I) \right) h(X|I) dX \\ & + \rho \sum_{\tau=1}^{t_g^{max}-1} \left[\int_{\mathcal{X}} \prod_{t=1}^{\tau-1} P_t(X, I, \mathcal{G}(I)) \frac{\partial P_\tau(X, I, \mathcal{G}(I))}{\partial \mathcal{G}(I)} \left(\mathcal{N}\mathcal{T}_{NPV}^{E\tau+1}(X, I) - \mathcal{N}\mathcal{T}_{NPV}^{D\tau}(X, I) \right) h(X|I) dX \right], \end{aligned}$$

where we let $P_0(X, I, \mathcal{G}(I)) = 1$.

In short term notation, similar to that in (3), we can write

$$\sum_{t=1}^{t_g^{max}} \frac{\partial Con_t(I)}{\partial \mathcal{G}(I)} \Big|_{E_t(I)} \Delta \mathcal{T}^{Con,t}(I) E_t(I)$$

where $Con_t(I)$ is the share of those enrollees with parental income I in period t , that continue studying to year $t + 1$. It is defined by

$$Con_t(I) = \frac{E_{t+1}(I)}{E_t(I)}$$

for $t = 1, 2, \dots, t_g^{max} - 1$ and

$$Con_{t_g^{max}}(I) = \frac{\int_{\mathcal{X}} \mathbb{1}_{V_j^E \geq V_j^H} \prod_{s=1}^{t_g^{max}} P_s(X, I, \mathcal{G}(I)) h(X|I) dX}{E_{t_g^{max}}(I)}.$$

where

$$E_1(I) = E(I) = \int_{\mathcal{X}} \mathbb{1}_{V_j^E \geq V_j^H} h(X|I) dX.$$

and

$$E_t(I) = \int_{\mathcal{X}} \mathbb{1}_{V_j^E \geq V_j^H} \prod_{s=1}^{t-1} P_s(X, I, \mathcal{G}(I)) h(X|I) dX.$$

⁴⁸Again, changes in dropout behaviour have no direct welfare effect due to the envelope theorem.

Finally, the changes in tax revenue are defined by:

$$\Delta \mathcal{T}^{Con,t}(I) = \frac{\int_{\mathcal{X}} \Delta \mathcal{T}^{Con,t}(X, I) \frac{\partial P(X, I, \mathcal{G}(I))}{\partial \mathcal{G}(I)} h(X|I) dX}{\int_{\mathcal{X}} \frac{\partial P(X, I, \mathcal{G}(I))}{\partial \mathcal{G}(I)} h(X|I) dX}$$

where

$$\Delta \mathcal{T}^{Con,t}(X, I) = \mathcal{N}\mathcal{T}_{NPV}^{Et}(X, I) - \mathcal{N}\mathcal{T}_{NPV}^{Dt}(X, I).$$

Hence, the equivalent to equation 3 is given by:

$$\frac{\partial E(I)}{\partial \mathcal{G}(I)} \times \Delta \mathcal{T}^E(I) + \sum_{t=1}^{t_y^{max}} E_t(I) \frac{\partial Con_t(I)}{\partial \mathcal{G}(I)} \Big|_{E_t(I)} \Delta \mathcal{T}^{Con,t}(I) - \tilde{E}(I) (1 - W^E(I)).$$

A.3 Proof of Proposition 1

The government's problem reads as

$$\begin{aligned} \max_{\mathcal{G}(I)} & \int_{\mathbb{R}_+} \int_{\underline{\theta}}^{\tilde{\theta}(I)} U((1 - \tau)y_H) d\tilde{H}(\theta|I) d\tilde{F}(I) \\ & + \int_{\mathbb{R}_+} \int_{\tilde{\theta}(I)}^{\bar{\theta}} U(((1 - \tau)y_H)(1 + \theta) - (\mathcal{F} - \mathcal{G}(I) - tr(I))) d\tilde{H}(\theta|I) d\tilde{F}(I) \\ & + \lambda \left[\int_{\mathbb{R}_+} \int_{\underline{\theta}}^{\tilde{\theta}(I)} \tau y_H dH(\theta|I) dF(I) + \int_{\mathbb{R}_+} \int_{\tilde{\theta}(I)}^{\bar{\theta}} (\tau y_H(1 + \theta) - \mathcal{G}(I)) dH(\theta|I) dF(I) - \bar{F} \right], \end{aligned}$$

where, as in Section 2, \tilde{H} and \tilde{F} denote the distributions of Pareto weights which integrate up to one and \bar{F} is some exogenous revenue requirement. All Pareto weights are non-negative.

The first-order condition for $\mathcal{G}(I)$ is given by:

$$h(\tilde{\theta}(I)|I) \Big| \frac{\partial \tilde{\theta}(I)}{\partial \mathcal{G}(I)} \Big| \left(\tau y_h \tilde{\theta}(I) - \mathcal{G}(I) \right) - \left(1 - H(\tilde{\theta}(I)|I) \right) (1 - W^E(I)) = 0, \quad (12)$$

where $W^E(I)$ is average social marginal welfare weight for enrollees with parental income I and formally given by:

$$W^E(I) = \frac{\int_{\tilde{\theta}}^{\bar{\theta}} U'(((1 - \tau)y_H)(1 + \theta) - (\mathcal{F} - \mathcal{G}(I) - tr(I))) d\tilde{H}(\theta|I) \tilde{f}(I)}{\lambda(1 - H(\theta|I))f(\theta)}.$$

Using

$$\frac{\partial \tilde{\theta}(I)}{\partial \mathcal{G}(I)} = -\frac{1}{y_H(1 - \tau)}$$

and inserting into (12) gives the first-order condition explained in the main text just before Proposition 1. Solving for $\mathcal{G}(I)$ gives Proposition 1.

Unweighted Utilitarianism with $U(\mathbf{x})=\mathbf{x}$. If we assume $\tilde{h}(\theta|I) = h(\theta|I)$ for all (θ, I) and further assume $U(x) = x$, then, we obtain

$$W^E(I) = \frac{1}{\lambda} \forall I,$$

hence the welfare weights are the same for all parental income groups. The value for λ is measures the marginal value of public funds and therefore depends on \bar{F} . To get an understanding of the optimal value for λ , consider the perturbation around the optimum, where $\mathcal{G}(I)$ is increased by $\delta \rightarrow 0$ for all I . This basically implies changing the lump sum component of $\mathcal{G}(I)$ and is equivalent to just integrating over (12). The impact on welfare is given by:

$$\int_I (1 - W^E(I)) \left(1 - H(\tilde{\theta}(I)|I)\right) dF(I) + \int_I h(\tilde{\theta}(I)|I) \left| \frac{\partial \tilde{\theta}(I)}{\partial \mathcal{G}(I)} \right| \left(\tau y_h \tilde{\theta}(I) - \mathcal{G}(I) \right) dF(I) = 0.$$

and hence for $W^E(I) = \frac{1}{\lambda} \forall I$:

$$\int_I \left(1 - \frac{1}{\lambda}\right) \left(1 - H(\tilde{\theta}(I)|I)\right) dF(I) + \int_I h(\tilde{\theta}(I)|I) \left| \frac{\partial \tilde{\theta}(I)}{\partial \mathcal{G}(I)} \right| \left(\tau y_h \tilde{\theta}(I) - \mathcal{G}(I) \right) dF(I) = 0.$$

Here we see that $\lambda = 1$ would be consistent with $\mathcal{G}(I) = \tau y_h \tilde{\theta}(I)$ for all I . Recall that the government budget constraint is given by:

$$\int_{\mathbb{R}_+} \int_{\underline{\theta}}^{\tilde{\theta}(I)} \tau y_H dH(\theta|I) dF(I) + \int_{\mathbb{R}_+} \int_{\tilde{\theta}(I)}^{\bar{\theta}} (\tau y_H (1 + \theta) - \mathcal{G}(I)) dH(\theta|I) dF(I) - \bar{F} = 0.$$

If the exogenous revenue requirement \bar{F} is such that the budget constraint holds for $\mathcal{G}(I) = \tau y_h \tilde{\theta}(I)$ for all I , then we obtain $\lambda = 1$ and the formula in Proposition 1 becomes

$$\mathcal{G}(I) = \tau (\mathcal{F} - tr(I)). \tag{13}$$

Assume that instead the budget constraint would be violated and this level of financial aid can not be financed. Then we have $\lambda > 1$ and hence $W^E(I) < 1 \forall I$. Generally, we could also have the case where $\lambda < 1$. E.g. assume that $\bar{F} = -\infty$. In this case, there would be infinitely many public funds available for financial aid and therefore the marginal value of public funds would be zero. But of course this is only of theoretical interest.

The fact that the marginal value of public funds is not equal to unity even though preferences are linear may seem in contrast to the optimal income tax literature, where it is a

standard result that the marginal value of public funds is equal to one for quasi-linear preferences and in other words the average welfare weights is equal to one, see e.g. Saez (2002). The reason is that the policy instruments that we consider are such that there is no lump sum element. While the financial aid schedule $\mathcal{G}(I)$ of course has an intercept $\mathcal{G}(0)$ that can optimally be chosen, this is no lump sum transfer in the classical sense because it only reaches college students and not individuals who forgo college. Therefore, varying this lump sum component also has incentive effects on the college decision and one cannot just pay out a dollar to everyone without affecting behavior.

A.4 Proof of Corollary 1

Differentiating (5) w.r.t. I yields:

$$\mathcal{G}'(I) = -\tau tr'(I) + (1 - \tau) \frac{\partial \left(\frac{1-H(\tilde{\theta}(I))}{h(\tilde{\theta}(I))} \right)}{\partial \tilde{\theta}(I)} (tr'(I) + \mathcal{G}'(I)) (1 - W^E)$$

where we used $\tilde{\theta}(I) = \frac{F-tr(I)-\mathcal{G}(I)}{(1-\tau)y_H}$. and therefore $\tilde{\theta}'(I) = \frac{-tr'(I)-\mathcal{G}'(I)}{(1-\tau)y_H}$. Solving for $\mathcal{G}'(I)$ we get

$$\mathcal{G}'(I) = \frac{-\tau tr'(I) + (1 - \tau) \frac{\partial \left(\frac{1-H(\tilde{\theta}(I))}{h(\tilde{\theta}(I))} \right)}{\partial \tilde{\theta}(I)} tr'(I) (1 - W^E)}{1 - (1 - \tau) \frac{\partial \left(\frac{1-H(\tilde{\theta}(I))}{h(\tilde{\theta}(I))} \right)}{\partial \tilde{\theta}(I)} (1 - W^E)}$$

which proves Corollary 1 since by assumption $tr'(I) > 0$ and log concavity of the skill distribution implies $\frac{\partial \left(\frac{1-H(\tilde{\theta}(I))}{h(\tilde{\theta}(I))} \right)}{\partial \tilde{\theta}(I)} < 0$.

