

Optimal Need-Based Financial Aid*

Mark Colas
University of Oregon

Sebastian Findeisen
University of Konstanz

Dominik Sachs
LMU Munich

April 8, 2020

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. We estimate a model of selection into college for the United States that comprises multidimensional heterogeneity, endogenous parental transfers, dropout, labor supply in college, and uncertain returns. We quantify optimal financial aid in the estimated model and find it 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, John Kennan, Kory Kroft, Paul Klein, Tim Lee, Lance Lochner, Normann Lorenz, Thorsten Louis, Alex Ludwig, Marti Mestieri, Joseph Mullins, Nicola Pavoni, Michael Peters, Emmanuel Saez, Chris Taber, Aleh Tsyvinski, Gregory Veramendi, Gianluca Violante, Matthew Weinzierl, Nicolas Werquin, and participants from many seminars.

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.

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. We show that effective fiscal costs of financial aid are increasing in parental income and are therefore lowest for those children that are targeted by the Pell Grant, as we explain below. Further, we show that fiscal benefits are decreasing in parental income. 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.

We start with a general model without imposing restrictions on the underlying heterogeneity in the population. 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 college enrollees, whom we refer to as inframarginal students. A higher share of inframarginal students implies a higher marginal cost of increasing financial aid; if the planner increases financial aid by one dollar, these inframarginal students must all be paid an additional dollar without any effect on these students’ enrollment decisions. These costs are scaled down by the marginal social welfare weights attached to these students, which reflect the utility gain of the additional dollar for these inframarginal students. Optimal financial aid increases in the share of marginal students, the students that are at the margin of attending college with respect to financial aid, and the fiscal externality per marginal student, the change in lifetime fiscal contributions causal to college attendance. These two factors jointly capture the marginal benefits of the subsidy. 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 (Hendren,

2016) or about its value at a particular parental income level. Knowledge of these elasticities 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 structural parameters and therefore are not policy-invariant.

The main approach of this paper is therefore to estimate a structural model of selection into college that allows us to compute this policy elasticity along the parental income distribution and for alternative policies and therefore to solve for the optimal financial aid schedule. In our quantitative structural model we account for earnings risk, dropout, labor supply during college and, importantly, 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.

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 consider two alternative social planners who do not value redistribution: a social planner that sets equal marginal social welfare weights (Saez and Stantcheva, 2016) on all students and a planner that is only interested in maximizing tax revenues. Both would choose 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.

To elucidate the channels which drive the progressivity result, we show that parental income is positively associated with the share of inframarginal students and negatively associated with the share of marginal students, implying that increases in financial aid for high income families are both high cost and low benefit.¹ Both these factors are strong forces for need-based financial aid. We perform a model-based decomposition to assess which features are most important in generating these associations and therefore in the progressivity of optimal financial aid. We find that the correlations between parental income and parental transfers, preferences, and ability all play important roles in the progressive optimal aid schedule.²

¹One may have expected that efficiency considerations would make a case against need-based financial aid because of the positive correlation between returns to college and parental income. This correlation is indeed positive in our empirical model, and the fiscal externality of the average marginal student with high parental income is higher than for the average marginal student with low parental income. However, this effect is dominated by the relationship between inframarginal and marginal students with parental income.

²The role played by parental transfers is particularly large, as they interact with borrowing constraints. When borrowing constraints are eliminated first, heterogeneity in parental transfers plays a negligible role.

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 extends 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 to 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 very close to the empirical 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.³ Lawson (2017) uses an elasticity approach to characterize optimal uniform tuition subsidies for all college students.⁴ 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 referred to in Section 3 and online appendix 4.6 contains a more detailed review.

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 specify our quantitative model as a special case of the general model presented in Section 2 and present our estimation approach. Section 4 presents optimal financial aid policies, and Section 5 decomposes the forces

³Bohacek 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. Findeisen and Sachs (2018) focus on the implications of no-commitment on the side of the policy maker.

⁴Our 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.

which lead to an optimal financial aid schedule. In Section 6 we discuss further robustness issues. Section 7 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.

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 each following year before graduating, decide whether to drop out or continue in college. 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 the expected value of entering the labor force

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

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 $V_t^D(\cdot)$ 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:

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

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

Agents who do not drop out of college choose consumption and labor supply. At the end of the period, they either graduate or continue in college. The value function for staying enrolled is therefore 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{w} 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 .⁵ 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. These transfers 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 .⁶ We allow the interest rate for college enrollees to vary by the agent's asset position (positive or negative) and by 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.⁷ 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 .⁸ An individual enrolls in college if $V^E(X, I) \geq V^H(X, I)$.

