

The Equilibrium Effects of Subsidized Student Loans*

Cauê Dobbin[§]
(Job Market Paper)

Nano Barahona[†]

Sebastián Otero[‡]

July 4, 2022

[\[Link to the latest version\]](#)

Abstract

We investigate the equilibrium effects of subsidized student loans on tuition costs, enrollment, and student welfare. Two opposing forces make the impact on tuition theoretically ambiguous. First, students with loans become less price-sensitive because they do not bear the total tuition cost, causing tuition to rise (*direct effect*). Second, loan programs tend to increase the market share of more price-sensitive students, reducing tuition (*composition effect*). We develop a model of the supply and demand for higher education and estimate it leveraging a large change in the availability of student loans in Brazil. We find that Brazil's current loan program raises prices by 1.6% and enrollment by 11% relative to a counterfactual without loans. We decompose the price effect into its direct (2.7% increase) and composition (1.1% decrease) components. Finally, we show that an alternative policy that gives loans only to low-income students raises enrollment by 16% relative to a counterfactual without loans. Most of the difference in enrollment between the two policies are due to price reductions coming from a stronger composition effect in the alternative policy.

Keywords: Education, student loans, equilibrium effects, private colleges

JEL Codes: H52 , H22, I22, I23, I24, G59, L11

*We are indebted to Caroline Hoxby, Rebecca Diamond, Melanie Morten, Isaac Sorokin, and Constantine Yannelis for invaluable guidance and support on this project. We also thank Claudia Allende, Luis Armona, Arun Chandrasekhar, Marcel Fafchamps, Neale Mahoney, Liran Einav, Pascaline Dupas, Matthew Gentzkow, and seminar participants at Stanford University for valuable comments and suggestions. We gratefully acknowledge financial support from the Stanford King Center on Global Development, the Stanford Center for Computational Social Sciences, the Stanford Institute for Economic Policy Research (SIEPR), the Haley-Shaw Dissertation Fellowship, and the NAEd/Spencer Dissertation Fellowship Program. All remaining errors are our own. [§]Stanford University. Email: dobbin@stanford.edu. [†]University of California, Berkeley. Email: nanobk@berkeley.edu. [‡]Stanford University. Email: sotero@stanford.edu.

1. Introduction

Governments worldwide offer subsidized loans to increase low-income students' access to higher education: in OECD countries, 10% of the public expenditure on higher education is on student loans (OECD, 2014), and several Latin-American countries have government-funded student loan programs (Marta Ferreyra et al., 2017). However, despite the popularity of these programs, policymakers have long been worried that they enable colleges to raise tuition costs and capture a large share of the invested public funds, undermining the policy's effectiveness (Bennett, 1987).

Conceptually, an expansion of a student-loan program has two opposite effects. On the one hand, students with loans become less price-elastic because tuition does not pass through entirely to them, leading to higher markups and prices. We refer to this mechanism as the *direct effect*. On the other hand, loan programs usually target low-income students, increasing their market share. Since low-income students are more price elastic in our and many settings, the average price elasticity of the market increases, reducing markups and prices. We refer to this mechanism as the *composition effect*. These opposing forces imply that the net impact on prices is ambiguous and depends on how the government targets loans, a fact the previous literature has not discussed (Long, 2004; Singell and Stone, 2007; Turner, 2012; Cellini and Goldin, 2014; Turner, 2017; Lucca et al., 2018; Kelchen, 2019).

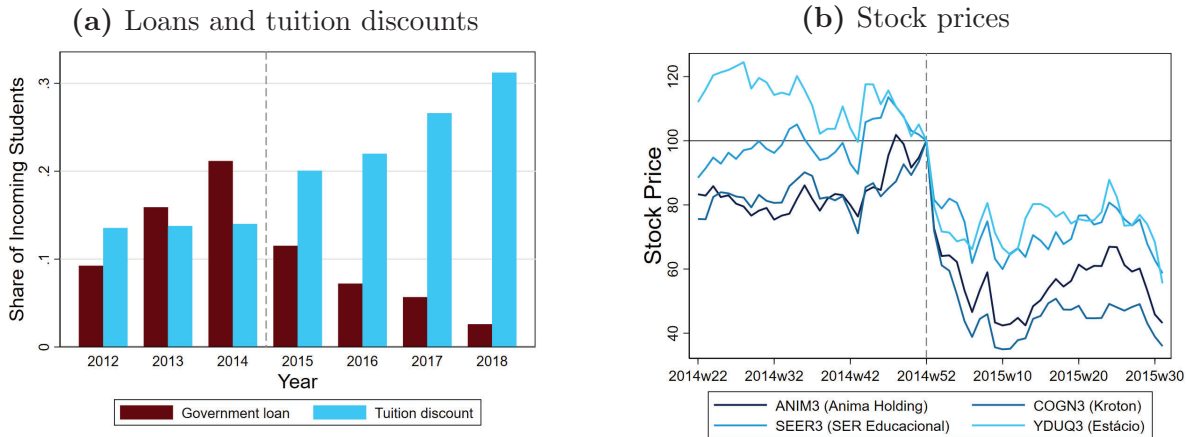
In this paper, we investigate the equilibrium effects of subsidized student loans. First, we develop a model of the supply and demand for higher education and highlight the key parameters governing the strength of the composition and direct effects. We then estimate these parameters empirically and show that the net effect of loans on prices depends on how loans are targeted. Finally, we use our estimated model to compare the outcomes of alternative policy designs in terms of enrollment and student welfare.

We show that the magnitudes of the direct and composition effects depend on three key parameters. First, the effect of loans on students' price sensitivity: if students who receive loans become much less sensitive to prices, the direct effect is strong. Second, the heterogeneity in price elasticity across students: the larger the difference in price elasticity between high- and low-income, the stronger the composition effect. Third, the accuracy of price discrimination: in the extreme, if price discrimination is perfect, colleges charge one individualized price to each student, with no composition effect.

We investigate the empirical relevance of these forces in the context of the Brazilian higher education market. We exploit a policy change that resulted in a drastic reduction in the availability of loans, and document how students and colleges reacted. In 2014, 21% of incoming students received federal loans; by 2018, only 3% did. Consequently, the number of tuition discounts soared: the share of incoming students with discounts increased from 14% to 31% between 2014 and 2018 (Figure 1a). Colleges with a large pre-policy share of students with loans drove the increase in discounts. The

reform in the federal loan program severely impacted private colleges and led to a massive drop in their stock prices (Figure 1b).

Figure 1: The 2015 reform in the federal loan program



Notes: Panel (a) shows the shares of incoming students in tuition-charging institutions receiving a federal loan (FIES) or a tuition discount each year. These two forms of aid are not mutually exclusive. Source: Census of Higher Education and FIES administrative records. Panel (b) shows the stock prices of four higher-education conglomerates that receive 30% of the students with federal loans. Stock prices are normalized to 100 on the day before the policy change was announced. In both panels, the vertical line marks the announcement of the change in FIES’s rules.

The Brazilian higher education market has several features that help identify the key parameters governing the equilibrium impacts of student loans. First, we observe yearly-level enrollment decisions and tuition costs, as well as individual-level income, which we use to estimate the difference in price elasticity between high- and low-income students. Second, we observe which students receive tuition discounts, and use this information to assess how colleges price discriminate. Third, the government allocates loans to students through a centralized mechanism with clear assignment rules based on eligibility thresholds. We leverage the discontinuity in loan availability at these thresholds as a natural experiment to estimate how loans change students’ price elasticities.

Guided by the trends observed in Figure 1, we develop a model of the supply and demand for higher education. The goals of the model are twofold. First, to put together the different pieces of empirical evidence and estimate the net effect of student loans on prices, considering both the direct and composition effects. Second, to compare the outcomes of alternative policy designs.

In our model, the market is composed of three actors. First, the government allocates loans to students. Second, on the supply side, colleges choose prices to maximize profits, considering how loans are distributed. Third, on the demand side, students make enrollment decisions after observing prices and loan availability. Next, we further detail each of these actors.

On the government side, the government allocates loans using two policy levers. First, the government sets degree-specific minimum scores on ENEM (the Brazilian standardized exam) necessary to receive a loan. Second, the government sets socioeconomic eligibility criteria to participate, and

non-eligible students cannot receive a loan to enroll in any degree, even if their score is high enough.

On the supply side, colleges are profit-maximizing multiproduct firms that offer a fixed set of degrees. They price discriminate between students by charging a full and a discounted price for each of their degrees. The first-order conditions of colleges' problem shed light on how loans affect equilibrium prices. The full price equals the marginal cost plus a markup, which is approximately given by the inverse of the average price elasticity of students paying the full price. The analogous is valid for the discounted price. Hence, an increase in the availability of loans affects prices in two ways. First, the direct effect: students with loans become less price elastic, raising markups and prices. Second, the composition effect: loan programs targeted at low-income students increase the market share of such students, who are more price elastic, increasing the average elasticity of the market, and thus reducing markups and prices.

This framework also highlights the role played by price discrimination. If discrimination were perfect, only high-income students would pay the full price, and only low-income would pay the discounted price. Therefore, since there are separate first-order conditions for the full and the discounted prices, the relative market share of high- and low-income students would not matter for price-setting, and there would be no composition effect.