A.5 Proof of Corollary 2

Differentiating (5) w.r.t. I yields:

$$\mathcal{G}'(I) = -\tau tr'(I) + (1 - W^E) (1 - \tau) \left[\frac{\partial \left(\frac{1-H(\tilde{\theta}(I)|I)}{h(\tilde{\theta}(I)|I)} \right)}{\partial \tilde{\theta}(I)} (tr'(I) + \mathcal{G}'(I)) - y_H (1 - \tau) \frac{\partial \left(\frac{1-H(\theta|I)}{h(\theta|I)} \right)}{\partial I} \Big|_{\theta=\tilde{\theta}(I)} \right]$$

Hence we obtain

$$\mathcal{G}'(I) = \frac{-\tau tr'(I) + (1 - W^E) (1 - \tau) \left[\frac{\partial \left(\frac{1-H(\tilde{\theta}(I))}{h(\tilde{\theta}(I))} \right)}{\partial \tilde{\theta}(I)} tr'(I) - y_H (1 - \tau) \frac{\partial \left(\frac{1-H(\theta|I)}{h(\theta|I)} \right)}{\partial I} \Big|_{\theta=\tilde{\theta}(I)} \right]}{1 - (1 - \tau) \frac{\partial \left(\frac{1-H(\tilde{\theta}(I))}{h(\tilde{\theta}(I))} \right)}{\partial \tilde{\theta}(I)} (1 - W^E)}$$

which proves Corollary 2 since by assumption $tr'(I) > 0$, log concavity of the skill distribution implies $\frac{\partial \left(\frac{1-H(\hat{\theta}(I))}{h(\hat{\theta}(I))} \right)}{\partial \hat{\theta}(I)} < 0$ and we assumed

$$\frac{\partial \left(\frac{1-H(\theta/I)}{h(\theta/I)} \right)}{\partial I} > 0 \quad \forall \theta, I.$$

A.6 Optimal Income Taxation

The planner's problem is the same as in (1) with the difference that the planner also optimally chooses the income tax schedule $T(\cdot)$. Notice that the formula for optimal financial aid policies is unaltered. We allow the tax function $T(\cdot)$ to be arbitrarily nonlinear in the spirit of Mirrlees (1971). We restrict the tax function to be only a function of income and to be independent of the education decision. This tax problem can either be tackled with a variational or tax perturbation approach (Saez, 2001; Golosov et al., 2014; Jacquet and Lehmann, 2016) or with a restricted mechanism design approach for nonlinear history-independent income taxes that we explore in Findeisen and Sachs (2017).

We here provide a heuristic version of the former approach within our model. For notational convenience, we consider the model of Section 2 with the assumption that individuals can only dropout at the beginning of period 3. We also assume that agents can only graduate in 4 years.

Consider an increase of the marginal tax by an infinitesimal amount $d\mathcal{T}'$ in an income interval of infinitesimal length $[y^*, y^* + dy]$. As a consequence of this reform, all individuals with $y > y^*$ face an increase of the absolute tax level of $d\mathcal{T}'dy$. The tax reform therefore induces a mechanical increase in welfare of

$$\begin{aligned} \Delta W_{MR}(y^*) &= \rho d\mathcal{T}' dy \sum_{t=1}^T \left(\frac{1}{1+r} \right)^{t-1} \int_{y^*}^{\infty} h_{t,H}(y) dy \times s_H \\ &+ \rho d\mathcal{T}' dy \sum_{t=3}^T \left(\frac{1}{1+r} \right)^{t-1} \int_{y^*}^{\infty} h_{t,D}(y) dy \times s_D \\ &+ \rho d\mathcal{T}' dy \sum_{t=5}^T \left(\frac{1}{1+r} \right)^{t-1} \int_{y^*}^{\infty} h_{t,G}(y) dy \times s_G \end{aligned}$$

through the tax revenue increase (in net present value). $h_{t,e}(y)$ is the density of income of individuals with education level e in period t and s_e is the overall share of individuals with education level e . Both, the income densities and the education shares are endogeneous w.r.t. to taxes, we get to this below.

Note that this increase in tax payment also has mechanical effects on individual utilities which adds up to the following welfare effect

$$\begin{aligned}\Delta W_{MU}(y^*) &= d\mathcal{T}' dy \sum_{t=1}^T \left(\frac{1}{1+r}\right)^{t-1} \int_{y^*}^{\infty} E(U_c|y_t = y) h_{t,H}(y) dy \times s_H \\ &+ d\mathcal{T}' dy \sum_{t=3}^T \left(\frac{1}{1+r}\right)^{t-1} \int_{y^*}^{\infty} E(U_c|y_t = y) h_{t,D}(y) dy \times s_D \\ &+ d\mathcal{T}' dy \sum_{t=5}^T \left(\frac{1}{1+r}\right)^{t-1} \int_{y^*}^{\infty} E(U_c|y_t = y) h_{t,G}(y) dy \times s_G.\end{aligned}$$

Now we turn to the endogeneity of education shares. First of all some individuals will change their initial enrollment decision. We define $\mathbb{1}_{H_j \rightarrow E_j}^{y^*}$ to take the value one if an individual of type j is marginal in the enrollment decision w.r.t. to a one dollar tax increase for earnings above y^* . Then, the welfare effect of individuals changing their enrollment decision due to a small increase in $T'(y^*)$ is given by:

$$\Delta W_E(y^*) = \rho d\mathcal{T}' dy \int_{\mathbb{R}_+} \int_{\mathcal{X}} \mathbb{1}_{H_j \rightarrow E_j}^{y^*} \left\{ \mathcal{N}\mathcal{T}_{NPV}^E(X, I) - \mathcal{N}\mathcal{T}_{NPV}^H(X, I) \right\} h(X|I) dX dI.$$

Similarly, the probability to continue college and not drop out is endogenous w.r.t. taxes, i.e. we have $P(X, I, \mathcal{G}(I), T(\cdot))$. The change in welfare due to the change in dropout behavior, with some abuse of notation, is simply given by:

$$\Delta W_D(y^*) = \rho \int_{\mathbb{R}_+} \int_{\mathcal{X}} \int_{y^*}^{\infty} \frac{\partial P(X, I, \mathcal{G}(I), T(\cdot))}{\partial \mathcal{T}(y)} dy \left\{ \mathcal{N}\mathcal{T}_{NPV}^G(X, I) - \mathcal{N}\mathcal{T}_{NPV}^D(X, I) \right\} h(X|I) dX dI.$$

Finally, an increase in the marginal tax rate also affects labor supply behavior for individuals within the interval $[y^*, y^* + dy]$. Individuals within this infinitesimal interval change their labor supply by

$$\frac{\partial y_t^*}{\partial \mathcal{T}'} d\mathcal{T}' = -\varepsilon_{y_t^*, 1-\mathcal{T}'} \frac{y_t^*}{1-\mathcal{T}'} d\mathcal{T}'.$$

Whereas this change in labor supply has no first-order effect on welfare via individual utilities by the envelope theorem, it has an effect on tax revenue, which is given by:

$$\begin{aligned}\Delta W_L(y^*) &= \frac{\mathcal{T}'(y^*)}{1-\mathcal{T}'(y^*)} y^* \varepsilon_{y, 1-\mathcal{T}'} \times d\mathcal{T}' \times \\ &\left(s_H \sum_{t=1}^T \left(\frac{1}{1+r}\right)^{t-1} h_{t,H}(y^*) + s_D \sum_{t=3}^T \left(\frac{1}{1+r}\right)^{t-1} h_{t,D}(y^*) + s_G \sum_{t=5}^T \left(\frac{1}{1+r}\right)^{t-1} h_{t,G}(y^*) \right).\end{aligned}$$

Since this reform must not have any non-zero effect on welfare if the tax system is optimal, we have to have

$$\Delta W_{MR}(y^*) + \Delta W_{MU}(y^*) + \Delta W_E(y^*) + \Delta W_D(y^*) + \Delta W_L(y^*) = 0 \quad (14)$$

which provides an implicit characterization of $\mathcal{T}'(y^*)$.

Finally, the optimal level for the lump-sum element of the tax schedule $T(0)$ is implicitly characterized by

$$\Delta W_{MR}(0) + \Delta W_{MU}(0) + \Delta W_E(0) + \Delta W_D(0) = 0.$$

This optimal tax approach is related to the formulas of Saez (2002) and Jacquet et al. (2013).

To implement this formula numerically, we follow a guess and verify approach. Hence, we start with a guess for the tax schedule and then evaluate (14).⁴⁹ We then slightly adjust $T'(y^*)$ to make (14) closer to zero (but keep $\Delta W_{MR}(y^*) + \Delta W_{MU}(y^*) + \Delta W_E(y^*) + \Delta W_D(y^*)$ fixed, i.e. we only adjust $\Delta W_L(y^*)$). We then calculated the new allocation for this adjusted schedule and evaluate (14) for income levels again and so on. We proceed until convergence.⁵⁰

B Estimation and Calibration

B.1 Current Tax Policies and Tuition

To capture current tax policies, we use the approximation of Heathcote et al. (2017), which has been shown to work well in replicating the US tax code. Since this specification does not contain a lump-sum element, we slightly adjust this schedule. We set the lump sum element of the tax code $T(0)$ to minus \$1,800 a year. For average incomes this fits the deduction in the US-tax code quite well.⁵¹ For low incomes this reflects that individuals might receive transfers such as food stamps.⁵²

For tuition costs, we take average values for the year 2000 from Snyder and Hoffman (2001) for the regions Northeast, North Central, South, and West, as they are defined in the NLSY. We also take into account the amount of money that is spent per student by public appropriations, which has to be taken into account for the fiscal externality. The average values are \$7,434 for annual tuition and \$4,157 for annual public appropriations per student. Besides these implicit subsidies, students receive explicit subsidies in the form of grants and tuition waivers. We estimate how this grant receipt varies with parental income and ability

⁴⁹In fact a more complicated version of (14) which accounts for dropout behavior in every period and also accounts for stochastic graduation.

⁵⁰In each iteration, we also optimally choose the financial aid schedule $\mathcal{G}(I)$ given the tax schedule in the respective iteration.

⁵¹Guner et al. (2014) report a standard deduction of \$7,350 for couples that file jointly. For an average tax rate of 25% this deduction could be interpreted as a lump sum transfer of slightly more than \$1,800.

⁵²The average amount of food stamps per eligible person was \$72 per month in the year 2000. Assuming a two person household gives roughly \$1,800 per year. Source: <http://www.fns.usda.gov/sites/default/files/pd/SNAPsummary.pdf>

in Appendix B.2 using information provided in the NLSY97. We find a strong negative effect of parental income on financial aid receipt. Additionally, we can capture merit-based grants by the conditional correlation of AFQT scores with grant receipt. Finally, we calibrate the exogenous budget element \bar{F} in the following way. For the current U.S. polices, we calculate the present value of financial aid spending and the present value of tax revenues collected from the cohorts that we consider (born between 1980 and 1984 from the NLSY97) and obtain \bar{F} from the difference between the two.

We categorize the following 4 regions:

- Northeast: CT, ME, MA, NH, NJ, NY, PA, RI, VT
- North Central: IL, IN, IA, KS, MI, MN, MO, NE, OH, ND, SD, WI
- South: AL, AR, DE, DC, FL, GA, KY, LA, MD, MS, NC, OK, SC, TN , TX, VA, WV
- West: AK, AZ, CA, CO, HI, ID, MT, NV, NM, OR, UT, WA, WY

We base the following calculations on numbers presented by Snyder and Hoffman (2001). Table 313 of this report contains average tuition fees for four-year public and private universities. According to Table 173, 65% of all four-year college students went to public institutions, whereas 35% went to private institutions. For each state we can therefore calculate the average (weighted by the enrollment shares) tuition fee for a four-year college. We then use these numbers to calculate the average for each of the four regions, where we weigh the different states by their population size. We then arrive at numbers for yearly tuition & fees of \$9,435 (North East), \$7,646 (North Central), \$6,414 (South) and \$7,073 (West). For all individuals in the data with missing information about their state of residence, we chose a country wide population size weighted average of \$7,434.