Before we derive the optimal financial aid formulas, we ease the upcoming notation a bit. 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 respect to financial aid because the value of not dropping out, $V_t^{ND}(\cdot)$, is increasing in 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 $P^C(X, I, \mathcal{G}(X, I)) = \sum_{t=1}^{T_G} P_t^E(X, I, \mathcal{G}(X, I)) Pr_t^{Grad}(X)$. 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

⁵We 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.

⁶This allows for the possibility that high ability students may graduate from college in shorter time, for example.

⁷See Cunha et al. (2005), Heckman et al. (2006) or Heckman and Navarro (2007).

⁸The term $v(X)$ is only experienced at the time of enrollment. In contrast, $U^E(\cdot)$, the flow utility for someone enrolled in college, is a flow each year the agent is enrolled in college. Therefore, $v(X)$ only affects the enrollment decision while $U^E(\cdot)$ affects enrollment and dropout decisions.

$t = 1, 2$.⁹ 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.¹⁰

Further, 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.

2.2 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) :

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

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:

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

⁹We provide the optimal policy formulas without these simplifying assumptions in online appendix 1.2. The intuition of these formulas are the same but the notation is considerably more cumbersome.

¹⁰We 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 online appendix 6.7.

Finally, we turn to students who ultimately graduate from college. The average fiscal contribution of graduates of type (X, I) is given by the net present value of tax payments minus grants received, weighted by the probability of graduating in a given year:

$$\begin{aligned} \mathcal{NT}_{NPV}^G(X, I) &= \frac{1}{\sum_{g=3}^{T_G} P_g^E(X, I, \mathcal{G}(I)) P_{r_g}^{Grad}(X)} \times \\ &\sum_{g=3}^{T_G} P_g^E(X, I, \mathcal{G}(I)) P_{r_g}^{Grad}(X) \left[\sum_{t=g+1}^T \left(\frac{1}{1+r} \right)^{t-1} \mathbb{E}(\mathcal{T}(y_t) | X, I, G) - \mathcal{G}(I) \sum_{t=1}^g \left(\frac{1}{1+r} \right)^{t-1} \right] \end{aligned}$$

where T_G is the latest possible graduation date. Finally, expected fiscal contributions of an individual that decides to enroll are given by the weighted average of the fiscal contributions of college graduates and college enrollees:

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

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 $\mathcal{G}(I)$. 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.¹¹ 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:

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

¹¹To complete the picture, in online appendix 6.6, 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.

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 $\mathcal{G}(I)$. We show in online appendix 1.1, it has the following impact on welfare:

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

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

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

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

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

where ρ is the marginal value of public funds, U_c^E is the marginal utility of consumption, and $\mathbb{1}_{V_s^{ND} \geq V_s^D}$ is an indicator for an individual choosing not to drop out of college in year s . Thus, $W^E(I)$ is a money-metric (weighted by Pareto weights $\tilde{h}(X|I)$ and adjusted by the number of years students receive financial aid) average marginal social welfare weight of students of parental income I . It captures by how much welfare, measured in units of public funds, increases if inframarginal students of this parental income level receive one more dollar of financial aid. If the social welfare function is Utilitarian, then $W^E(I)$ measure how much additional yearly financial aid is valued, in units of public funds, by the average enrollee in

this income group. 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 $\left(1 + \frac{\partial \text{tr}_t^E(X, I, \mathcal{G}(I))}{\partial \mathcal{G}(I)}\right)$. Ceteris paribus, the stronger the crowding out of transfers, the lower are these welfare weights since fewer of the additional grants effectively reach students.¹²

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

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

where $\mathbb{1}_{H_j \rightarrow E_j}$ 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 $\int_{\mathcal{X}} \mathbb{1}_{H_j \rightarrow E_j} h(X|I) dX = \frac{\partial E(I)}{\partial \mathcal{G}(I)}$. $\Delta \mathcal{T}^E(X, I)$ 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:

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

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.

¹²Note 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.

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.¹³ 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, we specify and estimate such a model. As a complement to the quantitative optimal policy analysis, we consider a simplified model in online appendix 2, for which we can derive closed-form solutions.

3 Quantitative Model and Estimation

We now present a fully specified version of our model presented in Section 2. We start by stating the functional form assumption in Section 3.1. The model estimation is described in Section 3.2. The performance of the estimated model is discussed in Section 3.3.

3.1 Quantitative Model

3.1.1 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 enrollment. 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 online appendix 3.3.

¹³Sometimes such formulas are labeled as sufficient statistics formulas. See Kleven (2018) for a discussion on the terminology in the literature.