On the demand side, students choose a degree, or the outside option, to maximize utility. Students' sensitivity to prices depends on whether they have a loan, representing several mechanisms through which loans change enrollment decisions. First, government loans alleviate credit constraints. Second, the interest rate is subsidized. Third, default rates are high; hence, the effective price paid by students with loans is lower than the actual price.

We estimate our structural model using data from the Brazilian higher education market, with pre- and post-policy data. We pay special attention to the three key parameters governing the magnitudes of the direct and composition effects: the difference in price elasticity between high- and low-income, the impact of loans on price elasticities, and the accuracy of price discrimination.

To estimate price elasticities, we rely on the variation of prices across time and degrees. To address the potential endogeneity of prices, we build an instrument that exploits the panel structure of the data and the ownership relations of multiregional college chains. This strategy leverages the fact that prices of multiregional firms are often highly correlated across regions due to managerial inertia. We find that low-income students are significantly more price elastic than high-income students, with a median elasticity of -5.5 for the low-income and -1.4 for the higher income.

To estimate the effects of loans on price elasticities, we leverage quasi-random variation created by the centralized mechanism that allocates loans to students. The mechanism determines a minimum score to receive a loan in each degree; hence, access to loans varies discontinuously at these thresholds. We find significant enrollment discontinuities at the thresholds and estimate that receiving a loan reduces price elasticities by 2.7 and 0.9 for low- and high-income students, respectively.

To estimate the accuracy of price discrimination, we leverage the granularity of our data, which

contains individual-level information on student income and tuition discounts. We find that low-income individuals are more likely to receive discounts, but targeting is far from perfect: 28% of the students in the top income ventile receive tuition discounts.

The main takeaways from the estimated parameters are the following. On the one hand, the difference in price elasticity between high- and low-income students is large, and colleges have limited ability to price discriminate, which implies that the composition effect is strong. On the other hand, loans substantially reduce price elasticities, indicating that the direct effect is also strong. Since both forces are relevant in this market, we rely on the structure of the model to compute their net effect under different policy designs.

We start by evaluating how well our model predicts the equilibrium outcomes of different policies, leveraging the 2015 reduction in loan availability. More specifically, we use our estimated structural model to predict what would have been the equilibrium in 2016 if the rules of the program had not changed. Our prediction indicates that if the rules had not changed, total enrollment, number of loans, and number of discounts would have followed the same trend they had from 2011 to 2014. Notice that the model was estimated using only data from 2014 and 2016; hence, previous trends were not among the targeted moments. Therefore, these results indicate that the model is well suited to predicting the equilibrium outcomes of different policy designs.

We then use the model to estimate the equilibrium effects of the current Brazilian student loan policy and decompose these impacts into three components: partial equilibrium responses, direct price effects, and composition price effects. In partial equilibrium, the loan program increases total enrollment by 17.0%, relative to a counterfactual without loans. The direct effect raises prices by 2.7%, on average, and, as a result, reduces enrollment by 10.4%. Finally, the composition effect reduces prices by 1.1% and, as a result, raises enrollment by 4.9%. The net result of these three forces is a 1.6% price increase and an 11.5% increase in total enrollment. In summary, we find that both the direct and the composition effects are responsible for prices responses that lead to substantial changes in enrollment. In particular, price reductions induced by the composition effect are responsible for 40% of the enrollment gains of the current Brazilian loan program.

Next, we build upon our previous findings to propose an alternative allocation of loans: giving loans only to low-income students. The reason for focusing on this specific policy is that low-income students are more price-sensitive than high-income ones. Hence targeting loans to low-income students can strengthen the composition effect and result in lower prices, leading to higher enrollment. Indeed, we show that an alternative policy that keeps the same budget but gives loans only to low-income students raises prices by 1.2% and enrollment by 16%. That is, the alternative policy increases enrollment by 40% more than the current one. Most of the difference in enrollment between the two policies are due to price reductions coming from a stronger composition effect in the alternative policy.

To evaluate the effects of loan programs beyond changes in total enrollment, we compute consumer

surplus under different policy designs. We calculate consumer surplus under a range of assumptions, and we find that the low-income-only alternative policy always results in higher consumer surplus than the current policy. The reason is that the alternative policy lowers prices. For example, under our baseline assumption, each \$1.00 invested in the current policy induces supply-side responses that decrease total consumer surplus by \$0.30. In contrast, supply-side responses would *increase* consumer surplus by \$0.09 per \$1.00 invested in a low-income-only alternative loan program. The difference in consumer surplus between the two policies is entirely due to price reductions coming from a stronger composition effect in the alternative policy.

We then show how these results depend on the extent of price discrimination. For this purpose, we show how the current and the alternative policies compare under perfect price discrimination. There is no composition effect in this scenario, and all price changes come from the direct effect. Therefore, under perfect price discrimination, the enrollment gains obtained from the alternative policy compared to the current one are much more modest. These patterns highlight that the outcomes of different policy designs depend on the market structure. When colleges have limited ability to discriminate, policies that increase the market-shares of low-income students reduce prices, leading to substantial enrollment gains. However, this is not the case under perfect price discrimination because there is no composition effect.

Finally, we discuss how governments could improve loan targeting in practice. Perfectly targeting low-income students might be challenging due to misreporting and fraud. Hence we simulate the outcomes of a feasible alternative: giving loans only to public high school students since they are, on average, poorer.¹ The patterns are very similar to those obtained giving loans only to low-income students, but the magnitudes are smaller. For example, targeting public high schools reaches 28% of the enrollment increase of perfectly targeting low-income students.

In summary, our results show that the price changes induced by loan programs have substantial consequences for enrollment and student welfare. Moreover, the magnitude and direction of the effects depend on how the government targets loans because of the composition effect. In particular, in the context of the Brazilian higher education market, if the government gave loans only to low-income students, average tuition would go down, resulting in large gains in enrollment and student welfare.

This paper builds upon several strands of the previous literature. It adds to a large body of research that studies how access to financial aid changes students' decisions of whether to enroll in higher education and in which degree (van der Klaauw, 2002; Hoxby, 2004; Angrist et al., 2014; Fack and Grenet, 2015; Londono-Velez and Rodriguez, 2020). Most related to our work is Solis (2017). He uses a regression discontinuity strategy—similar to ours—to estimate the effect of loan availability on enrollment in Chile. We advance his work in two ways. First, in our setting, loan eligibility cutoffs are degree-specific and are spread across the whole score distribution. Hence, we can estimate the

¹In our sample, 82% of public high school students are low income, compared to 42% in private high schools.

effect of loan availability for a much broader group of students. Second, we combine our reduced form results with a structural choice model to assess the impact of loans on price elasticities, which is a crucial input to understanding the equilibrium effects of loan programs.

Our work is also closely related to the literature on the supply-side responses to student financial aid (Long, 2004; Singell and Stone, 2007; Turner, 2012; Cellini and Goldin, 2014; Turner, 2017; Lucca et al., 2018; Kelchen, 2019; Eaton et al., 2020; Dobbin et al., 2021). Our contributions to this literature are two-fold. First, we show that student loans affect prices in two ways, the direct and composition effects and that the net result is ambiguous and depends on how loans are targeted. Second, we use our empirical findings to propose an alternative allocation of loans that achieves higher enrollment and student welfare.

We also contribute to the literature on price discrimination in higher education (Tiffany, 1998; Zerkle, 2006; Epple et al., 2006; Fillmore, 2016; Epple et al., 2019). We advance previous work in two ways. First, we investigate the interaction between price discrimination and student loans. Second, our study is in Brazil, whereas most previous work has focused on the United States, leading to major differences in how we model price discrimination. On the one hand, American colleges choose a personalized price for each student and explicitly target discounts based on information provided by students in their application and financial aid forms. On the other hand, Brazilian colleges have access to much less information. Hence, discounts are targeted through unexpected seasonal promotions and by posting different prices on different online platforms so that only certain students find the discounts. We develop a novel model of price discrimination that, while being computationally tractable, captures the main empirical patterns we observe in the Brazilian higher education market. Our setting is related to the model of discounts and consumer search in eBay, developed in Coey et al. (2020).

Finally, this paper is part of a growing body of research that studies educational markets through the lens of structural models (Ferreira and Kosenok, 2018; Bau, 2019; Neilson, 2019; Singleton, 2019; Allende, 2020; Neilson, 2021; Dinerstein and Smith, 2021; Armona and Cao, 2021). Our work builds upon this literature to understand the equilibrium effects of student loans and propose a more efficient allocation of these loans.

The remainder of the paper is organized as follows. Section 2 presents a simplified conceptual framework. Section 3 describes the setting and the data. In Section 4, we provide descriptive evidence on the moments that identify key parameters governing the equilibrium responses to subsidized loans. In Section 5, we introduce our structural model of the supply and demand for higher education. In Section 6, we estimate the model. In Section 7, we present our main counterfactual exercises and discuss policy implications. Section 8 concludes.

2. Conceptual framework

This section presents a simplified theoretical framework describing our main mechanisms. The framework has two goals. First, to provide a formal definition of the direct and composition effects. Second, to uncover the parameters determining how strong each effect is.