Tuition revenue of colleges typically only covers a certain share of their expenditure. Figures 18 and 19 in Snyder and Hoffman (2001) illustrate by which sources public and private colleges finance cover their costs. Unfortunately no distinction between two and four-year colleges is available. From Figures 18 and 19 we then infer how many dollars of public appropriations are spent for each dollar of tuition. Many of these public appropriations are also used to finance graduate students. It is unlikely that the marginal public appropriation for a bachelor student therefore equals the average public appropriation at a college given that costs for graduate students are higher. To solve this issue, we focus on institutions “that primarily focus on undergraduate education” as defined in Table 345. Lastly, to avoid double counting of grants and fee waivers, we exclude them from the calculation as we directly use the detailed individual data about financial aid receipt from the NLSY (see Section B.2). Based on these calculations we arrive at marginal public appropriations of \$5,485 (Northeast), \$4,514 (North Central), \$3,558 (South), \$3,604 (West) and \$4,157 (No information about region).

B.2 Estimation of Grant Receipt

Grants and tuition subsidies are provided by a variety of different institutions. Pell grants, for example, are provided by the federal government. In addition, there exist various state and university programs. To make progress, similar to Johnson (2013) and others, we go on to estimate grant receipt directly from the data.

Next, we estimate the amount of grants conditional on receiving grants as a Tobit model:

$$gr_i = \alpha^{gr} + f(I_i) + \beta_4^{gr} AFQT_i + \beta_5^{gr} depkids_i + \varepsilon_i^{gr}. \quad (15)$$

where $f(I_i)$ is a spline function of parental income and ε_i^{gr} represents measurement error. Besides grant generosity being need-based (convexly decreasing), generosity is also merit-based as $\hat{\beta}_4^{gr} > 0$ and increases with the number of other dependent children (besides the considered student) in the family.

Table 3: OLS for Grants

	AFQT	Dependent Children
Coefficient	39.40***	321.75**
Standard Error	(5.03)	(106.39)

N=968. * p ≤ 0.10, ** p ≤ 0.05, *** p ≤ 0.01.

B.3 Wage Estimation

We specify and estimate wage life-cycle paths as follows. Our procedure first estimates labor earnings life-cycle profiles and then calibrates the respective wage profiles based on those estimates in a second step. Specifically, we use the following functional form for earnings y :

$$\forall e = H, G : \log y_{it}^e = \beta_0^e + \beta_\theta^e \log \theta_i + \beta_{t1}^e t + \beta_{t2}^e t^2 + \beta_{t3}^e t^3 + v_i^{e*}. \quad (16)$$

We estimate separate parameters for high school graduates and college graduates.⁵³ The parameter β_θ^e captures different returns to ability for agents of a given education level. The extent to which the college wage premium is increasing in ability is determined by the ratio $\frac{\beta_\theta^G}{\beta_\theta^H}$. We find a ratio larger than 1, which implies a complementary relationship between initial ability and education. Our estimates can be found in Table 4. v_i^{e*} is a random effect that captures persistent differences in wages conditional on the agent's schooling choice. We assume that agents do not know the value of v_i^{e*} at the beginning of the model, but that its value is revealed as soon as the agents finish their education and enter the labor market. Uncertainty

⁵³Dropouts have the same wage parameters as high school graduates except for the constant term. This gives us a very good fit for the relative earnings of dropouts, consistent with the evidence in Lee et al. (2017).

over v_i^{e*} creates uncertainty over an agent's returns to college. After v_i^{e*} there is no further uncertainty about an agent's wage path.

The age earnings coefficients β_{t1}^e , β_{t2}^e and β_{t3}^e are education dependent but independent from gender. However, since we assume different labor supply elasticities for men and women, the implied wage life-cycle profiles will differ across gender because how a given earnings path maps into wages depends on the labor supply elasticity. The age coefficients are estimated from the NLSY79 since individuals from the NLSY97 are only observed until their mid-30s. In sum, this procedure pins down a stochastic distribution of potential life-cycle wage paths for each individual, which depend on gender, ability, and the education decisions. We demonstrate in Section 4.3 that we obtain life-cycle paths of earnings and wages which are consistent with the data.⁵⁴

We estimate the age coefficients β_{t1}^e , β_{t2}^e , β_{t3}^e using panel data from the NLSY79 since individuals in the NLSY97 are too young (born between 1980 and 1984) such that we can infer how wages evolve once individuals are older than 35. In the second step, we build the transformed variable $\widetilde{\log y_{it}^e} = \log y_{it}^e - \beta_{t1}^e t - \beta_{t2}^e t^2 - \beta_{t3}^e t^3$, which takes out age effects from yearly log incomes. Using the NLSY97, we estimate the relationship of log income with gender and log AFQT, estimating separate models and coefficients by education level. We use a random-effects estimator and assume normality, yielding education specific variances for v_i^e . The estimates are displayed in Table 4. There is a significant college premium in the model, although the high-school constant is larger, because we have used education dependent age profiles.

College Educated				
	Female	Log AFQT	Education Constant	Variance v_i
Coefficient	-0.14***	0.47***	3.06***	0.42
Standard Error	(0.02)	(0.07)	(0.35)	
High-School Educated				
	Female	Log AFQT	Education Constant	Variance v_i
Coefficient	-0.25***	0.31***	7.11***	0.36
Standard Error	(0.01)	(0.03)	(0.35)	

Table 4: Regressions: Income

Notes: Random effect models, estimated with NLSY9. Dependent variable is log yearly income, cleaned for age effects. Age effects are obtained by estimating a cubic polynomial on the NLSY79. These age coefficients are available upon request. N=10,165 (College) and N=19,955 (High-School). * $p \leq 0.10$, ** $p \leq 0.05$, *** $p \leq 0.01$.

⁵⁴We use these same parameter estimates to calculate life-cycle earnings for parents. We choose the idiosyncratic component of earnings, v_i^{e*} , to generate earnings at age 45 equal to the parental earnings levels we observe in the data.

Next, we explain how to go from the estimated income to the wage profiles. The reason why we do not estimate wage profiles directly is that we append Pareto tails to the income distribution on which more reliable information is available. Top incomes are underrepresented in the NLSY as in most survey data sets. Following common practice in the optimal tax literature (Piketty and Saez, 2013), we therefore append Pareto tails to each income distribution, starting at incomes of \$150,000. We set the shape parameter α of the Pareto distribution to 1.5 for all income distributions.

Next we describe the mapping from y to w as in Saez (2001). Given the utility function we assume with no income effects, in each year individuals solve a static labor supply problem where optimal labor supply in that year only depends on the current wage (which evolves over the life-cycle) and marginal tax distortions. It is easy to show that the first-order condition for an individual facing a marginal tax rate schedule is

$$\ln w = \frac{\epsilon + \tau}{1 + \epsilon} \ln y - \frac{1}{1 + \epsilon} \ln(\lambda(1 - \tau)),$$

if the tax function is of the form $T(y) = y - \rho y^{1-\tau}$. Using the estimates from the regression model, we can express the wage for a given type (age, gender, ability, education) as at age t :

$$\ln w_{it} = \frac{\epsilon + \tau}{1 + \epsilon} \left(\hat{\beta}_0^{es} + \hat{\beta}_\theta^e \log \theta_i + \hat{\beta}_t^e t + \hat{\beta}_{t2}^e t^2 + \hat{\beta}_{t3}^e t^3 + v_i^{e*} \right) - \frac{1}{1 + \epsilon} \ln(\lambda(1 - \tau)).$$

B.4 Value Functions During College

Agents do not graduate and remain in college with probability $(1 - Pr_t^{Grad}(\theta))$, which depends on the agent's ability level θ . Further, we allow the interest rate the agent receives in college to vary by the agent's assets (positive or negative) and by the agent's parental income, to reflect features of the Stafford loan program.

$$\begin{aligned} V_t^{E, \ell_t^E} \left(X, I, a_t, \varepsilon_t^{\ell_t^E} \right) = \max_{c_t} [U^E \left(c_t, \ell_t^E, X, \varepsilon_t^{\ell_t^E} \right) + \\ \beta \{ (1 - Pr_t^{Grad}(\theta)) \mathbb{E} [V_{t+1}^E(X, I, a_{t+1}, \varepsilon_{t+t})] + Pr_t^{Grad}(\theta) \mathbb{E} [V_{t+1}^W(X, e = G, a_{t+1}, w_{t+1})] \} \end{aligned} \quad (17)$$

subject to

$$c_t = \ell_t^E \underline{\omega} + a_t (1 + r(a_t, I)) - a_{t+1} - \mathcal{F}_{\text{Region}} + \mathcal{G}(X, I)$$

and

$$a_{t+1} \geq \bar{a}_{t+1}^E,$$

where $V_{t+1}^E(X, I, a_{t+1}, \varepsilon_{t+t})$ and ε_{t+t} are defined in the main body. The term $[V_{t+1}^W(X, e = G, a_{t+1}, w_{t+1})]$ is the expected value of being a college graduate in the workforce in year $t + 1$, where the expectation is taken over the permanent skill shock. We allow tuition, $\mathcal{F}_{\text{Region}}$, to depend on

the agent's region. This allows the model to capture differences in tuition across geographic regions and is also helpful for identifying the parameters of the model.

B.5 Details: Parent's Problem

The parent's problem begins when the parent turns 20 years old. Each year the parent receives income and makes consumption/saving decisions. We assume that all parents make transfers to their children at the year which corresponds to $t = 1$ for the child and an age of 43 for the parent.⁵⁵ Parents start the model with 0 assets and live until age 65.

For all years when the transfer is not given, the parent simply chooses how much to consume and save. Let V_t^P denote the parent's value function in year t . We can write this as

$$V_t^P(\tilde{X}, I, a_t^P) = \max_c \left[\frac{c^{1-\gamma}}{1-\gamma} + \beta V_{t+1}^P(\tilde{X}, I, a_{t+1}^P) \right],$$

subject to:

$$c = y_{it}^P + (1+r)a_{it}^P - a_{it+1}^P$$

where a_t^P is the parent's assets in year t and y_{it}^P is the parent's income in year t .⁵⁶ Note that a parent's state space does not include the child's idiosyncratic preference for college ε^E .

In the year of the transfer, the parent also receives utility from transfers. In this year, we write the parent's Bellman equation as

$$V_t^P(\tilde{X}, I, a_t^P) = \max_{c, tr^{hs}, tr^{col}} \left[\frac{c^{1-\gamma}}{1-\gamma} + F(tr^{hs}, tr^{col}, \tilde{X}, I) + \beta \mathbb{E} \left[V_{t+1}^P(\tilde{X}, I, a_{t+1}^P - tr^e) \right] \right]$$

subject to:

$$c = y_{it}^P + (1+r)a_{it}^P - a_{it+1}^P.$$

where tr^{hs} and tr^{col} are the transfers offered conditional on the child's education choice, and tr^e are the realized transfers.⁵⁷ As the parent must commit to transfers before the child's college preference shock is realized, the child's college choice and therefore the value of tr^e is stochastic at the time the parent chooses the transfer in the eyes of the parent. $F(tr^{hs}, tr^{col}, \tilde{X}, I)$ is

⁵⁵This will correspond to age 18 of the child if the parent gave birth to the child at age 25. This is the median age a mother gave birth to their child in the NLSY97.

⁵⁶We set the risk aversion for parents $\gamma = 1$ outside of estimation such that the estimate of the child's γ is identified only by decisions of the child, and is not identified by the amount of parental transfers given.

⁵⁷In the data, we follow Johnson (2013) we calculate transfers as the sum of monetary transfers and the monetary benefit of living at home. We assume that the monetary benefit of living at home is given exogenously and only the actual monetary transfers are included in the parent's budget constraint. We assume that the monetary benefit of living at home is equal to the average amount conditional on parental income and the child's education choice.

the expected utility the parent receives from the transfer schedule tr^{hs} , tr^{col} and is defined in the main text.