3.1.2 College Problem

We now consider the 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:

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

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 match differences in college-going rates across genders. We therefore parameterize the psychic cost term as

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

The term $\zeta^{\ell_t^E}$ is the common cost of working ℓ_t^E hours in college.¹⁴ Together, $\zeta^{\ell_t^E}$ and κ_X capture the deterministic cost of continuing in college and working ℓ_t^E , while the term $\varepsilon_t^{\ell_t^E}$ is the idiosyncratic and i.i.d. utility associated with staying in college and working ℓ_t^E . 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 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 $V_t^{E, \ell_t^E} \left(X, I, a_t, \varepsilon_t^{\ell_t^E} \right)$. For brevity, we do so in Appendix A.1, 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:

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

where ε_t is a vector of choice-specific 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 $V_t^D \left(X, I, a_t, \varepsilon_t^D \right) = \mathbb{E} \left[V_t^W \left(X, I, e = D, a_t, w_t \right) \right] - \delta + \varepsilon_t^D$. As in Section 2.1, an agent's problem at the beginning of the period is to choose whether or not to drop out: $V_t^E \left(X, I, a_t, \varepsilon_t \right) = \max \left\{ V_t^D \left(X, I, a_t, \varepsilon_t^D \right), V_t^{ND} \left(X, I, a_t, \varepsilon_t \right) \right\}$.

¹⁴We normalize $\zeta^0 = 0$ w.l.o.g.

At the beginning of the model, agents must decide whether to enter college or to enter the labor market directly. Let $v(X) = \varepsilon^E$ represent idiosyncratic taste for college enrollment 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 , which is a flow cost incurred every year the agent is enrolled. We assume that ε^E has a logistic distribution 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 $\tilde{X} = (\theta, s, \text{ParEdu}, \text{Region})$.

3.1.3 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.¹⁵ 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 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.¹⁶ For all years when the transfer is not given the parent simply chooses how much to consume and save.¹⁷ The parent's Bellman equation and details on the calibration of life-cycle parental earnings are given in Appendix A.2.¹⁸ In the main body, we only elaborate on the portion of the utility function that arises due to transfers.

¹⁵Note 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.

¹⁶If 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.

¹⁷The 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 3.3, our model performs well regarding the dynamics of dropout and graduation.

¹⁸We assume that parents exogenously provide transfers to the agent's siblings as well.

In the year of the transfer, the parent receives utility from transfers. Let $F(tr^H, tr^E, \tilde{X}, I)$ 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) .

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

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

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

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.

3.2 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.¹⁹ 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 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 academic year if they report being enrolled in college for at least six months. 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 3.2.1. Second, we calibrate current US tax and college policies, which we document in online appendix 3.1. Third, we estimate the parameters of the wage function, which we document in online appendix 3.3. Fourth, we estimate the parameters of the child’s and parent’s utility via maximum likelihood in Section 3.2.2.

3.2.1 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.²⁰ We make the assumption that students can only borrow through the public loan system. In

¹⁹We 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).

²⁰See 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 2)
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

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 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 level and choose $Pr_t^{Grad}(\theta)$ 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.

3.2.2 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, $\sigma^{e^{tr}}$. 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 online appendix 3.4.

In online appendix 3.5, we provide a formal discussion of the identification of psychic costs and give necessary conditions for non-parametric identification.²¹ In essence, variation in the conditional choice probabilities of working in college and dropping out for agents in their final year before graduating identifies the *differences* in psychic cost terms between the choices in college. For example, all else equal, a large number of agents choosing to work full-time in the terminal period of college would imply that the psychic costs of working full-time in college are low relative to the psychic costs of other options for college enrollees. Further, the probability of dropout in the penultimate year of college identifies the *levels* of psychic costs terms.²² All else equal, a higher number of agents dropping out in the penultimate year of college implies a higher level of psychic costs for college enrollees. How these conditional choice probabilities vary across agents with different levels of assets, parental transfers, and effective costs of college identify the idiosyncratic components of psychic costs. As in all discrete choice models, we need to make a normalization for scale. Scale here is set by the flow utility function from consumption of the numeraire good. Specifically, as there is no coefficient in front of this flow utility from consumption of the numeraire good ($\frac{e^{1-\gamma}}{1-\gamma}$) and because the price of the consumption good is normalized to one, the scale of the utility function is measured in the flow utility of consumption, where consumption is measured in dollars.

The maximum likelihood estimates are shown in Table 2 in Appendix A.3. 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. 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.

²¹See also Heckman and Navarro (2007), Cunha et al. (2005), and Navarro and Zhou (2017) for discussions on the identification of psychic costs in dynamic discrete choice models.

²²Identification of the level of utility, and not just the differences, is possible because college dropout is a *terminating action*. See Bajari et al. (2016) for a proof of identification of levels of utility in dynamic discrete choice models with a terminating action.