Consider a market with only one college (a single-product monopolist) charging price p and with marginal cost c . Students are divided into groups of consumer types x , and \mathcal{X} is the set of consumer types. Let N_x be the size of group x and $s_x(p, a_x)$ the probability that a student of group x enrolls in college. Finally, a_x is a continuous variable denoting the generosity of financial aid for group x . It can represent, for example, the value of a voucher given to individuals in this group or a percent subsidy on tuition payments. In our setting, a_x will represent the probability of a student from group x receiving a loan.

The problem of the college is given by:

$$p^* = \arg \max_p \sum_{x \in \mathcal{X}} N_x \cdot s_x(p, a_x) \cdot (p - c).$$

To decompose the effects of student financial aid on prices into direct and composition components, we compute the impacts of a marginal increase in financial aid generosity. Let η be the price elasticity of total demand, η_x the price elasticity of demand coming from group x , and λ a measure of the curvature of the demand curve.² In Appendix B, we show that the effect of a marginal increase of financial aid for students of type x on the equilibrium price p^* is:

$$\frac{1}{p^*} \frac{dp^*}{da_x} = \overbrace{\left(\frac{N_x s_x}{\sum_{\tilde{x} \in \mathcal{X}} N_{\tilde{x}} s_{\tilde{x}}} \right)}^{\text{scale}} \cdot \underbrace{\frac{1}{\eta^2}}_{\text{market power}} \cdot \overbrace{\frac{1}{2 - \lambda}}^{\text{curvature}} \cdot \left[\underbrace{\left(-\frac{\partial \eta_x}{\partial a_x} \right)}_{\text{direct effect}} + \overbrace{\left(\eta - \eta_x \right) \cdot \frac{1}{s_h} \frac{\partial s_x}{\partial a_x}}^{\text{composition effect}} \right]. \quad (1)$$

Equation (1) has five components. The first three determine the overall effect of financial aid on prices and are similar to the previous literature on pass-through and tax incidence under imperfect competition (Weyl and Fabinger, 2013). The last two determine the relative importance of the direct and composition effects and are the main focus of this paper.

The components determining the overall price effects are the following. First, *scale* is the relative size of the group for which financial aid is being expanded. If the group is small, then price responses are small. Second, *market power* depends on the price elasticity of demand. If demand is inelastic,

² $\eta_x \equiv -\frac{p}{N_x s_x} \frac{\partial(N_x s_x)}{\partial p}$; $\eta \equiv -\frac{p}{\sum_{x \in \mathcal{X}} N_x s_x} \frac{\partial \left(\sum_{x \in \mathcal{X}} N_x s_x \right)}{\partial p}$; $\lambda \equiv \frac{\left(\sum_{x \in \mathcal{X}} N_x s_x \right) \frac{\partial^2 \left(\sum_{x \in \mathcal{X}} N_x s_x \right)}{\partial p^2}}{\left[\frac{\partial \left(\sum_{x \in \mathcal{X}} N_x s_x \right)}{\partial p} \right]^2}$.

price responses are stronger. Third, price effects are stronger when the *curvature* of the demand curve is larger.

Now let us discuss the objects determining the relative importance of the direct and composition effects. The *direct effect* depends on how much financial aid reduces the price elasticity of demand $\left(\frac{\partial \eta_x}{\partial a_x}\right)$. The larger the reduction, the more the direct effect increases prices. The *composition effect* depends on two elements. First, how much financial aid increases demand for education $\left(\frac{1}{s_h} \frac{\partial s_x}{\partial a_x}\right)$. The larger the increase, the stronger the composition effect. Second, the relative price elasticity of the group receiving financial aid $(\eta - \eta_x)$. If the group is more price elastic than average, the composition effect reduces prices and vice versa. The reason is that increasing the demand of a group that is more price elastic than average increases the price elasticity of the overall demand curve.

In Appendix B.2, we introduce price discrimination to this model. The main takeaway is that price discrimination weakens the composition effect. The intuition is the following. With price discrimination, the composition effect depends on the price elasticity of the group receiving financial aid relative to the overall price elasticity of *individuals paying the same price as the targeted group*. Consider an extreme example in which the college can perfectly discriminate between consumer types, that is, each group pays a different price p_x . In this case, the elasticity of overall demand with respect to p_x is equal to the price elasticity of group x . Hence, there is no composition effect.

3. Background

3.1 Higher education in Brazil: Overview

The Brazilian higher education market is composed of both public and private institutions. In 1997, a series of regulations facilitated the expansion of the private sector by allowing the entrance of for-profit colleges.³ As a result, enrollment in private institutions has more than tripled in the last two decades (Appendix Figure A.1). In 2016, 80% of the 11 million students enrolled in higher education, and 87% of the 2,400 colleges were in the private sector. Despite this massive expansion, only 15% of the population between 25 to 64 years currently has a higher education degree, compared to 37% in OECD countries (OECD, 2019).

The public and private sectors operate very differently. Public colleges are tuition free and in general they are more prestigious and of higher quality. For example, public institutions have, on average, 30 students per faculty, whereas the private ones have 200. Consequently, public colleges are highly oversubscribed and selective, whereas the private ones are not. The median public degree has five applicants per vacancy, with over twenty applicants per vacancy in the 10% most selective degrees. On the other hand, over 90% of the degrees in the private sector fill fewer than 80% of their spots. Moreover, in Appendix C, we show that there is no evidence that private colleges select

³Law 9.394, of December 20, 1996; and Decree 2.207, of April 15, 1997.

students based on their score. In this context, affordability is the main obstacle to attending a private college, which was the motivation for the creation of the federal student loan program, the focus of this study.

Finally, note that in Brazil, students enroll in a specific degree, defined as a major-college combination. For example, Economics at Pythagoras University. Also, most degrees take four years, with a few majors taking longer, such as Engineering (usually five years).

3.2 ENEM: The Brazilian National Standardized Exam

The Brazilian National Standardized Exam (ENEM) takes place once a year, across the whole country, and over 4 million students take the exam every year. The exam consists in five parts: four multiple choice exams (math, languages, natural sciences, and human sciences) and one essay.

ENEM is high stakes for students and its results are the main component of several selection processes. For example, the platform that allocates spots in public universities is based on ENEM scores, students who dropped out of high school can receive a high school diploma with a score above a certain threshold, and taking ENEM is required to receive a federal scholarship to study abroad.

Most importantly, for our purposes, the centralized mechanism that allocates subsidized loans to students attending private colleges is also based on ENEM scores. In Section 3.3, we describe this allocation in more detail.

3.3 Public financial aid: The federal student loan program

In 1999, the Brazilian federal government combined several smaller student aid initiatives into the *Higher Education Finance Fund* (FIES). This program provides students enrolled in private colleges loans that cover 100% of their tuition costs. Until 2009, the program remained relatively small. In 2010, the government restructured FIES and access became virtually unrestricted. That is, both the need-based and the merit-based eligibility conditions were extraordinarily generous, and nearly all students were eligible. Consequently, the number of new contracts skyrocketed, from fewer than twenty thousand in 2009 to more than seven hundred thousand in 2014.

FIES requires large investments from the federal government for two reasons. First, interest rates are highly subsidized: in 2014, the FIES rate was 3.5% per year, whereas the one on federal bonds was 9%, and the mean rate for personal loans was 101%. Second, default is widespread: 30% of the contracts in the repayment phase were delinquent in 2014.

Hence, due to budget limitations, the program was once again restructured in 2015, but this time to limit access. After the reform, the number of new contracts plummeted, reaching around 100,000 in 2017, a 7-fold drop from the 2014 level (Figure 1a). The policy change was unexpected and had a massive impact on the private higher-education sector, causing a sharp drop in the stock prices of education conglomerates (Figure 1b).

The reform limited access to loans in two ways:

First, it imposed a maximum per capita family income of 2.5 times the federal minimum wage. However, this limit was not binding for most of the relevant population. Appendix Figure A.2 presents the income distribution among incoming students and shows that in the last year before the policy change (2014), only 5.6% of the students enrolled in private institutions and 2.5% of the ones receiving FIES loans were above this limit. Hence, the maximum income limit had little practical consequence.

Second, and most importantly, the regulation imposed a cap on the number of students receiving FIES in each degree. The cap follows a deterministic rule based on the degree’s quality, field of study, and location. The demand for loans is above the cap for most degrees, and loans are distributed to students through an iterative deferred acceptance mechanism⁴ based on ENEM scores, creating degree-specific score cutoffs to participate in the program. Appendix Figure A.3, Panel (a), shows that the distribution of cutoffs is very wide, with a mean of 495 points and a standard deviation of 45 points, whereas the overall distribution of scores has a mean of 511 points and a standard deviation of 71 points. Therefore, the eligibility cutoffs are relevant for students across the whole score distribution.

Panel (b) shows that the overall score distribution is smooth, and there is no evidence of manipulation. In particular, Panel (c) highlights that there is no evidence of manipulation across the 450-points threshold, the minimum score to receive a FIES loan.

Selected students are supported for as long as they remain enrolled in the degree and must begin to pay back the loan 18 months after graduating or dropping out. The benefit is lost if the student transfers to a different degree.

For clarity, note that admission to private degrees is not centralized, the process described above only distributes loans. As discussed in Section 3.1, most private sector degrees are not selective. So, if they pay out of pocket, students can still enroll even if their score is too low to receive a loan.