We assume parents must also pay transfers to the agent's siblings. Therefore, if the child has $nsibs$ siblings, the child's parents also pay $nsibs \times \bar{tr}(nsibs, I)$ out of their budget to the other siblings, where $\bar{tr}(nsibs, I)$ is the predicted level of parental monetary transfers for children with $nsibs$ and parental income of I , unconditional of the child's education choice. We predict $\bar{tr}(nsibs, I)$ by regressing monetary transfers on parental income separately for each number of siblings we observe in the data.

Parent's Earnings Profile Calibration We assume that parental earnings are determined by a similar process to the child's earnings. Specifically, parental earnings are given by

$$\forall e = H, G : \log y_t^P = \beta_{t1}^{ParEdu} \text{ParAge}_t + \beta_{t2}^{ParEdu} \text{ParAge}_t^2 + \beta_{t3}^{ParEdu} \text{ParAge}_t^3 + v^P.$$

where ParAge_t is the parent's age in period t . The age coefficients, β_{t1}^{ParEdu} , β_{t2}^{ParEdu} , and β_{t3}^{ParEdu} are taken from the child's earnings regression. We assume that the parent's age coefficients are given by the college age coefficients if at least one parent has attended college, otherwise the parent's age coefficients are given by the age coefficients for a child that has not attended college.

The term v^P represents persistent, idiosyncratic differences in earnings across parents. We assume that we observe the parental income variable I when parents are 40 years old. Therefore, we must have $y_{40} = I$ for each parent we observe in data. We therefore choose v^P such that the predicted parental income at age 40 is equal to the observed parental income variable I . We can write this as

$$v^P = \log I - (\beta_{t1}^{ParEdu} \text{ParAge}_t + \beta_{t2}^{ParEdu} \text{ParAge}_t^2 + \beta_{t3}^{ParEdu} \text{ParAge}_t^3).$$

B.6 Likelihood Function

Assume that the econometrician observes transfers $tr_i^{e,o}$, which differ from true transfers, tr_i^{e*} , by an error term e^{tr} . Further, we assume this error term is normally distributed: $e^{tr} \sim N(0, \sigma^{e^{tr}})$. We suppress all dependencies for notational convenience. Then, given parameters Γ , the likelihood contribution of an agent who graduates from college after T_i^E years, has a sequence of work in college decisions of $\{\ell_{it}^E\}_{t=1}^{T_i^E}$, and has observed college transfers $tr_i^{E,o}$ is⁵⁸

⁵⁸The probability of this event in fact also depends on the graduation probabilities Pr_t^{Grad} . But these are just constant factors in the likelihood, which is why refrain from putting them here.

$$\begin{aligned} \mathcal{L}_i \left(e_i = G, tr_i^{E,o}, \{\ell_{it}^E\}_{t=1}^{T_i^E} | \Gamma \right) = \\ Pr(E) f^N \left(\frac{tr_i^{E*} - tr_i^{E,o}}{\sigma^{etr}} \right) \frac{1}{\sigma^{etr}} \left(\prod_{t=1}^{T_i^E} Pr(\ell_{it}^E) \right), \end{aligned} \quad (18)$$

where f^N is the standard normal PDF, and where the probability of initially enrolling in college, $Pr(E)$, and the choice probability of not dropping out and working ℓ_{it}^E in college, $Pr(\ell_{it}^E)$, are given by the extreme-value choice probabilities as

$$Pr(E) = \frac{\exp(\tilde{V}^E/\sigma^E)}{\exp(\tilde{V}^E/\sigma^E) + \exp(\tilde{V}^H/\sigma^E)}$$

and

$$Pr(\ell_{it}^E) = \frac{\exp(\tilde{V}_t^{E,\ell_{it}^E}/(\sigma^{\ell^E} \lambda)) \left(\sum_{\ell \in \{0,PT,FT\}} \exp(\tilde{V}_t^{E,\ell}/(\sigma^{\ell^E} \lambda)) \right)^{\lambda-1}}{\left(\exp(\tilde{V}^D - \delta/\sigma^{\ell^E}) \right) + \left(\sum_{\ell \in \{0,PT,FT\}} \exp(\tilde{V}_t^{E,\ell}/(\sigma^{\ell^E} \lambda)) \right)^{\lambda}},$$

where σ^E and σ^{ℓ^E} are parameters governing the variance of the enrollment shock and college working shock, respectively, and λ is a nesting parameter and where value functions with tildes represent the value function minus the idiosyncratic preference draws.

The likelihood contribution of an agent who drops out in year T_i^D , has a sequence of work in college decisions of $\{\ell_{it}^E\}_{t=1}^{T_i^{dropout}-1}$, and has observed college transfers $tr_i^{E,o}$ is

$$\begin{aligned} \mathcal{L}_i \left(e_i = D, tr_i^{E,o}, \{\ell_{it}^E\}_{t=1}^{T_i^D-1} | \Gamma \right) = \\ Pr(E) f^N \left(\frac{tr_i^{E*} - tr_i^{E,o}}{\sigma^{etr}} \right) \frac{1}{\sigma^{etr}} \left(\prod_{t=1}^{T_i^D-1} Pr(\ell_{it}^E) \right) Pr(D_{T^D}), \end{aligned} \quad (19)$$

where the probability of dropping out, $Pr(D_{T^D})$, is given by the extreme value choice probabilities as

$$Pr(D_{T^D}) = \frac{\exp(\tilde{V}_{T^D}^D/\sigma^{\ell^E})}{\left(\exp(\tilde{V}_{T^D}^D/\sigma^{\ell^E}) \right) + \left(\sum_{\ell \in \{0,PT,FT\}} \exp(\tilde{V}_{T^D}^{E,\ell}/(\sigma^{\ell^E} \lambda)) \right)^{\lambda}}.$$

The likelihood function of an agent who enters the labor force directly and is observed with transfers $tr_i^{H,o}$ is given by

$$\mathcal{L}_i \left(e_i = H, tr_i^{H,o} | \Gamma \right) = (1 - Pr(E)) f^N \left(\frac{tr_i^{H*} - tr_i^{H,o}}{\sigma^{etr}} \right) \frac{1}{\sigma^{etr}}. \quad (20)$$

We therefore choose the parameters Γ to maximize the log likelihood:

$$\max_{\Gamma} \sum_i \log \mathcal{L}_i(\cdot | \Gamma).$$

B.7 Identification

B.7.1 Identification of Main Parameters

The parameter γ and the parameter governing the variance of the college-enrollment preference shock, σ^E , play crucial roles in our analysis as they determine the extent to which increasing financial aid affects the college enrollment decision. Higher values of σ^E and lower values of γ imply a smaller elasticity of enrollment with respect to increases in financial aid. These parameters are jointly identified by the relationship between enrollment and parental income of otherwise similar individuals. Enrolling in college will generally imply lower net income while enrolled in college and higher income later in life. To the extent that borrowing constraints are effective and that parental transfers are increasing in parental income, children from lower-income backgrounds will not be able to smooth consumption and therefore will have lower consumption in their early life. The parameter γ determines the cost of not being able to smooth consumption early in life. A high value of γ therefore implies low college enrollment for individuals close to the borrowing constraint.

Furthermore, exclusion restrictions in the grant function help us to identify the elasticity of college going with respect to financial aid. Tuition varies by region but region does not enter the earnings function or utility function. Therefore, similar to Heckman et al. (1998), variation in tuition levels creates variation in the value of college enrollment which helps us to identify γ and σ^E . In Section B.7.2 we use a simple example to build intuition on how these two parameters are separately identified.

Additionally, the extent to which poor students are more likely to work than rich students will be governed by γ ; this tells us how much more students who are close to the borrowing constraint are willing to work relative to those who are not. As such, we can identify γ by comparing the labor supply decisions of poor students with those from rich students.

The psychic cost of college, κ_X , is identified by different rates of attending college by ability, gender, and parental education, after controlling for differences in utility coming from consumption and differences in future earnings. The parameters governing the value of working in college, ζ , σ^{ℓ^E} , and λ , are jointly identified by variation in college labor supply choices across agents and across periods. Specifically, the parameter vector ζ is identified by different

rates of working in college after controlling for differences in utility coming from consumption and differences in future earnings. The parameter governing the variance of the labor supply shock, σ^{ℓ^E} , is identified by variation in the timing of working in college decisions. For example, suppose that $\sigma^{\ell^E} = 0$. Then the labor supply decisions of identical agents would be exactly the same in each period. A larger σ^{ℓ^E} implies more variation in the labor supply choices of identical agents and across periods. The nesting parameter λ , is identified by the substitution patterns across labor supply decisions and dropping out. The warm-glow parameters, ϕ and c_b , are identified by the relationship between parental income and parental transfers. A larger value of ϕ increases the derivative of parental transfers on parental income. Decreasing c_b increases the level of transfers overall. Warm-glow utility only depends on the amount of transfers given, not on other things that may enter the child's problem (i.e., ability, tuition, number of siblings). The degree to which parental transfers respond to different children's characteristics will instead be determined by the strength of the altruism motive. Therefore, any differences in parental transfers across student characteristics will identify ω . For example, if students who face higher tuition levels generally receive higher parental transfers, this will identify ω and give us a sense of how much we expect parental transfers to be crowded out by financial aid. Parents' paternalism parameters, ξ_0 and ξ_{ParEd} , are identified by the ratio of college parental transfers to high school parental transfers. A higher value of ξ_0 implies higher transfers for children going to college relative to transfers for children entering the labor force directly. Finally, the parameter governing the standard deviation of observed parental transfers, σ^{etr} , is identified by the variance in observed parental transfers of identical agents.

B.7.2 Identification in a Simplified Model

In this section we use a simple model to show how the parameter governing the the agents' risk aversion, γ , is separately identified from the variance of the idiosyncratic taste for college σ^E .

Consider a simple static model where agents make a discrete choice over options j , representing schooling options. For an agent i , each option has an associated income Y_{ij} which could represent, for example, how much income an agent has left for consumption net of tuition costs. Let an agent's choice specific utility be given by:

$$u_{ij} = \frac{c_{ij}^{1-\gamma}}{1-\gamma} + \kappa_j + \sigma \varepsilon_{ij}$$

where κ_j is a parameter that is common to all agents who choose option j and ε_{ij} is an extreme value-type I preference parameter. Collectively, we can think of $-(\kappa_j + \varepsilon_{ij})$ as the psychic cost of option j . As this is a static model, let $c_{ij} = Y_{ij}$.

Note that $(1-\gamma)$ shows up in two places in the utility function. First it shows up as the denominator of $\frac{c_{ij}^{1-\gamma}}{1-\gamma}$ and therefore scales utility from consumption. This scaling factor

alone is not separately identified from the variance of the error term σ . However, $(1 - \gamma)$ also shows up as the exponent on consumption, and therefore dictates the curvature of utility from consumption. This is crucial for separately identifying γ from σ . Essentially if γ is large (and $(1 - \gamma)$ small), then agents with lower levels of income will behave differently than agents with higher levels of income.

Suppose we can divide agents i into several different “regions”. All agents in a given region r face the same menu of incomes across options. We will therefore write all choice specific incomes as Y_{rj} . Note that agents across regions all have same preferences but face different income levels, representing differences in tuition across regions, differences in wage income in college, or differences in parental transfers.