3.3 Model Performance and Relation to Empirical Evidence

3.3.1 Model Fit

Enrollment. Figure 1 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 online appendix 4.1, we show that the fit is equally good for graduation rates and when we examine enrollment rates separately by gender.

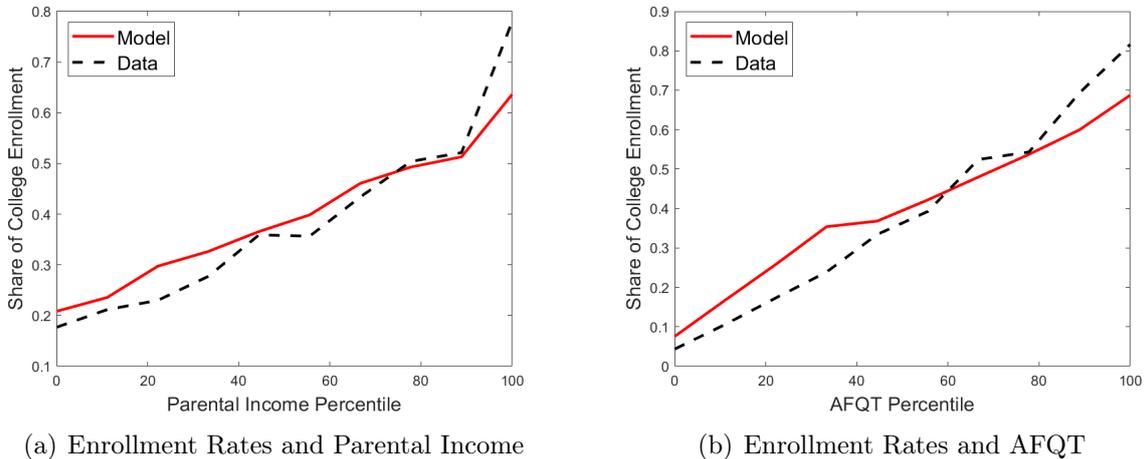


Figure 1: Enrollment Rates

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.

Parental Transfers. Differences in parental transfers across parental income levels can play a role in generating differential college-going rates across income groups. 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.²³ 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 school transfers are

²³We plot the fit of our model with respect to parental transfers in online appendix 4.3.

increasing globally in parental income, parental transfers for high school graduates in the data are decreasing for the highest-income children.²⁴

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.²⁵ We observe a weak negative relationship between parental income and working during college in the model and the data.

Earnings and College Premia. Online appendix 4.5 analyzes the performance of the model with respect to earnings dynamics. The simulated mean earnings across ages are very close to those in the data. 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 online appendix 4.6.

4 Results: Optimal Financial Aid

4.1 Baseline

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.²⁶ Additionally, we work under the constraint that financial aid is nonnegative everywhere.²⁷ Figure 2(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

²⁴A 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.

²⁵Note that average hours of work are calculated using data from the entire year and thus include work during summer break.

²⁶At 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

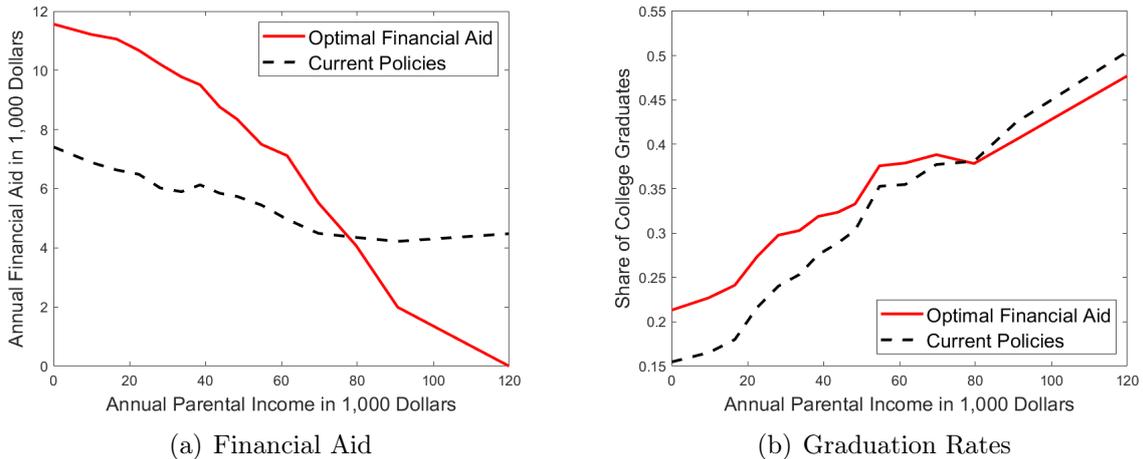


Figure 2: Optimal versus Current Financial Aid

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.

in financial aid policies is mirrored in the change in college graduation, as shown in Figure 2(b). The total graduation rate increases by 2.8 percentage points to 32.8%.