3.4 Private financial aid: Tuition discounts

Tuition discounts are widespread in the private sector of the Brazilian higher education market. In 2016, 22% of the incoming students in private colleges received a tuition discount, and the average discount was 33% of the full price. Incoming students who receive tuition discounts usually keep them for the duration of the program, provided they stay in good academic standing.

These discounts are not scholarships to high-achieving students but part of price discrimination policies targeting individuals that could not pay otherwise. Indeed, we find that students receiving discounts have, on average, 25% lower per capita family income, 0.45 lower exam scores (in standard deviations), and are 25% less likely to have college-educated parents.

⁴See Bo and Hakimov (2019) for the formal properties of this mechanism.

Colleges do not allocate discounts by explicitly targeting specific demographics. Instead, they rely on a discrimination strategy like the one followed by airlines and hotels. That is, they impose a search cost to find discounts as a way of targeting the most price-sensitive students. For example, a large share of the discounts is distributed through online marketplaces. These are websites where colleges can post special offers to enroll at discounted prices (example in Appendix Figure A.4). They operate very similarly to travel fare aggregators, such as Expedia or Hotwire. Offers are temporary and often stay online only for a few weeks, creating a cost to find them. We have obtained access to the administrative records of QueroBolsa, the largest marketplace in the country, responsible for 18% of all tuition discounts. Using this data, we find that offers are indeed temporary: 25% of them stay on the platform for less than one week and 50% for less than three weeks.

Finally, it is important to note that FIES regulations stipulate that tuition discounts may not be offered solely based on student loan status. That is, colleges cannot explicitly target tuition discounts based on whether the student has a government loan.

3.5 Data

3.5.1. Individual-level data: To paint a complete picture of students' educational history, financial aid access, and demographics, we combine different sources of data.

The *Census of Higher Education* is a survey run by the National Educational Research Institute (INEP) that covers the universe of students enrolled in higher education. This dataset allows us to observe the college and degree in which students are enrolled and which of these students receive tuition discounts. We are not able to observe, however, the magnitude of the discounts.

The *FIES Administrative Records*, provided by the Ministry of Education, covers the universe of student loans provided by the federal government. We observe which students are participating in the program each year.

Finally, the *ENEM Administrative Records* covers the universe of the students taking the national standardized exam, ENEM, each year. It includes a detailed report of each student's performance, as well as an extensive socio-economic survey. Among other demographics, we observe students' family income, whether they attended a public or private high-school, and whether their parents attended college. Moreover, students also answer questions about why they are taking the exam. Most importantly, it asks students whether they plan to apply for a FIES loan.

We use individual-level identifiers, an encrypted version of students' social security numbers (CPF), to merge all three datasets.

3.5.2. Degree-level price data: To obtain a comprehensive description of private colleges' pricing practices, we combine three different data sources.

First, a nationally representative survey conducted by *Hoper*, a company that provides consulting

services to colleges. It reports the posted price, without any discounts, of the degrees in their sample.

Second, in the *Administrative Records of the National Education Fund (FNDE)*, we observe the payments made by the government to students participating in FIES and in other financial aid programs. From this dataset, we can infer what is the full price and the distribution of discounts for participating degrees.

Third, we have obtained access to the *Administrative Records of QueroBolsa*, the largest degree-search platform, responsible for 18% of the all the tuition discounts offered to incoming students in 2018. We observe the tuition paid by all students that enrolled through the platform. This is the first time academic researchers had access to this data.

We combine these sources in the following way:

The posted price of degrees is taken from the *Hoper* and *FNDE* datasets. To minimize measurement error, we take the average of the two sources when there is an overlap.

Discounts are taken from the *QueroBolsa* and *FNDE* sources. Notice that we observe several different discounts in these sources, but they are not identified and, therefore, cannot be merged with the individual-level datasets. Hence, we take the average among all the observed discounts and call it the degree's discounted price.

To maximize coverage, we collapse these prices into pre-policy change (2013 and 2014) and post-policy change (2016 and 2017) prices. With this procedure, 98% and 80% of the degrees are included in our full and discounted price datasets, respectively. The coverage of our discount data improves for degrees that offer discounts to more students. We cover 86%, 90%, and 93% of the degrees that offer discounts to at least 10%, 25%, and 50% of their students, respectively.

4. Descriptive evidence

In this section, we present the moments in the data that identify the key parameters governing the equilibrium responses to changes in loan availability, in both the demand (Section 4.1) and supply (Section 4.2) sides.

4.1 Descriptive evidence: Demand

Loan availability affects equilibrium prices in two ways. First, loans reduce price sensitivity, which leads to higher prices (direct effect). Second, loans change the composition of the market, which has an ambiguous effect on prices (composition effect). If the participation of more (less) price-sensitive students increases, prices go down (up).

The magnitude of these two forces depends on the parameters of the demand curve. The strength of the direct effect depends on how much receiving a loan reduces students' sensitivity to prices. The strength of the composition effect depends on two factors: first, the variation of price sensitivity

across students; second, the accuracy of targeted discounts. In this section, we show the patterns in the data that identify each of these three parameters. In Section 6, the moments described here will be used to estimate the demand-side parameters of our structural model.

We begin by discussing our strategy to estimate price elasticities. For this, we rely on the variation of prices across time and between degrees. To address the potential endogeneity of prices, we build an instrument that relies on the panel structure of the data and on the ownership relations of multiregional college chains. More specifically, we instrument the price of degree j with the average price in period t of degrees owned by the same chain as degree j but located in a different region (\tilde{p}_{jt}). This strategy explores the fact that prices of multiregional firms are often highly correlated across regions, due to managerial inertia, as shown in DellaVigna and Gentzkow (2019).

A common concern with this identification strategy is that demand shocks might also be correlated across regions. To address this issue, we include region-year and degree fixed effects in our estimates. Hence, only time-varying chain-specific shocks would bias our estimates. One example of such shocks would be if certain college chains specialize in majors for which demand is growing. To investigate whether differential trends in demand are driving our results, we create a proxy for the national trends in the demand for each field of study and use it as a control. More specifically, for each degree j , we proxy the national trends in demand for its field as the total enrollment of degrees of the same field of study as degree j , but in other regions (\tilde{N}_{jt}).

In Section 6, we show how we use this strategy to estimate our structural model. However, in a simplified model without price discrimination or loans, students' price sensitivity α can be directly estimated with the following two-stage least squares (2SLS) regression:

First Stage:

$$p_{jt} = \beta_0^I + \beta_1^I \cdot \overbrace{\tilde{p}_{jt}}^{\text{Instrument}} + \beta_2^I \cdot \log \tilde{N}_{jt} + \gamma_j + \gamma_{mt} + \epsilon_{jt}^I \quad (2)$$

Second Stage:

$$\log N_{jt} = \beta_0^{II} + \alpha \cdot \overbrace{p_{jt}}^{\text{Endogenous regressor}} + \beta_2^{II} \cdot \log \tilde{N}_{jt} + \gamma_j + \gamma_{mt} + \epsilon_{jt}^{II}$$

where N_{jt} is the number of students enrolled in degree j in period t , γ are fixed effects, and ϵ are residuals.

The estimates of Equation 2 are in Table 1, separately for students with family income above and below three times the national minimum wage (low- and high-income). Note that the elasticities estimated by 2SLS are substantially more negative than those obtained with a simple OLS regression, suggesting that prices are strongly correlated with demand shocks and reinforcing the importance of using an instrument. Moreover, low-income students are much more price elastic than high-income ones. This suggests that the composition effect of loans might be a major force in determining prices.

Table 1: Estimated price elasticities

Outcome: log Number of incoming students						
	High-income			Low-income		
	(1) OLS	(2) 2SLS	(3) 2SLS	(4) OLS	(5) 2SLS	(6) 2SLS
Average Price Elasticity	0.08 (0.06)	-0.27 (0.27)	-0.38 (0.26)	-0.14 (0.10)	-1.78*** (0.49)	-1.89*** (0.49)
log Students in Same Field			0.94*** (0.10)			1.13*** (0.10)
Degree Fixed Effect	X	X	X	X	X	X
Year-Region Fixed Effect	X	X	X	X	X	X
Observations	4180	4180	4180	4574	4574	4574
Regions	35	35	35	35	35	35
1st Stage F-statistic		82	83		85	85

Notes: This table shows the estimates of regression (2). Each observation is a degree-year. The sample includes two years (2014 and 2016) and is restricted to the regions with at least 5,000 students. A region is a commuting zone (*Microregião*), as defined by the Brazilian Institute of Geography and Statistics (IBGE). Average price elasticities are calculated as the average of $\alpha \cdot p_j \cdot (1 - s_j)$, where j is a degree, p are prices, and s are market shares. Standard errors presented in parentheses are clustered at the degree level. * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$.