The probability that an agent in region r chooses option j is given by:

$$P_{rj} = \frac{\exp\left(\frac{Y_{rj}^{1-\gamma}}{\sigma(1-\gamma)} + \frac{\kappa_j}{\sigma}\right)}{\sum_{j'} \exp\left(\frac{Y_{rj'}^{1-\gamma}}{\sigma(1-\gamma)} + \frac{\kappa_{j'}}{\sigma}\right)}$$

We can then write the log ratio of choice probabilities of choosing option j over k for agents in region r :

$$\log \frac{P_{rj}}{P_{rk}} = \frac{1}{\sigma} \left[\left(\frac{Y_{rj}^{1-\gamma}}{(1-\gamma)} + \kappa_j \right) - \left(\frac{Y_{rk}^{1-\gamma}}{(1-\gamma)} + \kappa_k \right) \right] \quad (21)$$

Now consider the difference in this log ratio between agents in region r and agents in region \hat{r} :

$$\log \frac{P_{rj}}{P_{rk}} - \log \frac{P_{\hat{r}j}}{P_{\hat{r}k}} = \frac{1}{(1-\gamma)\sigma} [(Y_{rj}^{1-\gamma} - Y_{rk}^{1-\gamma}) - (Y_{\hat{r}j}^{1-\gamma} - Y_{\hat{r}k}^{1-\gamma})] \quad (22)$$

Before completing the proof, we consider a simple example here to build intuition. Suppose agents in region \hat{r} are poorer by a positive amount τ for both options, relative to agents in region r . That is: $Y_{\hat{r}j} = Y_{rj} - \tau$ and $Y_{\hat{r}k} = Y_{rk} - \tau$. Then $\gamma = 0$ would imply that $\log \frac{P_{rj}}{P_{rk}} - \log \frac{P_{\hat{r}j}}{P_{\hat{r}k}} = 0$ — the ratio of choice probabilities in the poor region is equal to the ratio in the rich region. Larger γ implies that the poorer region will choose the relatively higher income option more relative to the richer region.

Consider now a third region, \tilde{r} . We now also have

$$\log \frac{P_{rj}}{P_{rk}} - \log \frac{P_{\tilde{r}j}}{P_{\tilde{r}k}} = \frac{1}{(1-\gamma)\sigma} [(Y_{rj}^{1-\gamma} - Y_{rk}^{1-\gamma}) - (Y_{\tilde{r}j}^{1-\gamma} - Y_{\tilde{r}k}^{1-\gamma})] \quad (23)$$

combined with 22, this gives us two equations and two unknowns. Further, taking the ratio of 22 and 23 yields:

$$\frac{\log \frac{P_{rj}}{P_{rk}} - \log \frac{P_{\hat{r}j}}{P_{\hat{r}k}}}{\log \frac{P_{rj}}{P_{rk}} - \log \frac{P_{\tilde{r}j}}{P_{\tilde{r}k}}} = \frac{(Y_{rj}^{1-\gamma} - Y_{rk}^{1-\gamma}) - (Y_{\hat{r}j}^{1-\gamma} - Y_{\hat{r}k}^{1-\gamma})}{(Y_{rj}^{1-\gamma} - Y_{rk}^{1-\gamma}) - (Y_{\tilde{r}j}^{1-\gamma} - Y_{\tilde{r}k}^{1-\gamma})}$$

Table 5: Maximum Likelihood estimates

		Estimate	Standard Error
College Utility: $U^E(c, \ell) = \frac{c^{1-\gamma}}{1-\gamma} - \kappa_{\theta,d} - \zeta^{\ell^E} + \varepsilon_{it}^{\ell^E}$			
Curvature of Utility	γ	1.9	0.0057
No work	ζ^0	0	(Normalization)
Part time	ζ^1	0.057	3.6*
Full time	ζ^2	0.18	4.1*
Standard Deviation of Enrollment Shock	$\sigma^{\{E\}}$	7.8*	.069*
Dropout cost	δ	1.2	1.7*
Standard Deviation of Working Shock	$\sigma^{\{\ell^E\}}$	0.40	.0078*
Nesting Parameter	λ	0.59	0.38*
Psychic Cost: $\kappa_{\theta,d} = \kappa_0 + \kappa_{\theta} \log(\theta_i) + \kappa_{\text{fem}} \mathbb{I}(s = \text{female}) + \kappa_{\text{ParEd}} \text{ParEdu}_i$			
Constant	κ_0	0.44	0.15*
Ability Interaction	κ_{θ}	-8.3*	0.014*
Female Dummy	κ_{fem}	-0.083*	0.00075*
Parental Education	κ_{ParEd}	-1.7*	0.0049*
Parental Utility from Transfers:			
$F(tr^H, tr^E, \Omega_i) = \underbrace{\omega EV(\Omega_i tr^H, tr^E)}_{\text{Altruism}} + \mathbb{E} \left[\underbrace{(\xi_0 + \xi_{\text{ParEd}} \text{ParEdu}_i) \mathbb{1}_E}_{\text{Paternalism}} + \underbrace{\phi \frac{(c_b + tr^e)^{1-\gamma}}{1-\gamma}}_{\text{Warm Glow}} \right]$			
Altruism	ω	3.0	0.021
Prestige Constant	ξ_0	0.27	0.0017
Parent's Education Interaction	ξ_{ParEd}	0.69*	0.00024*
Warm Glow Strength	ϕ	0.09	0.0016
Warm Glow Level	c_b	-3.4 **	0.059**

* we display 10,000 times the parameter value.

** we display the parameter value divided by 10,000.

which identifies γ . σ is therefore identified by 22. Finally, we must normalize one of the $\kappa_j = 0$. Then all of the κ_j terms are identified by 21.

B.8 ML Estimates

The parameter estimates are contained in Table 5.

B.9 Graduation Rates and Enrollment by Gender

Figure 11 shows the college graduation rates as a function of parental income and ability in the model and in the data. The model is able to replicate these moments well.

Figure 12(a) shows the college enrollment rates for male and female students as a function of parental income in the model and in the data. Figure 12(b) shows the college enrollment rates for male and female students as a function of ability in the model and in the data. We can see that the model is able to replicate these moments quite well.

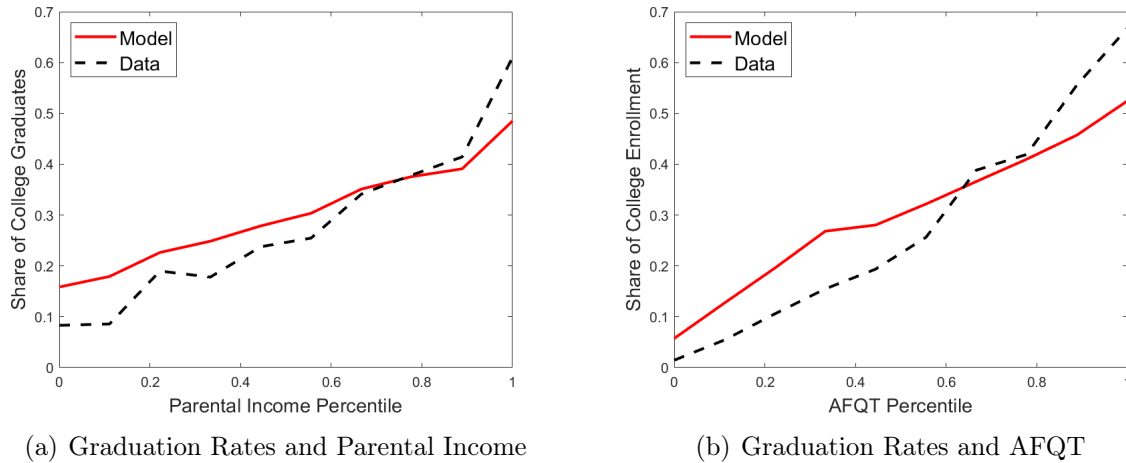


Figure 11: Graduation and Enrollment Rates

Notes: The solid (red) line shows simulated enrollment shares by parental income and AFQT percentile. This is compared to the dashed (black) line which shows the shares in the data.

B.10 Earnings Profiles Model


Figure 13 shows the simulated average for college graduates and high school graduates as a function of age.

B.11 Untargeted Moments


Responsiveness of Enrollment to Grant Increases. Many papers have analyzed the impact of increases in grants or decreases in tuition on college enrollment. Deming and Dynarski (2009) survey the literature. The estimated impact of a \$1,000 increase in yearly grants (or a respective reduction in tuition) on enrollment ranges from 1 to 6 percentage points, depending on the policy reform and research design. A more recent study by Castleman and Long (2016) looks at the impact of grants targeted to low-income children. Applying a regression-discontinuity design for need-based financial aid in Florida (Florida Student Access Grant), they find that a \$1,000 increase in yearly grants for children with parental income around \$30,000 increases enrollment by 2.5 percentage points.

Simulating a \$1,000 increase in financial aid for all individuals in our model leads to a 1.69 percentage point increase in overall enrollment rates and a 2.06 percentage point increase for students near the studied discontinuity in Castleman and Long (2016). Overall, our simulated elasticities are fairly consistent with these reduced-form estimates. This gives us confidence in our maximum likelihood estimates, especially given that these reduced form estimates were not targeted in estimation.

Importance of Parental Income. A well-known empirical fact is that individuals with higher parental income are more likely to receive a college degree (see also Figure 2). However,



(a) Enrollment Rates and Parental Income



(b) Enrollment Rates and AFQT

Figure 12: Graduation and Enrollment Rates by Gender

Notes: The panel on the left shows the relationship between enrollment rates and parental income in the model and in the data for females and males. The panel on the right shows the relationship between enrollment rates and ability in the model and in the data for females and males.




Figure 13: College and High School Graduate Earnings Profiles.

it is not obvious whether this is primarily driven by parental income itself or by variables correlated with parental income and college graduation. Using income tax data and a research design exploiting parental layoffs, Hilger (2016) finds that a \$1,000 increase in parental income leads to an increase in college enrollment of 0.43 percentage points. To test our model, we increased parental income for each individual by \$1,000 and obtained increases in college enrollment by 0.18 percentage points. Our model predicts a moderate direct effect of parental income, smaller but in line with Hilger (2016).

Returns for Marginal Students. We find a return to one year of schooling of 12.1% for marginal students. This reflects that marginal students are of lower ability on average than inframarginal students and is also in line with Oreopoulos and Petronijevic (2013). A clean way to infer returns for marginal students is found in Zimmerman (2014). In his study, students are marginal with respect to academic ability, measured by a GPA admission cutoff. He finds

that these students have earnings 22% higher than those just below the cutoff, when earnings are measured 8 to 14 years after high school graduation. We perform a similar simulation and make use of the fact that the NLSY also provides GPA data. In fact, our model gives a return to college of 26.3%, measured 8 to 14 years after high school graduation, for students with a GPA in this neighborhood.⁵⁹

C Additional Results

C.1 Marginal and Inframarginals Evaluated at Current Financial Aid Levels

In Section 6.1, we plotted the share of marginal enrollees and inframarginal enrollees at a flat financial schedule for a number of model specifications. In this section, we repeat this exercise but plot the share of marginal enrollees and inframarginal enrollees at the current financial aid schedule.

The results are very similar to those presented in Section 6.1. The relationship between parental income and share of inframarginal students has become weaker (and eventually becomes negative), reflecting that the current financial aid schedule is decreasing in parental income, see Figure 14(a). Further, children with high income parents are more likely to be marginal with respect to financial aid relative to graph is in Section 6.1, again reflecting that they receive less financial aid than children with low income parents, see Figure 14(b).

C.2 An Alternative Decomposition

In Section 6.1, we perform a model-based decomposition exercise to better understand which drive the optimal progressivity result. In this appendix, we perform a similar decomposition but alter the order in which we change various components to the model. In particular, we first remove the relationship between parental income and parental transfers, before proceeding to remove the relation between parental income and ability and the relation between parental education and the psychic costs of college. As before, all changes to the model specification are cumulative.