4.2 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.²⁸ 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.²⁹ The results are in shown with the blue dashed-dotted line in Figure 3(a). The optimal financial aid schedule is slightly less progressive than the optimal financial with a Utilitarian welfare function.

online appendix 6.7, we show that our main result also extends to the case in which the merit-based elements are chosen optimally. We outline the algorithm we use to calculate optimal policies in online appendix 7.

²⁷Relaxing this, one would get a negative subsidy at high parental income levels but nothing substantial changes in terms of results.

²⁸In fact, $1 - W^E(I)$, 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 $1 - W^E(I)$ 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.

²⁹We outline the algorithm we use to calculate optimal policies with no desire for redistribution in online appendix 7.2.

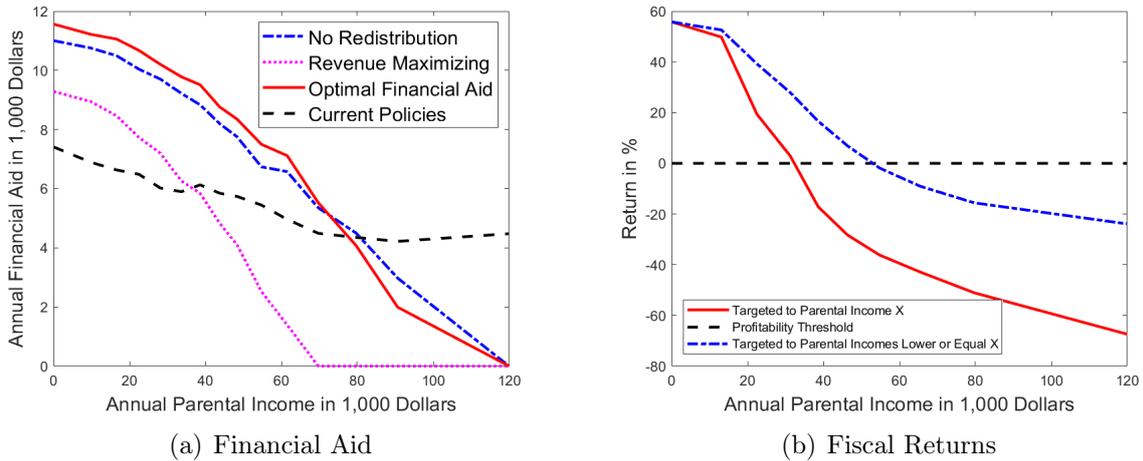


Figure 3: Financial Aid Policies with no Redistribution Motive and Fiscal Returns on Increase in Financial Aid

Notes: The dashed-dotted (blue) line shows the optimal schedule for a social planner with no redistribution motive. The dotted (magenta) 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), 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 . 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 .

4.3 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? The dotted magenta line in figure 3(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 4.4.

4.4 Self-Financing Reforms

An increase in financial aid can be self-financing if properly targeted. The solid red line in Figure 3(b) 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 3(b). 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.

5 Why Are Optimal Policies Progressive?

We have just shown in Section 4 that optimal financial aid is considerably more progressive than the current US policies and 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)

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

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.³⁰ Figure 4(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

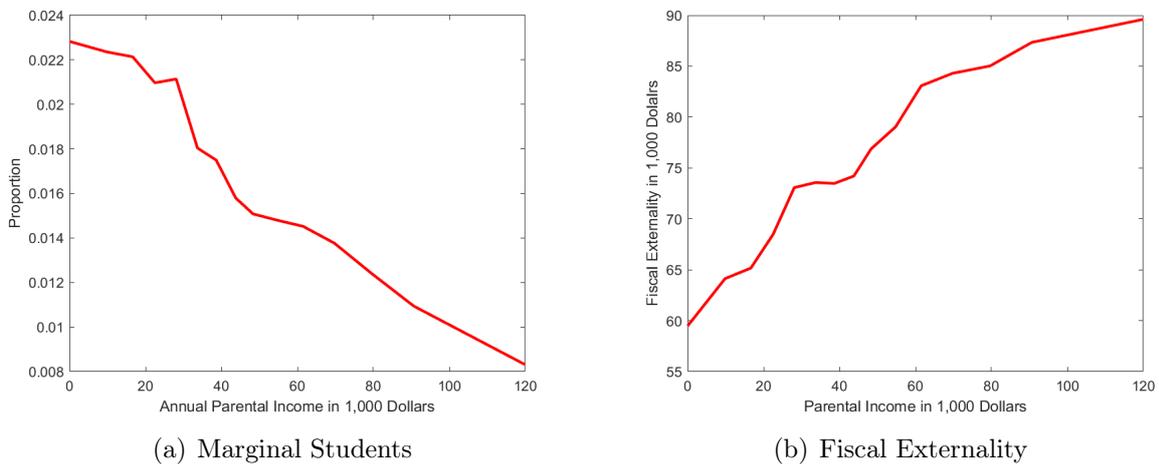


Figure 4: Marginal Students and Tax Revenue Changes

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.