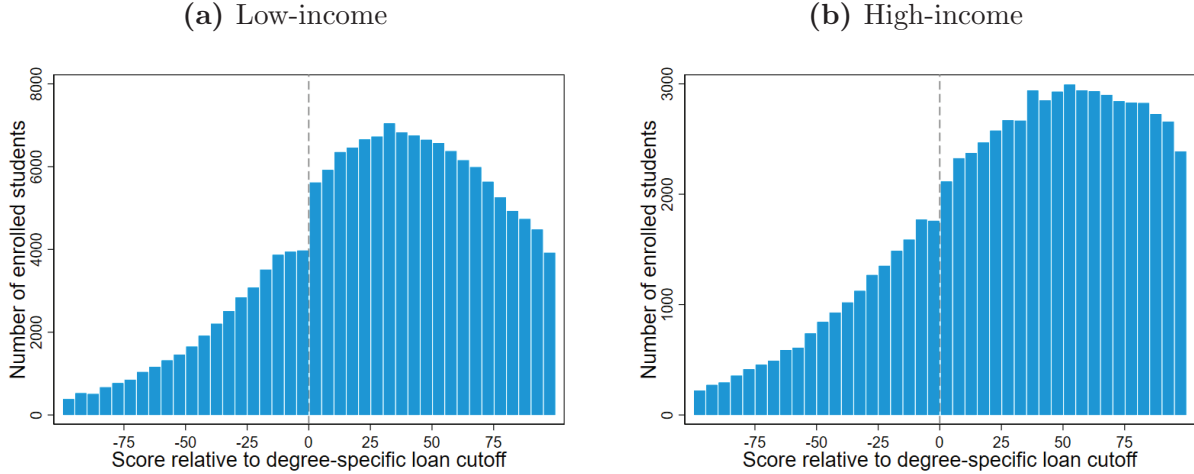
Finally, we see in columns (3) and (4) that changes in demand for different fields of study are indeed correlated across regions. However, controlling for these demand changes does not affect our price elasticity estimates. This suggests that correlated demand shocks are not driving our results and supports our identification strategy.

We now provide evidence on how access to loans changes students' enrollment decisions. Identification is complicated because loan recipients are not randomly selected. To overcome this issue, we leverage quasi-random variation created by the centralized mechanism that allocates loans to students as described in Section 3. The mechanism determines a minimum score to receive a loan in each degree; hence, access to loans varies discontinuously at these thresholds.

We explore this discontinuity by studying how the probability of enrollment changes at the thresholds. In the absence of score manipulation, students right above or below the cutoff to receive a loan for a given degree are, in expectation, identical. Indeed, as discussed in Section 3.3 and shown in Appendix Figure A.3, the overall distribution of scores is smooth and there is no evidence of manipulation. Hence, if access to loans does not change enrollment decisions, we should see students above and below the cutoff enrolling in the degree with the same frequency. In other words, the distribution of student scores, relative to the loan cutoff of the degree in which they enrolled, should be smooth.

Figure 2 shows the distribution of scores relative to the degree-specific cutoffs to receive a federal loan in the degree in which each student is enrolled, separately for low- and high-income students.

Figure 2: Access to student loans and enrollment decisions



Notes: This figure shows the distribution of relative scores, defined as the score of a student minus the cutoff to receive a FIES loan in the degree in which she is enrolled, for low- and high-income students. Sample: all students enrolled in degrees participating in FIES in 2016. Source: Census of Higher Education, ENEM administrative records, and FIES administrative records.


There is a clear jump at the threshold, indicating that students are self-selecting into degrees for which they could get loans. This shows that access to loans substantially changes students' enrollment decisions. Moreover, the jump is much more pronounced for low-income students, suggesting that this group is more affected by access to loans.

Finally, we now discuss the targeting of tuition discounts. Figure 3 shows the share of students receiving discounts, as a function of their per capita family income. We see that low-income students are much more likely receive discounts: 60% of the students in the bottom income ventile have discounts, compared to 26% of the ones in the top. However, most of this difference is driven by certain colleges offering more discounts. Also in Figure 3, we see that, within the same college, students in the bottom ventile are only 5 percentage points more likely to receive a discount than the ones in the top ventile. The regression estimates corresponding to the fitted lines are in columns (1) and (2) of Table 2.

We draw two lessons from these patterns. First, colleges' discount-targeting technology is imperfect, resulting in a weak correlation between income and discounts within each college. Second, low-income students, being more price elastic, select into colleges that offer more discounts, resulting in a strong correlation between income and discounts across colleges. Both of these patterns will be present in our structural model.

An alternative explanation for the weak within-college correlation between income and discounts would be that colleges are targeting other characteristics. For example, some of the tuition discounts might be scholarships to high-scoring students, who are, on average, also high-income. However, Table 2 shows that this is not the case: high-scoring students are actually less likely to receive

Figure 3: Tuition discounts and family income



Notes: This figure shows the share of students receiving tuition discounts as a function of their family per capita income. The sample includes all students who took ENEM in 2015 and enrolled in a in-person tuition-charging college that offered discounts in 2016. Data on enrollment and discounts from the Census of Higher Education and income from the ENEM administrative records.

discounts. Moreover, all the other observed student characteristics are also only weakly correlated with discounts. For example, we see in Table 2 that within a college, public high-school students are only 3.3 percentage points more likely to receive discounts.

Table 2: Tuition discounts and student characteristics

	Outcome: $100 \times \mathbb{1}\{\text{Has Tuition Discount}\}$			
	(1)	(2)	(3)	(4)
log per Capita Income	-8.5*** (0.7)	-1.1*** (0.3)	-5.5*** (0.6)	-0.2 (0.2)
Score (normalized)			-4.5*** (0.4)	-0.9** (0.3)
Parents Have College Education			-3.6*** (0.6)	-0.0 (0.3)
Public High School			1.6* (0.8)	3.3*** (0.4)
College Fixed Effect		X		X
Observations	173194	173194	173194	173194
Colleges	421	421	421	421

Notes: This table shows the OLS estimates of a regression of an indicator variable of receiving tuition discounts on a series of student characteristics. Each observation is a student. The sample includes all students who took ENEM in 2015 and enrolled in a in-person tuition-charging college that offered discounts in 2016. Data on enrollment and discounts from the Census of Higher Education and data student characteristics from the ENEM administrative records. Standard errors presented in parentheses are clustered at the degree level. * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$.

Taken together, the results of this section can teach us the following about the effects of student loan programs. Since loans reduce students price elasticities, there is scope for a direct effect of loans on equilibrium prices. Moreover, there are two indications that the composition effect might also play a major role. First, discount targeting is imperfect. This means that students with different price elasticities often pay the same price, which is the reason why a composition effect exists in the first place. Second, low-income students are much more price elastic than high-income students. This means that, if loans increase the market share of low-income students, the average price elasticity of the market will increase, potentially leading to lower markups and prices.

4.2 Descriptive evidence: Supply

In our structural model, we will estimate the supply-side parameters governing how colleges respond to policy changes using a revealed preference argument. That is, we explore the observed responses to the reduction in loan availability that occurred in 2015 and find the parameters that rationalize these responses.

More specifically, the contraction in the loan program was followed by a drastic expansion in tuition discounts. In this section, we describe this expansion and provide evidence that it was indeed a response to the policy change. In Section 6, we use the moments described here to estimate the supply-side parameters of our structural model.

Figure 1, Panel (a), describes the overall trends before and after the policy change and shows that, as FIES shrank, colleges themselves offered more aid. In 2014, 22% of the incoming students in tuition-charging institutions received federal loans whereas 15% received a tuition discount. In 2018, these numbers were 3% and 31%, respectively.

These trends indicate that private institutions responded to reduction in student loan availability by offering more discounts. However, other contemporaneous shocks might be driving the patterns seen in Figure 1. To investigate whether the patterns observed since 2015 are indeed a response to the policy change, we compare the outcomes of degrees that were differentially exposed to the contraction. The policy imposed a cap on the number of students receiving a loan in each degree, following a rule based on degrees' quality, region, and field of study. The rule did not consider the number of students with loans each degree previously had. Therefore, degrees that received more students with loans until 2014 were disproportionately affected by the policy.

Guided by this intuition, we measure exposure to the policy change as the share of students with loans in each degree in 2012, and use an event study approach in which we estimate the following regression:


$$S_{jt}^D = \sum_{\tau=2012}^{2017} \beta_{\tau} \cdot S_{j2012}^L \cdot \mathbb{1}\{t = \tau\} + \gamma_j + \gamma_{ht} + \epsilon_{jt} , \quad (3)$$

where j is a degree and t is a year, S_{jt}^D is the share of students with discounts in year t , S_{j2012}^L is the share of students with loans in 2012, γ_j are degree fixed effects, and ϵ_{jt} are residuals. Mirroring the rule that allocate loans to degrees, we also include quality-region-field-year fixed effects (γ_{ht}). The sample includes three pre-policy (2012-2014) and four post-policy (2015-2017) years.

The OLS estimates of β_{τ} are in Figure 4. Until 2014, degrees with a larger or smaller fraction of students with loans followed a parallel trend. After 2015, degrees that previously had more students with loans began offering more discounts. This suggests that the increase in discounts seen since 2015 was indeed a response to the policy change.

Note, however, that all degrees were affected by the reduction in loan availability through general equilibrium forces, even the degrees that did not enroll any students with loans. Hence, we can not directly estimate the total effect of the policy change from the event study shown in Figure 4 because there is no pure control group. To assess the equilibrium effects of subsidized loans taking these endogenous responses into account, we present our model of the supply and demand for higher education in the next section.

Figure 4: Tuition discounts and the reduction in loan availability



Notes: This figure shows the OLS estimates of β_τ in equation (3). The vertical dashed line marks the reduction in the availability of subsidized student loans. All degrees in private colleges are included in the sample. Each observation is a degree-year, and a degree is a major-college pair. The brackets represent 5% confidence intervals. Source: Census of Higher Education and FIES administrative records.