We first analyze the determinants of the positive relation between college enrollment and parental income in Figure 15(a) and the negative relationship between share of marginal students and parental income in Figure 15(a). The simulated relationships at a flat financial aid

⁵⁹Finally, we do not account for differing rates of unemployment and disability insurance rates. Both numbers are typically found to be only half as large for college graduates (see Oreopoulos and Petronijevic (2013) for unemployment and Laun and Wallenius (2016) for disability insurance). Further, the fiscal costs of Medicare are likely to be much lower for individuals with a college degree. Lastly, we assume that all individuals work until 65 not taking into account that college graduates on average work longer (Laun and Wallenius, 2016). These facts would generally strengthen the case for an increase in college subsidies.




Figure 14: Model-Based Decomposition for Marginal and Inframarginal Students at Current Grant Schedule

Notes: We plot the share of college enrollees and marginal college enrollees given the current US aid schedule for different model specifications. The solid red line represent the baseline model. For the dashed black line we simulate a model version for which we remove the correlation between ability and parental income. For the dashed-dotted blue line we simulate a model version for which we additionally remove the correlation between the psychic costs and parental education. For the dotted pink line we simulate a model version for which additionally removes labor market riskiness; i.e. education decisions are made with no uncertainty about future wages. For the turquoise line with crosses we simulate a model version for which we set parental transfers to the mean parental transfers in the data, conditional on education.

schedule are shown in the solid lines in the two figures. In this baseline case, college enrollment rates are strongly increasing in parental income while the share of marginal students are strongly decreasing in parental income. Next, in the turquoise lines, we set parental transfers exogenously to the mean levels for enrollees and non-enrollees and assume no families are eligible for subsidized Stafford loans. From Figure 15(a) we can see that the positive relation between college enrollment and parental income weakens slightly. The relation between parental income and share of marginal enrollees, however, flattens completely. The black dotted line and the blue dash-dotted line show the cases in which we remove the correlation between parental income and ability and in which we remove the relation between parental education and psychic costs, respectively. After removing these two factors there is no interesting heterogeneity between parental income groups. Removing these two relationships both weak the relationship between parental income and enrollment. In both these simulations, the gradient between parental income and share of marginal students remains flat.

We now simulate the respective optimal financial aid schedule under each model specification in Figure 16. When we remove the relationship between parental income and parental transfers (the turquoise line in Figure 16), the optimal financial aid schedule flattens. This flattening of the optimal schedule occurs because the relationship between parental income

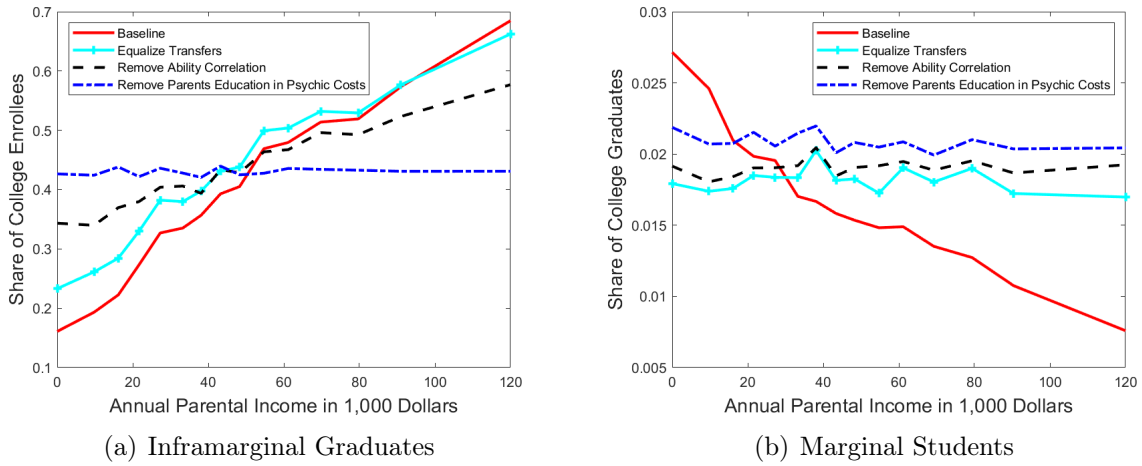


Figure 15: Alternative Model-Based Decomposition for Marginal and Inframarginal Students at Flat Grant Schedule

Notes: We plot the share of college enrollees and marginal college enrollees given a flat financial aid schedule for different model specifications. The solid red line represent the full model. For the turquoise line with crosses we simulate a model version for which set parental transfers to the mean parental transfers in the data, conditional on education. For the dashed black line we simulate a model version for which we remove the correlation between ability and parental income. For the dashed-dotted blue line we simulate a model version for which we additionally remove the correlation between the psychic costs and parental education.

and share of marginal enrollees is flat. However, there optimal aid is still positive—ranging from above \$6,500 for the poorest families to below \$4,500 for the wealthiest families. The optimal aid schedule is progressive because high income children are still much more likely inframarginal. When we remove the positive ability-income correlation (the black dashed line) and the relationship between parental education and psychic costs we flatten the relationship between parental income and share of inframarginal students. The optimal aid schedule flattens as a result.

C.3 Decomposition with Removal of Borrowing Constraints

In Figures 17(a) and 17(b), we perform the same decomposition we performed in Section 6.1 but additionally remove borrowing constraints before equalizing parental transfers. We additionally assume that no families are eligible for subsidized Stafford loans throughout the decomposition. Figure 18 shows the resulting optimal financial aid under each model specification.

As before, the red line shows the baseline case, the black dotted line shows the case where we remove the ability correlation, the blue dash-dotted line shows the case where we remove the correlation between psychic cost of parental education, and the dotted pink line show the case with no labor market uncertainty. These lines tell essentially the same story as the decomposition in Section 6.1.




Figure 16: Optimal Financial Aid for Different Model Specifications

Notes: For each model specification (see Figure 15), we illustrate the respective optimal financial aid schedule.

The green dotted lines show the case in which we remove borrowing constraints. As a result, the number of inframarginal students increases for all income groups, as student no longer have to deal with borrowing constraints in college. Additionally, the share of marginal enrollees drops substantially for all parental income groups. As students are no longer affected by borrowing constraints in college, the marginal benefit of additional financial aid decreases substantially.

However, despite the fact that both the gradients of marginal and inframarginal enrollees are flat, the optimal aid is still slightly decreasing in parental income. This is because, once the correlations of parental income with marginal and inframarginal students have been shut down, the differences in marginal social welfare weights play a role. We find that at the flat financial aid schedule, the marginal social welfare weight of the poorest children is roughly 20% higher than that of the richest students. Essentially, given that enrollment is so unresponsive to financial aid, the social planner allocates financial aid to agents with the highest marginal social welfare weights. This leads to a slightly progressive financial aid schedule.

Equalizing parental transfers on top of this removes these differences in marginal social welfare weights and therefore leads to an flat optimal aid schedule.

C.4 The Role of Borrowing Constraints

Figure 19(a) shows the optimal financial aid policies when we have abolished borrowing constraints. We first remove borrowing constraints and keep the current financial aid system. This will increase college enrollment and imply a windfall fiscal gain for the government. In a second step, we choose optimal financial aid but restrict the government to not use this windfall gain. Figure 19(b), shows the implied graduate patterns.

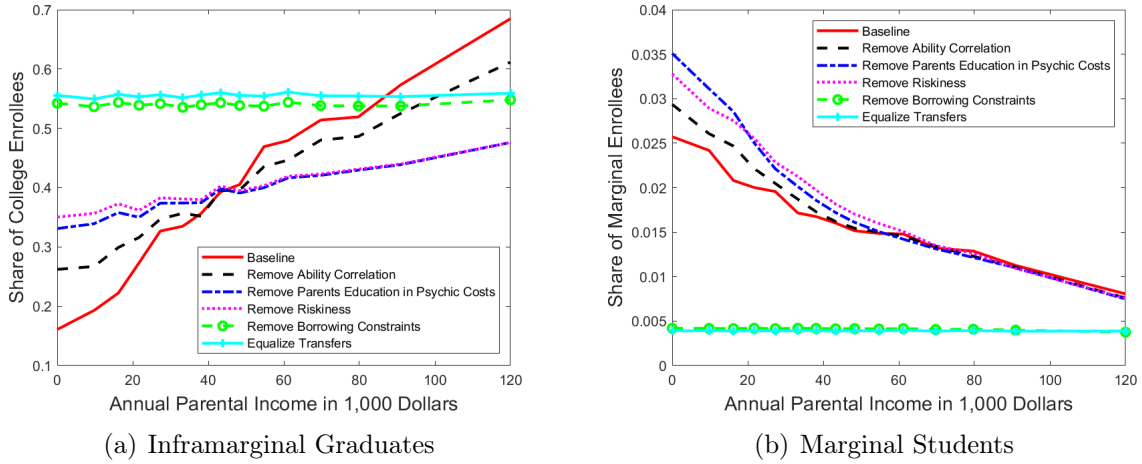


Figure 17: Model-Based Decomposition for Marginal and Inframarginal Students at Flat Grant Schedule with Removal of Borrowing Constraints

Notes: We plot the share of college enrollees and marginal college enrollees given a flat financial aid schedule for different model specifications. We assume no subsidized Stafford loans for all specifications. The solid red line represent the full model at the flat financial aid schedule. For the dashed black line we simulate a model version for which we remove the correlation between ability and parental income. For the dashed-dotted blue line we simulate a model version for which we additionally remove the correlation between the psychic costs and parental education. For the dotted pink line we simulate a model version for which on top removes any riskiness; i.e. education decisions are made under perfect foresight. For the dashed green line with circles we simulate a model version for which on top we remove all borrowing constraints. For the turquoise line with crosses we simulate a model version for which set parental transfers to the mean parental transfers in the data, conditional on education.

C.5 Varying Borrowing Constraints

To get a sense of how varying borrowing constraints would affect our main conclusions, we have re-estimated a version of the model in which the borrowing limit depends on parental resources. Here, it was very hard for us to get guidance on what would be a reasonable way to have exogenous borrowing constraints depend on parental income and ability of the child. Hence, we have decided to report a very simple and transparent case in the paper: we assume that children whose both parents have a college degree can borrow twice the amount of the Stafford loan limit. Admittedly, this is ad-hoc in two ways. The first ad-hoc decision is to separate children along the parental education dimension. Our motivation was that parental education strongly correlates with both parental earnings and child’s ability. The second ad-hoc decision we faced was: how much more can these children with highly educated parents borrow? We here decided to just double the amount in the case that we report.

The optimal utilitarian financial aid with parental education dependent borrowing constraints are shown in Figure 20. The shape is slightly different from the baseline optimal




Figure 18: Optimal Financial Aid for Different Model Specifications

Notes: For each model specification (see Figure 17), we illustrate the respective optimal financial aid schedule.

schedule, as changes in the borrowing constraints lead to changes share of marginal students.⁶⁰ However, the optimal financial aid is still highly progressive.

C.6 Details: Endogenous Ability

We assume that initial ability θ_0 is distributed as:

$$\ln \theta_0 = \beta_0 + \beta_1 \ln I + u$$

where u is normally distributed. We choose β_0 , β_1 , and the variance of u to match the mean and variance of log childhood ability and covariance of log childhood ability and log parental income from Agostinelli and Wiswall (2016).