³⁰We only found a quantitatively very small contribution of the completion effect and therefore focus on the other two effects.

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, as shown in Figure 4(b), the fiscal externality $\Delta\mathcal{T}^E(I)$ is increasing in parental income because (i) marginal enrollees from higher income households have higher returns³¹ 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 4.2 and 4.3 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 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.

5.1 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 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 provide a model-based decomposition by removing model features one-by-one to better understand which factors drive these two relationships. For example, we will simulate a version of the model in which we remove the correlation between parental income and ability and recalculate the shares of inframarginal students and the shares of marginal students. This will help us to isolate the effect of these features of the model on the increasing share of inframarginal students and the decreasing share of marginal students. 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.³²

³¹The 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.

³²In online appendix 5.2, we consider an alternative decomposition in which we equalize parental transfers first. Additionally, in online appendix 5.3, we conduct a decomposition where we remove the borrowing constraint before we equalize parental transfers.

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. For each decomposition, 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 online appendix 5.1.

Inframarginal Students We start by focusing on the relationship between parental income and college enrollment in Figure 5(a).³³ 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.³⁴ One factor that leads to this positive relationship is the correlation between parental income and ability. To understand the contribution of this 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 5(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.³⁵ 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

³³Alternatively, we could have focused on graduation instead of enrollment. The implications are very similar.

³⁴Note that this relationship is stronger than the one in Figure 1(a) where the current financial aid schedule is used as benchmark.

³⁵Furthermore, we set κ_0 so that the average psychic cost of going to college is unchanged.

draw. Removing the riskiness of college leads to a further flattening of the relationship between parental income and college enrollment. However, 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.³⁶ Differences in parental transfers therefore play an important role in explaining differences in college going rates, especially given that agents face borrowing constraints. We conclude that all components play an important role for the increasing share of inframarginal students with the exception of the risk channel.

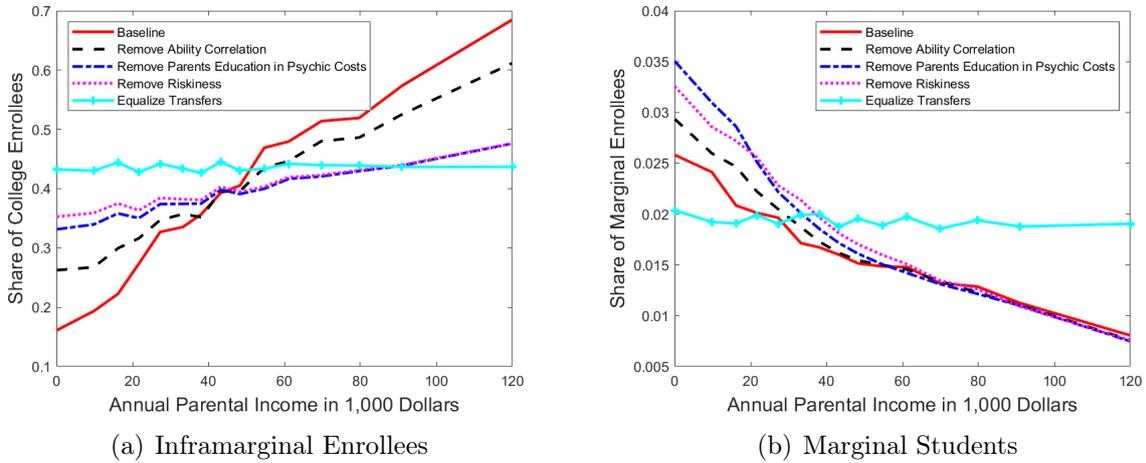


Figure 5: Model-Based Decomposition for Marginal and Inframarginal Students

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 remove additionally the correlation between the psychic costs and parental education. For the dotted pink line we additionally remove labor market riskiness; i.e. education decisions are made with no uncertainty about future wages. For the turquoise line with crosses we set parental transfers to the mean parental transfers in the data, conditional on education.

Marginal Students We turn to the determinants of the negative relation between parental income and the share of marginal enrollees (again considering a \$1,000 increase in financial aid) in Figure 5(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 parental income and ability and in which we remove the relation between

³⁶The line is not totally flat because individuals still differ in gender and region and these variables are not distributed in exactly the same way for each parental income group.