5. Structural model

We now present our equilibrium model of the supply and demand for higher education. The goal of the model is to simulate student loan programs with different targeting rules and how they compare in terms of enrollment and student welfare. Hence, we design the model to capture how students and colleges react to changes in the availability of loans.

On the demand side, students' enrollment and degree choice decisions depend on tuition costs, college characteristics, and loans accessibility. With this framework, we can predict how students would behave under different policy designs. Moreover, we can construct the demand curve faced by each college—an essential input to understanding supply-side responses.

On the supply side, colleges choose prices and discounts to maximize profits. The availability of student loans affects these decisions in two ways. First, loans make students less price elastic, which leads to higher prices. Second, loans change the composition of students who participate in the market, which has an ambiguous effect on prices. If the participation of more (less) price-sensitive students increases, prices go down (up).

We present our model in two parts. In this section, we introduce a general version of the model with fewer parametric assumptions to highlight the key mechanisms we want to capture. Later, in Section 6.1, we discuss the additional assumptions we make before taking the model to the data. In particular, in the current section, we allow for fully flexible heterogeneity across students; that is, we

allow all parameters to be student specific. In Section 6.1, we differentiate between observable and unobservable student characteristic and impose restrictions on parameter heterogeneity.

5.1 Setting

The market is composed by three actors: students, colleges, and the government. In each period $t \in \mathcal{T}$, the government acts first and chooses the rules determining the allocation of student loans. Second, given these rules, colleges choose a price schedule for each of their degrees. Third, students choose a degree to enroll into.

The government allocates degree-specific loans. That is, the government sets $L_{ij} \in \{0, 1\}$ which denotes whether student i receives a loan if enrolled in degree j .

Colleges are indexed by $f \in \mathcal{F}$ and offer a fixed set of degrees \mathcal{J}_f . Each degree is a major-college pair—for example, Economics at Pythagoras University. The set of all available degrees is $\mathcal{J} = \bigcup_{f \in \mathcal{F}} \mathcal{J}_f$. Colleges price discriminate between students by choosing student-specific prices p_{ij} for each of its degrees. There are no capacity constraints and all degrees are available to all students.⁵

Finally, there is a continuum of students, indexed by i . The set of students in each period t is denoted \mathcal{I}_t . Each student enrolls in one degree and d_{ij} is a choice indicator. Choosing the outside option ($d_{i0} = 1$) represents attending a public institution, an institution in another region, or not enrolling in higher education.

5.2 Demand: degree choice

Students choose a degree to maximize their utility:

$$d_{ij} = \mathbb{1}(U_{ij} \geq U_{ik}, \forall k \in \mathcal{J}),$$

where U_{ij} is the utility of student i enrolled in degree j in period t , and is given by:

$$U_{ij} = \beta_i^h \mathbf{h}_j + \alpha_{iL_{ij}} p_{ijt} + \xi_{ijt}, \quad (4)$$

where p_{ij} is a student-specific price of degree j , $\alpha_{iL_{ij}}$ measures price sensitivity, \mathbf{h}_j are fixed degree characteristics, β_i^h represents preference heterogeneity for these characteristics, and ξ_{ijt} is a student-specific demand shock.

This framework departs from the standard differentiated-product demand model in two ways. First, since tuition discounts are widespread in the Brazilian higher education market (Section 3.4), we allow prices to be student specific. Second, the price sensitivity parameter ($\alpha_{iL_{ij}}$) depends on whether the student has a loan, which captures several mechanisms through which loans change

⁵Over 90% of the degrees in the private sector fill less than 80% of their spots and, in Appendix C, we show that there is no evidence that private colleges select students based on their score.

enrollment decisions. First, loans alleviate credit constraints. Second, the interest rate is subsidized, and default rates are high. Hence, the effective price paid by students with loans is lower than the actual price.

5.3 Supply: pricing decisions

Each college f offers a set \mathcal{J}_f of degrees and chooses prices to maximize profits. Our model departs from the standard Bertrand differentiated product framework by allowing colleges to choose not only one price per degree, but rather a price schedule. As discussed in Section 3.4, tuition discounts are widespread in the Brazilian higher education market. Moreover, as shown in Section 4.2, the number of students with discounts grew dramatically after the reduction in the availability of loans. Hence we include price discrimination in our model to capture these patterns and to correctly predict how colleges would behave under different policy designs.

Colleges have a limited ability to price discriminate. As discussed in Section 3.4, they do not explicitly target discounts based on students' characteristics, but rather by offering seasonal promotions and posting special deals in certain platforms. Hence, we assume that each student i has a propensity to find promotions, denoted B_i^D , and students with higher B_i^D pay a lower price for the same degree. Due to data limitations,⁶ we further assume there are only two prices per degree. That is, colleges choose a full price (p_{jt}^F) and a discounted price (p_{jt}^D) for each of their degrees. They also choose how hard it is to get a discount, in the form of a discount threshold (\bar{D}_{ft}). Hence, the price paid by each student is given by:

$$p_{ij} = \begin{cases} p_{jt}^F, & \text{if } B_i^D < \bar{D}_{ft} \\ p_{jt}^D, & \text{if } B_i^D \geq \bar{D}_{ft} \end{cases} . \quad (5)$$

College f 's maximization problem in each period t is given by:

$$\begin{aligned} & \max_{\{p_{jt}^F, p_{jt}^D\}_{j \in \mathcal{J}_f}, \bar{D}_{ft}} \sum_{j \in \mathcal{J}_f} \left\{ \overbrace{\Pi_{jt}(\mathbf{p}_t)}^{\text{profits}} + \overbrace{\Omega_{jt}(\mathbf{p}_t)}^{\text{structural error}} \right\} \\ \Pi_{jt}(\mathbf{p}_t) & \equiv \int_{i \in \mathcal{I}_t} s_{ij}(\mathbf{p}_i) \cdot [p_{ij} - c_{jt}] di \\ \Omega_{jt}(\mathbf{p}_t) & \equiv s_{jt}(\mathbf{p}_t) \cdot \kappa_{jt} \cdot s_{jt}^D(\mathbf{p}_t) - s_{jt}(\mathbf{p}_t) \cdot \psi \cdot (\bar{D}_{ft} - \omega_{ft})^2 \\ & \text{subject to: } p_{jt}^D \leq p_{jt}^F, \forall j \in \mathcal{J}_f, \end{aligned} \quad (6)$$

where $\mathbf{p}_i \equiv \{p_{ij} | j \in \mathcal{J}\}$ is the vector of prices faced by student i , $\mathbf{p}_t \equiv \{\mathbf{p}_i | i \in \mathcal{I}_t\}$ is the set of all the prices in the market, $s_{ij}(\mathbf{p}_i) \equiv P(d_{ij} = 1 | \mathbf{x}_i, \mathbf{p}_i)$ is the probability that student i enrolls in degree j ,

⁶As discussed in Section 3.5.2, we just observe the full and the average discounted of each degree.

$s_{jt}(\mathbf{p}_t) \equiv P(d_{ij} = 1 | \mathbf{p}_t, i \in \mathcal{I}_t)$ is the market share of degree j , $s_{jt}^D(\mathbf{p}_t) \equiv P(D_{ij} = 1 | \mathbf{p}_t, d_{ij} = 1, i \in \mathcal{I}_t)$ is share of degree j 's students with discounts, and c_{jt} is the marginal cost of enrolling and educating each student.

The structural error (Ω_{jt}) allows colleges to depart from pure profit maximization and has two components. First, colleges receive an extra utility k_{jt} from enrolling a student with discount, representing the fact that some of colleges have social goals. Second, ω_{jt} further specifies such social goals and represents the type of student the college prefers to offer discounts to. The parameter ψ governs to which degree colleges discount policies are driven by these non-profit considerations. On the extreme, if $\psi = \infty$, then $\bar{D}_{ft} = \omega_{jt}$, regardless of market conditions.

The structural error also captures misspecification on colleges' marginal costs. For example, degree search platforms (see Section 3.4) charge a fee for each student who enrolls through their websites. Hence, colleges using these platforms face a larger marginal costs for students receiving discounts, which is captured in our model by κ_{jt} . More generally, the structural error captures all the factors that determine colleges' discount policies beyond differences in demand. This is analogous to the role played by marginal costs in standard Bertrand differentiated products empirical models: capturing unexplained variation in prices .

The main goal of our model is to predict the outcomes of different loan policy designs, and marginal costs and structural errors will be kept fixed in our counterfactuals. Therefore, a key assumption behind our framework is that colleges' exogenous reasons to offer discounts are not affected by changes in the loan program. An example of a violation of this assumption would be if a reduction in federal financial aid led to an increase in the social pressure for colleges to offer more support to their students. If this were the case, expanding public financial aid would crowd out private aid to a larger extent than predicted by our model.

Finally, notice that we take the set of available degrees and the quality of these degrees as fixed. Hence, market structure will be invariant to changes in the loan program in our counterfactuals. This assumption reflects the observation that the considerable reduction in loan availability that occurred in 2015 did not change market structure, at least in the short run. Indeed, in 2014, only 0.3% (1.3%) of students were enrolled in colleges (degrees) that closed by 2015. In Appendix D, we provide further evidence that the policy did not lead to an increase in exit or reduction in quality.