We need to calibrate the parameters of the childhood ability production function:

1. A - TFP of parental production function.
2. γ_1 - weight on initial ability
3. γ_2 - weight on parental investment
4. γ_3 - interaction term
5. σ^u - variance of u .

⁶⁰As we have shown earlier, relaxing borrowing constraints for all students reduces the progressivity of the optimal aid schedule. That force is still present here, as some low income students have two college educated parents. However, this force is partially muted by the fact that parental education is increasing in parental income. As such, the optimal aid schedule here is more progressive than the case with relaxed borrowing constraints for all individuals, but slightly less progressive than the baseline case with equal borrowing constraints for all students.

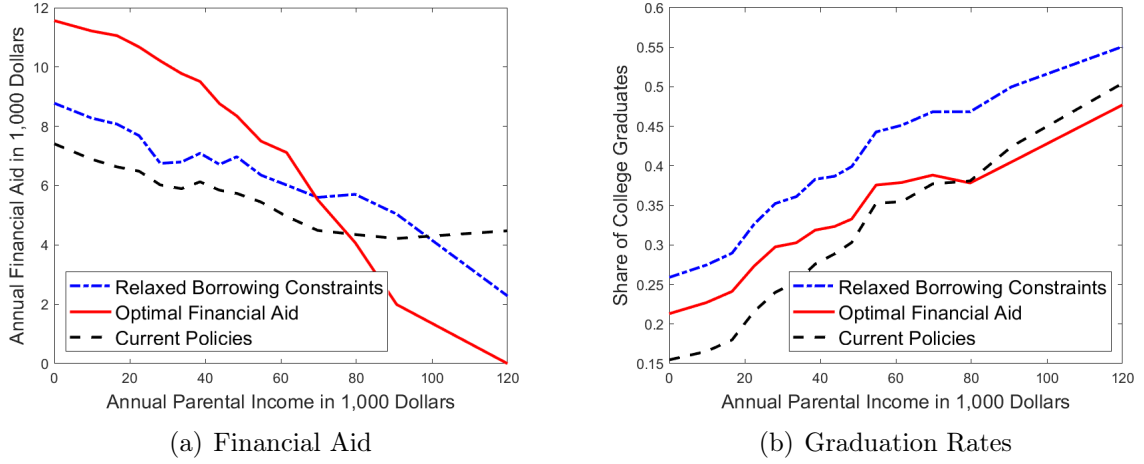


Figure 19: Financial Aid and Graduation with Free Borrowing

Notes: The dashed-dotted (blue) line shows the optimal schedule with no borrowing constraints. Optimal financial aid with a Utilitarian welfare function and with borrowing constraints and current financial aid are also shown for comparison in Panel (a). In Panel (b) we display the college graduation share by parental income group for each of the three scenarios.

Agostinelli and Wiswall (2016) estimate a translog production function of the following form:

$$\ln \theta_{t+1} = \ln A_t + \gamma_{1t} \ln \theta_t + \gamma_{2t} \ln I_t + \gamma_{3t} \ln \theta_t \cdot \ln I_t + \eta_{\theta,t}$$

for $t = 0, 1, 2, 3$. By combining these four equations, we can derive a single equation for end of childhood ability $\ln \theta_4$ as a function of initial ability $\ln \theta_0$, parental investment in each period $\ln I_t$, the yearly shocks $\eta_{\theta,t}$, and the technology parameters.

Specifically, after some algebra we can write

$$\ln \theta_4 = \ln \theta_0 (\gamma_{30} \ln I_0 + \gamma_{10}) (\gamma_{13} + \gamma_{33} \ln I_1) (\gamma_{12} + \gamma_{32} \ln I_2) (\gamma_{11} + \gamma_{31} \ln I_1) + f(I, A, \gamma)$$

where $f(I, A, \gamma)$ is a function that depends on investment and the other parameters but not directly on initial ability $\ln \theta_0$.

We can further rearrange this equation to yield

$$\ln \theta_4 = \tilde{\gamma} \ln \theta_0 + g(\theta_0, I_0, I_1, I_2, I_3) + f(I, A, \gamma)$$

where

$$\tilde{\gamma} = \gamma_{10} \gamma_{11} \gamma_{12} \gamma_{13}$$

and

$$g(\theta_0, I_0, I_1, I_2, I_3) = \ln \theta_0 (\gamma_{30} \ln I_0 + \gamma_{10}) (\gamma_{13} + \gamma_{33} \ln I_1) (\gamma_{12} + \gamma_{32} \ln I_2) (\gamma_{11} + \gamma_{31} \ln I_1) - \tilde{\gamma} \ln \theta_0$$




Figure 20: Optimal Financial Aid with Parental Education Dependent Borrowing Constraints

Notes: The dashed-dotted (blue) line shows the optimal schedule with parental income dependent borrowing constraints. Optimal financial aid with in the baseline case and current financial aid are also shown for comparison in Panel (a). In Panel (b) we display the college graduation share by parental income group for each of the three scenarios.

We set γ_1 equal to the product of the coefficients on lagged ability from Agostinelli and Wiswall (2016) $\gamma_{1,1}\gamma_{1,2}\gamma_{1,3}\gamma_{1,4} \approx 2$. This approximation will be true if the terms on the interaction terms in Agostinelli and Wiswall (2016) are close to zero. Agostinelli and Wiswall (2016) estimate $\gamma_{30} = -0.105(0.066)$, $\gamma_{31} = -0.005(0.019)$, $\gamma_{32} = -0.003(0.013)$, $\gamma_{33} = 0.003(0.010)$. None of the estimates are statistically different from 0 at 95% confidence level and only the first one at a 90% confidence level. Therefore, we think calibrating $\gamma_1 = 2$ seems like a reasonable choice.

Then we have four parameters, A , γ_2 , σ^ι , and γ_3 . We choose these parameters to match the four following moments:

1. Mean of θ
2. Variance of θ
3. Covariance of θ and parental income I .
4. From Agostinelli and Wiswall (2016): The effect on realized years of schooling of a monetary transfer to parents is roughly ten times larger for parents in the 10th percentile of the income distribution compared to those in the 90th percentile.

Loosely speaking, the covariance of θ and I helps to pin down the importance of parental monetary investments γ_2 . The variance of θ helps to pin down the variance to shock of ability production, σ^ι . The differential effect of monetary transfers for rich and poor parents helps to pin down the interaction between parental investment and initial ability, γ_3 . Finally, the average ability level helps to discipline the TFP of the production function, γ_1 .

Finally, we need to translate these measures of final ability, which are in the units used in Agostinelli and Wiswall (2016) into our measure of ability, which is based on AFQT scores.

Let $\hat{\theta}$ represent end of childhood ability as measured in the units used in Agostinelli and Wiswall (2016). We assume that our measures of ability θ is a linear projection of this log skill measure

$$\theta = \alpha_0 + \alpha_1 \ln \hat{\theta}$$

where we choose α_1 and α_0 to match the mean and variance of our AFQT measure. Therefore, when we simulate the model, we first simulate childhood ability in the units used in Agostinelli and Wiswall (2016). Then we translate the measures of ability in Agostinelli and Wiswall (2016) to the ability measures we use in this paper.

Changes in Childhood Ability Figure 21 shows the change in the relationship between parental income and ability as a result of switching from the current financial aid system to the optimal system with endogenous ability. Ability is measured in percentiles of AFQT scores where percentiles are evaluated at their current levels. We can see that switching to the optimal aid schedule leads to substantial increases in child ability, especially for children in the lower end of the parental income distribution.




Figure 21: Ability Levels with Endogenous Ability

Notes: This figure shows the relationship between parental income and ability in the optimal system with endogenous ability and under the current financial aid system. Ability is measured in percentiles of the AFQT distribution before financial aid is re-optimized.

C.7 Endogenous Ability with Parental Borrowing Constraints

One issue with the preceding analysis is that we have assumed that parents do not face borrowing constraints. Poor parents may be borrowing constrained while their children are young and therefore may not be able to increase investment in their children in response to changes in financial aid. To explore how borrowing constraints would affect the optimal policy, we assume that $P\%$ of parents without a college education cannot increase their investment in their children while the remainder of parents may choose their investment level without this

constraint.⁶¹ The optimal policy for a range of values of P is displayed in Figure 22. We can see that the optimal progressivity of the system decreases as we increase the percentage of low-education families who are borrowing constrained. However, the optimal schedule remains more progressive than the current schedule in all cases.




Figure 22: Financial Aid, Graduation and Ability Levels with Endogenous Ability and Parental Borrowing Constraints

Notes: In Panel (a), each line shows the optimal financial aid with endogenous ability when P percent of low-education parents are borrowing constrained and therefore cannot adjust their child's ability in response to changes in financial aid. In Panel (b) we display the college graduation share for each of these scenarios. Panel (c) shows the relationship between parental income and ability in each scenario. Ability is measured in percentiles of the AFQT distribution before financial aid is re-optimized.

C.8 General Equilibrium Effects on Wages

Our analysis abstracted from general equilibrium effects on relative wages. Accounting for these effects would imply that the effects of financial aid on enrollment might be mitigated in

⁶¹Caucutt and Lochner (2017) find that 20% of parents with a high school degree and young children are borrowing constrained. Of course, borrowing constraints will also affect the investment decisions of parents who are not at the borrowing limit.

the long run: if more individuals go to college, the college wage premium should be expected to decrease because of an increase in the supply of college educated labor (Katz and Murphy, 1992). This in turn would mitigate the initial enrollment increase. To investigate the role of general equilibrium effects on our results, we recalculate the optimal financial aid schedule under the assumption that wages are determined in equilibrium. We assume firms use a CES production function that combines total efficiency units of labor supplied by skilled and unskilled workers, implying that wages are determined by the ratio of skilled to unskilled labor. We assume an elasticity of substitution between skilled and unskilled workers of 2.

We assume identical perfectly competitive firms use CES production functions which combine skilled and unskilled labor. Therefore, wages are determined as a function of the ratio of the total skilled labor to the total unskilled labor.

Let P^U and P^S denote the endogenously determined efficiency wages for unskilled and skilled workers, respectively, where skilled workers are those with a college degree and unskilled workers are high school graduates. We allocate half of college dropouts to each of the skill groups, as is common in the literature (e.g. Card and Lemieux (2001)). Suppose an agent's wages can be written as the product of her efficiency wage and her quantity of efficiency units of labor supplied: $w_{it} = P^{sk} H_{it}$, where $sk \in \{unskilled, skilled\}$ denotes skill level and H_{it} denotes agent i 's level of human capital.⁶²

We assume perfectly competitive labor markets. Production at the representative firm is a CES function combining skilled and unskilled labor:

$$Y = A \left(\lambda S^{(\sigma-1)/\sigma} + (1 - \lambda) U^{(\sigma-1)/\sigma} \right)^{\sigma/(\sigma-1)}$$

where A is total factor productivity, λ is the factor intensity of skilled labor, and σ is the elasticity of substitution between skilled and unskilled labor. We assume $\sigma = 2$. S and U represent the total amount of human capital units supplied by skilled and unskilled workers. We assume the economy is in a long run steady-state equilibrium, and that the economy consists of identical overlapping cohorts. Therefore, as cohorts are identical, the total labor supply in the steady-state equilibrium is equal to the total amount of labor supplied over the life-cycle for a given cohort.