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 that the enrollment share of low income families moves closer to 50%: 50% corresponds to the mode of the logistic 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. This last step shows the differences in parental transfers play the most important role in generating the negative gradient between parental income and share of marginal students. Households with lower parental transfers are closer to the borrowing constraints and therefore more responsive to increases in financial aid.³⁷

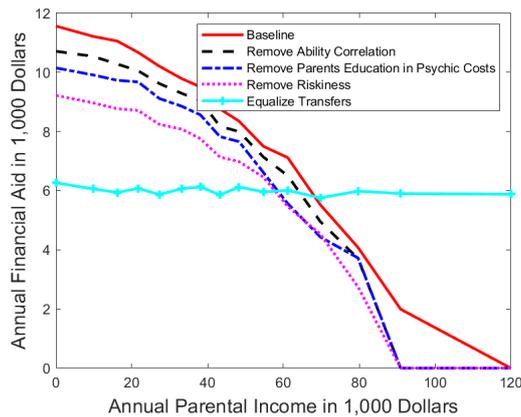


Figure 6: Optimal Financial Aid for Different Model Specifications

Notes: For each model specification (see Figure 5), we illustrate the respective optimal financial aid schedule.

5.2 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 share of marginal students. To understand what these factors imply for the optimal financial aid schedule, we therefore now remove model features one-by-one and recalculate the optimal financial aid schedule. For example, we start by calculating the optimal financial aid in an environment in which we remove the correlation between parental income and ability. This

³⁷As said, in online appendix 5.3, we consider a decomposition where we remove the borrowing constraint before we equalize parental transfers. In the absence of borrowing constraints, parental transfers play a much smaller roll in differences in share of marginal students across income groups.

allows us to measure the role of the correlation between parental income and child’s ability on the optimal aid schedule. For each alternative model specification, we recalculate the government’s exogenous revenue requirements, \bar{F} , by simulating the model with the current financial aid schedule under the alternative model specification. We then calculate the optimal financial aid schedule given this new level of the exogenous revenue requirements. Figure 6 shows the implied optimal policies for each model specification.

The red line shows the optimal aid schedule in the baseline model. The next two lines show the optimal aid when we remove the correlation between parental income and ability and then additionally remove the correlation between parental income and psychic costs. The optimal policies are still relatively progressive in those models. Although enrollment rates flatten (see the lines in Figure 5(a)), there are offsetting effects as low-income children are now more likely to be marginal (see the lines in Figure 5(b)). Therefore, the positive relation of parental income with ability and preferences for college are not the main drivers of the financial aid results.

The next line shows the optimal aid in which we additionally remove riskiness. As shown in the previous section, both the inframarginal and marginal relationships become flatter when we remove the riskiness of college. Therefore removing riskiness leads to a more progressive optimal aid schedule, because there are no offsetting effects in this case. Finally, the turquoise crossed line shows the optimal aid when we have equalized parental transfers. This line shows an almost zero slope, as transfers are equalized and we are in a world where parental income plays no more role. Given that agents face borrowing constraints, differences in parental transfers play an important role in determining the relationship between parental income and share of marginal students. The effects for marginal students and inframarginal students therefore again work in tandem, pushing towards flat aid.³⁸

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. In the presence of borrowing constraints, 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.³⁹

³⁸Again, the fact that financial aid 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.

³⁹As mentioned, in online appendix 5.3, 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 the decomposition where we remove borrowing constraints first, before removing parental transfers, the role of parental transfers is much more limited.

6 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 online appendix 6, because of space constraints.

6.1 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 4.3. 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). As in Section 5.2, we first recalculate the government's exogenous revenue requirements, \bar{F} , by simulating the model with the current financial aid schedule but without borrowing constraints. We then calculate the optimal financial aid schedule given this new level of the exogenous revenue requirements. As we show in online appendix 6.1, 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 online appendix 6.2.

6.2 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. We now 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 based on the translog functional form employed in Agostinelli and Wiswall (2016)⁴⁰

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

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 parental transfers 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 online appendix 6.3. We also show that our model is consistent with quasi-experimental estimates of the effect of family income on a child’s ability (Dahl and Lochner, 2012).

The optimal financial aid schedule and graduation rates with endogenous ability are shown in Figure 7. Panel 7(a) shows the new optimal financial aid schedule when ability is endogenous. Compared to the baseline case when ability is exogenous, the 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 7(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.⁴¹

⁴⁰Agostinelli 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.

⁴¹We show childhood ability as a function of parental income under the current financial aid schedule and under the optimal schedule in online appendix 6.3. Low income parents may face borrowing constraints when their child is young (Caucutt and Lochner, 2017; Lee and Seshadri, 2019). In online appendix 6.4, we calculate the optimal aid schedule with endogenous ability with low income parents that may be borrowing constrained when their child is young. The optimal progressivity decreases as we tighten borrowing constraints but the remains more progressive than the current schedule in all cases.