To better understand what drives colleges pricing decisions, we now present the first order conditions of colleges' problem. To present the intuition more concisely, we show here the first-order conditions of a single product firm and omit t subscripts. The general case is in Appendix E. With respect to prices, we have:

$$p_j^F = \underbrace{c_j}_{\text{marginal cost}} + \underbrace{\frac{E_i [s_{ij} | B_i^D < \bar{D}_j]}{E_i \left[-\frac{\partial s_{ij}}{\partial p_j^F} | B_i^D < \bar{D}_j \right]}}_{\text{markup}} + \underbrace{\psi \cdot (\bar{D}_f - \omega_f)^2}_{\text{structural error}}, \quad (7)$$

$$p_j^D = \underbrace{c_j}_{\text{marginal cost}} + \underbrace{\frac{E_i [s_{ij} | B_i^D \geq \bar{D}_j]}{E_i \left[-\frac{\partial s_{ij}}{\partial p_j^D} | B_i^D \geq \bar{D}_j \right]}}_{\text{markup}} + \underbrace{\left[\psi \cdot (\bar{D}_f - \omega_f)^2 - \kappa_j \right]}_{\text{structural error}}. \quad (8)$$

Pricing policies are similar to the standard Bertrand differentiated product model. That is, prices are set at the marginal cost plus a markup that depends on student's price sensitivity $\left(\frac{\partial s_{ij}}{\partial p_{ij}}\right)$. There are two key differences from the standard framework. First, when setting full-price markups, colleges consider only the price sensitivity of students who would pay their full price, as shown in the integration limits inside the conditional expectations in Equation (7). The analogous is true for the discounted price. This creates a difference between the full and discounted price markups, which drives the magnitude of the discounts. Second, there is a structural error in the price equation, and the difference between the structural errors on the full and discounted price equations is κ_j . Hence, this parameter captures the variation in the magnitude of the discounts beyond what is explained by differences in markup.

Equations (7) and (8) show the two channels how loan availability can affect prices. First, loans reduce price sensitivity $\left(\frac{\partial s_{ij}}{\partial p_{ij}}\right)$, which leads to higher prices. This is the direct effect. Second, loans change market shares (s_{ij}), which has an ambiguous effect on prices. If the market shares of the more price-sensitive students increase, prices go down. If the market shares of the less price-sensitive students increase, prices go up. This is the composition effect. The magnitude of these two forces depends on the parameters of the model. The strength of the direct effect depends on how much receiving a loan reduces students' sensitivity to prices. The strength of the composition effect depends on how much price sensitivity variation there is among students paying the same price, which in turn depends on two factors. First, how much price sensitivity varies in the overall population. Second, how well discounts are targeted. If targeting is very precise, only students with similar price sensitivity will pay the same price and the composition effect will be muted.

Now, let us consider the first order condition with respect to the discount threshold:

$$\bar{D}_f = \frac{1}{2\psi} \cdot \underbrace{E[\pi_{if}^F - \pi_{if}^D | B_i^D = \bar{D}_f]}_{\text{profit effect}} + \underbrace{\omega_f}_{\text{structural error}} \quad (9)$$

where π_{if}^F is the profit college f receives from student i if it charges the student the full price and π_{if}^D is the analogous for the discounted price.

The first element in Equation (9) is the marginal effect of increasing \bar{D}_f on expected profits. Making it harder to get a discount (increasing \bar{D}_f) affects profits in two opposite ways: losses from students who do not enroll without a discount versus gains from charging a higher price. If this profit effect is positive, the college chooses \bar{D}_f above its exogenous preference (ω_f), and vice versa. Note that the introduction of loans changes demand, which affects colleges' expected profits. Depending on how loans and discounts are targeted, the effects on the profits realized from full-price (π^F) and

discounted-price (π^D) students might be different. Hence, Equation (9) describes how colleges change their discount policies (\bar{D}_f) in response to variations in loan availability.

5.4 Government: loan allocation

As discussed in Section 3.3, the government allocates loans to students in two ways. First, it establishes a series of socioeconomic eligibility conditions. Second, it chooses a set of cutoffs $\{\bar{r}_{jt}\}_{j \in \mathcal{J}, t \in \mathcal{T}}$ such that student i can only receive a loan to enroll in degree j if their score is high enough ($r_i \geq \bar{r}_{jt}$).

Formally, L_{ij} denotes whether student i receives a loan if she enrolls in degree j and is given by:

$$L_{ij} = \begin{cases} 0, & \text{if } B_i^L = 0 \text{ or } r_i < \bar{r}_{jt} \\ 1, & \text{if } B_i^L = 1 \text{ and } \bar{r}_{jt} \leq r_i \end{cases}, \quad (10)$$

where B_i^L denotes whether student i uses a loan when enrolled in a degree for which their score is high enough. We can have $B_i^L = 0$ for two reasons: because the student does not meet the socioeconomic eligibility criteria or does not take up the loan. Students might not take up loans because of stigma, lack of information, missing documents, etc. For conciseness, we will refer to B_i^L as an indicator of whether student i is a loan taker.

The government uses both policy levers (score cutoffs and socioeconomic eligibility) to allocate loans in a way that promotes access to higher education at a low fiscal cost. However, this is challenging because it requires correctly estimating how loans affect students' enrollment choices and how colleges respond to changes in loan availability. Since loans change students' price elasticities, their allocation affects colleges' pricing decisions, which have important consequences for enrollment and student welfare. In Section 7, we use this framework to simulate the outcomes of alternative policies and show how the government could use these two instruments to achieve more desirable allocations.

6. Structural estimation

We now describe how we estimate the parameters of the model introduced in Section 5. We begin in Section 6.1 by presenting additional parametric assumptions we impose before taking our model to the data. We then present our identification and estimation strategies for demand (Section 6.2) and supply (Section 6.3) parameters. Finally, Section (6.4) presents the estimated parameters.

We estimate the model using data from one pre- and one post-policy change year: 2014 and 2016. Due to computational constraints, we restrict the sample to the Rio de Janeiro commuting zone, one of the largest markets in the country.⁷ We define the market as the universe of students

⁷We define a commuting zone as a *Microregião*, following the Brazilian Institute of Geography and Statistics

taking the centralized exam (ENEM) in Rio de Janeiro in each of these years. Table 3 shows the descriptive statistics of our structural estimation sample and how it compares with the rest of the country. We see that Rio de Janeiro is wealthier, with a larger share of high-income students and higher average exam scores. However, the most pronounced difference is that the market for private higher education in Rio de Janeiro is local. Only 0.3% of the students who take ENEM in Rio de Janeiro enroll in a private college in another market, whereas the country average is 3%. That is, Rio de Janeiro is an isolated market, making it ideal for studying equilibrium effects.

6.1 Parameterization

6.1.1. Demand: Before taking the model to the data, we impose the following restrictions to the components of the utility function:

$$\begin{aligned}\alpha_{iL_{ij}} &= \alpha_{w_i L_{ij}} , \\ \xi_{ijt} &= \gamma_{w_{ij}} + \gamma_{w_{it}} + \xi_{w_{ijt}} + \epsilon_{ij} , \\ \beta_i^h \mathbf{h}_j &= r_i \beta_{w_i}^r \cdot R_j ,\end{aligned}\tag{11}$$

where w_i indicates whether student i 's family income is above or below three times the national minimum wage (we refer to these groups as low- and high-income), r_i is the student's score on ENEM, R_j is the average score of the students in degree j in the baseline year (2014), γ_{w_j} are degree fixed effects, γ_{wt} are year fixed effects, and ϵ_{ijt} is an individual idiosyncratic preference shock assumed to follow a type-I extreme value distribution.

We allow price sensitivity (α) to depend, fully flexibly, on the student's family income and on whether the student receives a loan. The unobserved demand shifter (ξ), the fixed effects (γ), and the score parameter (β^r) are also income specific.

We also allow for preference heterogeneity based on scores ($r_i \beta_{w_i}^r \cdot R_j$). Notice that R_j is kept fixed and will not change when computing counterfactual equilibria. Hence, it is a proxy for fixed degree characteristics that appeal to high-scoring students, and it is not supposed to capture preferences for high-scoring peers. Note that peer effects are likely to be less relevant in our context since these are commuter schools—only 0.05% of the students live on-campus—, with few opportunities for interaction outside of the classroom—only 16% of the students participate in extracurricular activities.

Finally, there is no cost to choose the outside option ($p_{i0t} = 0$) and its utility is given by:

$$U_{i0} = r_i \beta_{w_i}^r \cdot R_0 + \epsilon_{i0} .$$

Notice that we allow the value of the outside option to depend on scores. The outside option

(IBGE).