Therefore, we can write:

$$S = \sum_i \sum_t H_{it} \ell_{it} \mathbb{I}(sk_i = skilled)$$

and

⁶²We normalize units of human capital such that $H_{it} = 1$ is an efficiency unit of labor is defined as the labor supplied by a male worker whose log wages at age 18 are equal to the constant of the wage equation. Therefore, the constants of the wage functions for skilled and unskilled workers are equal to the logs of the efficiency wages for skilled and unskilled workers.

$$U = \sum_i \sum_t H_{it} \ell_{it} \mathbb{I}(sk_i = \text{unskilled})$$

Efficiency wages are given by the first order conditions of the firm's profit maximization problem:

$$P^S = A (\lambda S^{(\sigma-1)/\sigma} + (1 - \lambda) U^{(\sigma-1)/\sigma})^{1/(\sigma-1)} \lambda S^{-1/\sigma}$$

and

$$P^U = A (\lambda S^{(\sigma-1)/\sigma} + (1 - \lambda) U^{(\sigma-1)/\sigma})^{1/(\sigma-1)} (1 - \lambda) U^{-1/\sigma}.$$

These two functions determine wages endogenously as functions of labor supply.

The optimal financial aid schedule and graduation rates with general equilibrium wages are shown in Figures 23(a) and 23(b). We can see that the overall amount of aid has decreased slightly as the fiscal externality of college has been scaled down by general equilibrium wage effects. However, the optimal aid schedule with endogenous wages is just as progressive as in the case with exogenous wages. Thus, while general equilibrium wages dampen the effectiveness of financial aid overall, they do not lead to dramatic changes in the *relative* benefit of financial aid increases for students of different parental income levels. Hence, whereas the overall (average) generosity of the optimal financial aid schedule is slightly lower, the implications for how financial aid should vary with parental income are unchanged.⁶³

C.9 Jointly Optimal Financial Aid and Income Taxation

The size of the fiscal externality of college education depends on the tax and transfer system in place. Our structural estimates took the current US tax system as given. An interesting question to ask is how optimal subsidies change when the tax schedule is chosen optimally. To address this, we enrich the optimal policy space such that the planner can also pick a nonlinear tax function $T(y)$ as is standard in the public finance literature (Piketty and Saez, 2013).⁶⁴

First, the optimal formulas for the subsidy schedule are unchanged and still given by the formulas in Section 2. In Appendix A.6, we show what the endogenous extensive education

⁶³Our results are, hence, consistent with the important earlier paper(s) by Heckman et al. (1998). They find that GE effects dampen the effectiveness of tuition subsidies, and in our case the average level of financial aid is also affected.

⁶⁴We abstract from education dependent taxation; for such cases please see Findeisen and Sachs (2016) and Stantcheva (2017).

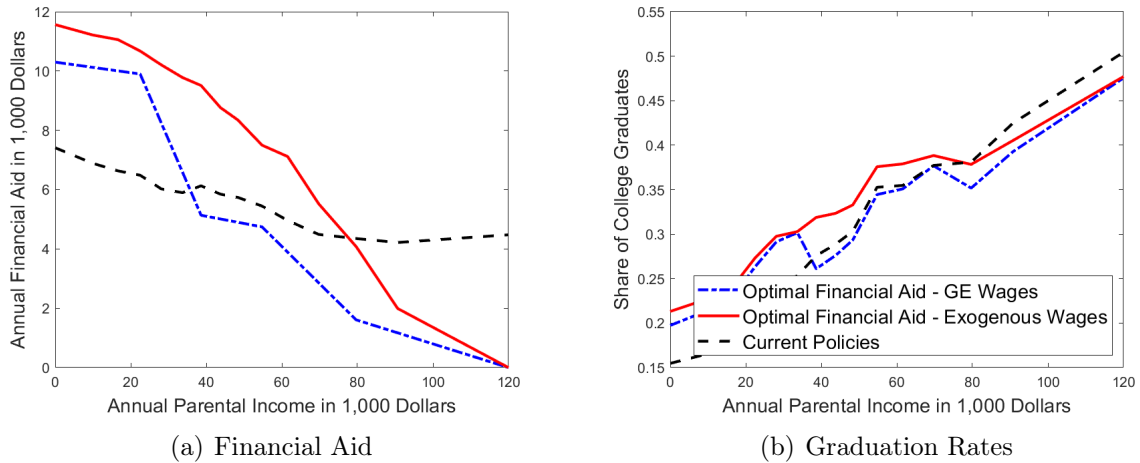


Figure 23: Financial Aid and Graduation with General Equilibrium Wages

Notes: The dashed-dotted (blue) line shows the optimal schedule when wages are determined in equilibrium. Production is CES between skilled and unskilled workers with an elasticity of substitution of 2. Optimal financial aid with exogenous wage rates and current financial aid are also shown for comparison in Panel (a). In Panel (b) we display the college graduation share by parental income group for each of the three scenarios.

margin implies for optimal marginal tax rates.⁶⁵ For the sake of brevity, we discuss the theory only in the Appendix and now move on to the quantitative implications of optimal taxes. We assume that agents are borrowing constrained and the government only (besides the tax schedule) maximizes the need-based element of the financial aid schedule. Results are barely changed if borrowing constraints are relaxed and/or the merit-based element is chosen optimally as well.

Figure 24(a) displays optimal average tax rates in the optimal as well as in the current US system. Average tax rates are higher for most part of the income distribution. As Figure 24(b) shows, this is driven by higher marginal tax rates throughout but especially at the bottom of the distribution, a familiar result from the literature (Diamond and Saez, 2011). In unreported results, we find that the direct effect of taxes on enrollment decisions, which we discussed in Section 3, is very small. In particular, it does not overturn the optimal U-shaped pattern of optimal tax rates nor does it influence the optimal top tax rate which is still mainly determined by the interaction of the labor supply elasticity and the Pareto parameter of the income distribution (Saez, 2001).

Figure 25(a) illustrates optimal financial aid in the presence of the optimal tax schedule. First, notice that financial aid is significantly higher on average compared to the case with the current US tax code. Higher income tax rates increase the fiscal externality, which increases the optimal level of the college subsidy (i.e. financial aid). Second, strikingly, the progressivity of

⁶⁵The formula is therefore related to the formulas of Saez (2002) and Jacquet et al. (2013), where the extensive margin is due to labor market participation, or Lehmann et al. (2014) where the extensive margin captures migration.

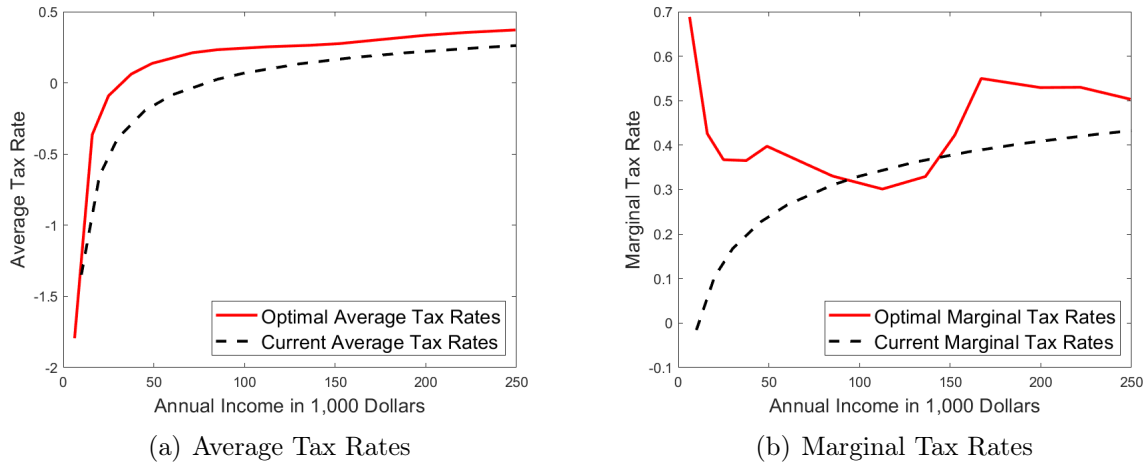


Figure 24: Optimal versus Current: Average and Marginal Tax Rates

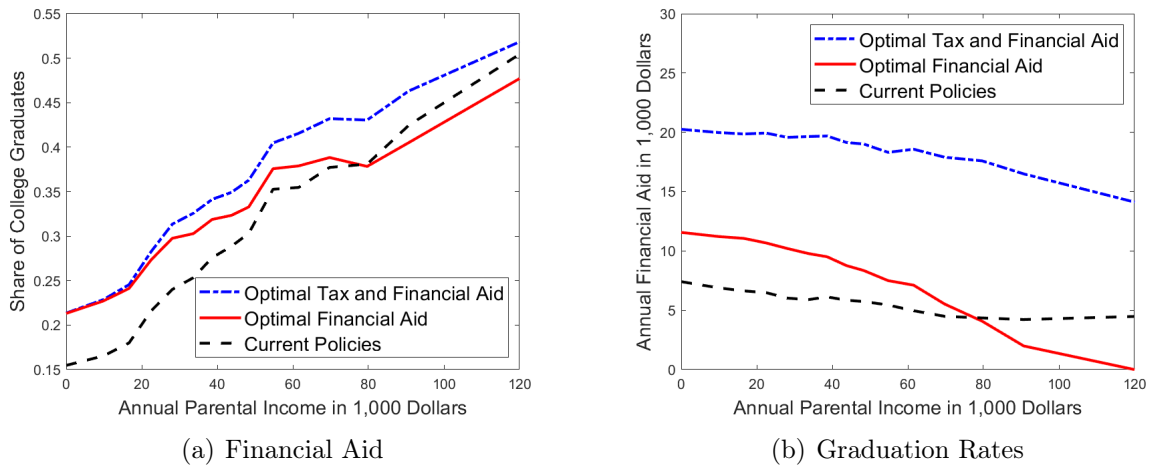



Figure 25: Financial Aid and Graduation with Optimal Tax Schedule

Notes: The dashed-dotted (blue) line shows the optimal schedule when the tax schedule is also chosen optimally. Optimal financial aid with the current tax schedule and current financial aid are also shown for comparison in Panel (a). In Panel (b) we display the college graduation share by parental income group for each of the three scenarios.

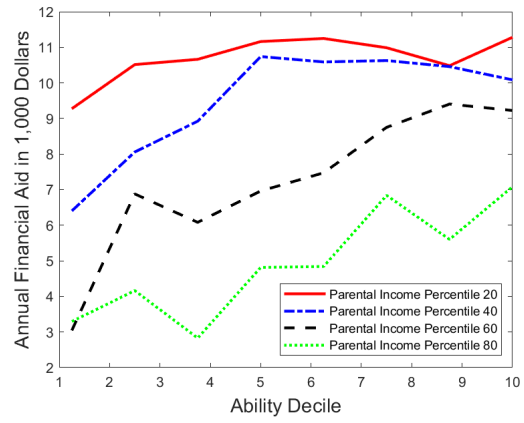
optimal financial aid policies is preserved. Progressive taxation does not change the desirability of progressive financial aid policies.

C.10 Merit-Based Financial Aid

Up to now, we have assumed that the merit-based element of financial aid policies stays unaffected. We now allow the government to optimally choose the gradient in merit and parental income. Figure 26(a) shows that – if optimally targeted also in terms of merit – financial aid policies can be more generous. The progressive nature however is even slightly reinforced.



(a) Financial Aid



(b) Illustration of Optimal Merit Based Element

Figure 26: Optimal Need and Merit Based Financial Aid

Notes: The dashed-dotted (blue) line shows the optimal financial aid for students with median ability as a function of income when the merit-based component of financial aid is also chosen optimally. Optimal financial aid with exogenous wage rates and current financial aid are also shown for comparison in Panel (a). In Panel (b) the merit based component of the optimal aid schedule.

Figure 26(b) shows how optimal financial aid is increasing in AFQT. Interestingly, the slope is almost independent of parental income.