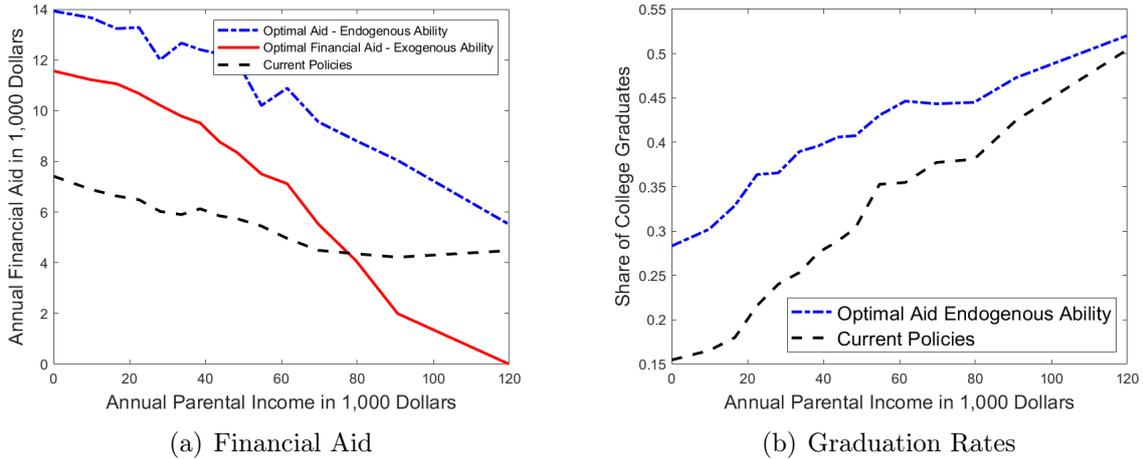


Figure 7: Financial Aid and Graduation with Endogenous Ability

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.

7 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):⁴² 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. 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.

⁴²Diamond 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)."

A Appendix

A.1 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.

$$V_t^{E, \ell_t^E} (X, I, a_t, \varepsilon_t^E) = \max_{c_t} [U^E (c_t, \ell_t^E, X, \varepsilon_t^E) + \beta \times \{ (1 - Pr_t^{Grad}(\theta)) \mathbb{E} [V_{t+1}^E (X, I, a_{t+1}, \varepsilon_{t+1})] + Pr_t^{Grad}(\theta) \mathbb{E} [V_{t+1}^W (X, e = G, a_{t+1}, w_{t+1})] \}] \quad (5)$$

subject to the budget constraint and the borrowing constraint, where $V_{t+1}^E (X, I, a_{t+1}, \varepsilon_{t+1})$ and ε_{t+1} 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 $t + 1$. We allow tuition, $\mathcal{F}(X)$, 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.

A.2 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.⁴³ 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

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

subject to:

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

where a_t^P is the parent's assets in year t and y_{it}^P is the parent's income in year t .⁴⁴ 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

⁴³This 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.

⁴⁴We 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.

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

subject to:

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

where tr^{hs} and tr^{col} are the transfers offered conditional on the child's education choice, and tr^e are the realized transfers.⁴⁵ 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. $F(tr^{hs}, tr^{col}, \tilde{X}, I)$ is 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.

A.3 Maximum Likelihood Estimates

The parameter estimates are contained in Table 2.

⁴⁵In 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.

Table 2: 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.

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*.
- 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.
- BAJARI, P., C. S. CHU, D. NEKIPELOV, AND M. PARK (2016): “Identification and semi-parametric estimation of a finite horizon dynamic discrete choice model with a terminating action,” *Quantitative Marketing and Economics*, 14, 271–323.
- 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.
- 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.
- (2018): “Education policies and taxation without commitment,” *The Scandinavian Journal of Economics*, 120, 1075–1099.
- 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. 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.

- HENDREN, N. (2016): “The Policy Elasticity,” in *Tax Policy and the Economy, Volume 30*, Chicago: University of Chicago Press.
- 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.
- JOHNSON, M. T. (2013): “Borrowing Constraints, College Enrollment, and Delayed Entry,” *Journal of Labor Economics*, 31, 669–725.
- 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.
- 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.
- NAVARRO, S. AND J. ZHOU (2017): “Identifying Agent’s Information Sets: an Application to a Lifecycle Model of Schooling, Consumption, and Labor Supply,” *Review of Economic Dynamics*, 25, 58–92.
- 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*.
- SAEZ, E. AND S. STANTCHEVA (2016): “Generalized Social Marginal Welfare Weights for Optimal Tax Theory,” *American Economic Review*, 106, 24–45.
- 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.