Table 3: Descriptive statistics of structural estimation sample

	Brazil		Rio de Janeiro	
Number of regions	356		1	
Number of private colleges	1,847		72	
Number of private degrees	23,701		1,385	
	Low-income	High-income	Low-income	High-income
<i>Number of students:</i>				
Total	6,743,732	2,204,709	333,728	154,808
In outside option	6,214,545 (92%)	1,927,889 (87%)	300,182 (89%)	134,519 (86%)
<i>Outside-option breakdown:</i>				
Not enrolled	5,563,596 (90%)	1,579,805 (82%)	280,603 (93%)	115,524 (86%)
Public college	294,147 (4.7%)	217,873 (11%)	11,980 (4.0%)	15,891 (12%)
Online degree	149,863 (2.4%)	26,241 (1.4%)	5,501 (1.8%)	1,533 (1.1%)
Not in home region	167,291 (2.7%)	87,946 (4.6%)	789 (0.3%)	714 (0.5%)
Missing tuition data	39,648 (0.6%)	16,024 (0.8%)	1,309 (0.44%)	857 (0.64%)
<i>Covariates:</i>				
Score (normalized)	-0.22	0.68	-0.09	0.90
Plans to use loan	53%	68%	42%	25%
Public High School	90%	53%	79%	37%
Parents have college educ.	7%	39%	7%	41%
<i>Financial aid:</i>				
Has federal loan	27%	15%	24%	10%
Has tuition discount	18%	15%	39%	31%

Notes: This table shows the descriptive statistics of our structural estimation sample. A *region* is a commuting zone (*Microregião*), as defined by the Brazilian Institute of Geography and Statistics (IBGE). *Score* is the score in ENEM, *Plan to use a loan* is the answer to a survey question about whether the student plans to apply to FIES, *Public High School* indicates whether the student went to a public high school, *Parents have college educ.* indicates whether any of the students' parents has at least one year of higher education. Average covariates are calculated among all students in the sample, and average financial aid usage among students not in the outside option.

represents both attending a public institution and not attending college; consequently, it will have a different value for high-scoring and low-scoring students.

6.1.2. Financial aid targeting: Our model features two types of financial aid: public (subsidized loans) and private (discounts). Whether a student receives financial aid is determined by unobserved characteristics, B_i^D and B_i^L . As defined in Section 5.3, B_i^D is student i 's propensity of finding discounts. As defined in Section 5.4, B_i^L denotes whether student i is a loan taker. We parameterize:

$$\begin{aligned} B_i^D &= \rho^D(\mathbf{x}_i) + \eta_i^D \\ B_i^L &= \mathbb{1}\{\rho_i^L(\mathbf{x}_i) \geq \eta_i^L\} , \end{aligned} \tag{12}$$

where \mathbf{x}_i is a vector of observable student characteristics, ρ^D and ρ^L are functions with parameters to be estimated, and η_i^D and η_i^L are idiosyncratic errors, independent of each other.

We assume that η_i^D follows a logistic distribution and ρ^D is linear:

$$\rho^D(\mathbf{x}_i) = \boldsymbol{\beta}^D \mathbf{x}_i ,$$

where $\boldsymbol{\beta}^D$ is a parameter to be estimated.

We estimate ρ_i^L non-parametrically. Hence we can normalize, without loss of generality, η_i^L to be uniform $[0, 1]$, so that $P(B_i^L = 1 | \mathbf{x}_i) = \rho_i^L(\mathbf{x}_i)$.

We use different sets of covariates to estimate ρ^L and ρ^D . For ρ^L , we use the student's family income, whether the student attended a public high school, whether at least one of the student's parents attended college, and the answer to a survey question asking whether the student plans to apply for a FIES loan. We collapse family income into four bins divided at two, three, and seven times the federal minimum wage. More price-sensitive students have a stronger incentive to search for discounts, and, in our model, price sensitivity depends only on income. Hence, we use only family income to estimate ρ^D , divided in the same four bins.

6.2 Estimation: Demand

Two groups of parameters determine the demand colleges face: first, the parameters of students' utility function, described in Equations (4) and (11); second, the functions governing financial aid targeting, defined in Equation (12).

In Section 6.2.1, we describe how we estimate the preference parameters, taking financial aid targeting as given. In Section, 6.2.2, we do the reverse: estimating the targeting functions, taking preferences as given. Since the two groups have to be estimated together, we perform these two steps in a loop, until it converges to a point that jointly satisfies the identification conditions presented in both Sections 6.2.1 and 6.2.2.

6.2.1. Students’ preferences: We estimate the parameters of students’ utility function using a method of moments estimator. The model is exactly identified, so our estimation finds the parameters that make the moment conditions hold exactly. Even though all parameters are jointly estimated, each moment influences more strongly the estimation of one of the parameters. We now present each of these moments, its associated parameter, and the corresponding identification assumptions.

We begin by describing how we estimate the parameters governing preferences for degrees characteristics that are exogenously determined and will be kept fixed in the counterfactuals. Time-varying demand shocks (ξ_{wjt}) and degree (γ_{w_ij}) and year (γ_{w_it}) fixed effects are jointly estimated by matching the market shares of each degree, in each period, for each income group. We normalize $\sum_{t \in \mathcal{T}} \xi_{wjt} = 0$, for all (w, j) , and $\xi_{w0t} = 0$, for all (w, t) .

The score parameter (β_w^r) is estimated by matching the correlation between the average scores in each degree in the periods included in our sample with the average scores in the baseline period (R_j).⁸

We now move to the estimation of the price sensitivity parameters. We estimate two different sensitivity parameters for each student type: with and without loans. We jointly estimate these parameters using two moments, one that identifies the average price sensitivity and another that identifies the difference between the sensitivity with and without loans.

To estimate students’ average price sensitivity, we rely on variation of prices across time and degrees. As discussed in Section 4.1, to address the potential endogeneity of prices, we build an instrument that relies on the panel structure of the data and on the ownership relations of multiregional higher education chains. More specifically, we instrument the prices of degree j with \tilde{p}_{jt} , the average price in period t of degrees owned by the same chain as degree j but located in a different region. Our identifying assumption is:

$$E[\Delta \tilde{p}_{jt} \cdot \Delta \xi_{wjt}] = 0.$$

That is, conditional on degree and year fixed effects, demand shocks are not correlated within degrees owned by the same chain. This is similar to the approach proposed by Hausman (1994) and applied by Nevo (2001) and DellaVigna and Gentzkow (2019). To provide support for this identification strategy, we have shown in Section 4.1 that prices of degrees of the same chain are strongly correlated across regions, and found no evidence that common demand shocks drive this pattern.

To estimate the difference in price sensitivity between students with and without loans, we leverage the discontinuity in enrollment across the loan eligibility threshold, described in Section 4.1 and shown in Figure 2. More specifically, we match the size of this discontinuity, in the model and in the data, separately for low-income and high-income students.

⁸As discussed in Section 6.1.1, R_j is meant to capture fixed degree characteristics that are appealing to high-score students, not preferences for high-score peers. Hence, we do not use an instrument for R_j and it will be kept fixed in our counterfactuals.

As discussed in Section 3.3, there is no evidence of manipulation in ENEM scores. Hence, differences in enrollment across the loan eligibility threshold must come from students who receive and who do not receive loans making different choices. Exploiting this insight, we show that, in a simplified model without discounts and in which all students are loan takers:

$$\alpha_{w1} - \alpha_{w0} = \frac{1}{J} \sum_{j \in \mathcal{J}} \frac{\lim_{r \rightarrow \bar{r}_j^+} \log \frac{s_{jr}}{s_{0r}} - \lim_{r \rightarrow \bar{r}_j^-} \log \frac{s_{jr}}{s_{0r}}}{p_j}, \quad (13)$$

where J is the number of degrees in the market, s_{jr} is the market share of degree j among students with score r . The proof is in Appendix F.

Equation (13) shows how the discontinuity in enrollment across the loan eligibility threshold identifies our parameter of interest—the difference in price elasticity between students with and without loans ($\alpha_{w1} - \alpha_{w0}$). Its right-hand side is the average discontinuity in enrollment across all degrees. Each discontinuity is normalized by the degree’s price because the demand for more expensive degrees is more affected by changes in price sensitivity. The left-hand side gives us our parameter of interest. However, we cannot use this closed form estimator for two reasons. First, the data is sparse: s_{jr}^L is zero for most entries. Second, Equation (13) only holds in a simplified model in which there is no price discrimination and in which all students are loan takers. Hence, in our full model, $\alpha_{w1} - \alpha_{w0}$ must be jointly estimated with the remaining parameters, in particular the ones governing the allocation of financial aid (ρ^L and ρ^D). Therefore, we aggregate the discontinuities across all degrees and match the following moment in the data:

$$\log \left[\sum_{j \neq 0} \left(\sum_{i | 0 < r_i - \bar{r}_{jt} < bw^r} s_{ij} \right) \right] - \log \left[\sum_{j \neq 0} \left(\sum_{i | -bw^r < r_i - \bar{r}_{jt} < 0} s_{ij} \right) \right],$$

where bw^r is a small bandwidth, 10 points (0.14 standard deviations) in our baseline specification.

6.2.2. Financial aid targeting: In this section, we describe how we estimate the parameters determining which students receive financial aid, for both private (discounts) and public (loans) aid. The challenge is that we only observe the financial aid status of the realized student-degree pairs; that is, we do not know whether students would receive financial aid if they were enrolled in a different degree.

Moreover, receiving financial aid changes students’ enrollment decisions. This means that students who receive aid are overrepresented among the ones who enroll in higher education. Therefore, extrapolating financial aid availability from the observed to the unobserved student-degree pairs would lead to an upward bias.

We overcome this issue by computing the magnitude of the bias using students’ preference parameters, estimated in Section 6.2.1; that is, we know how much having financial aid changes students’

enrollment decisions. Therefore, we can calculate how overrepresented students with access to financial aid are among the ones who enroll.

Let us first consider the estimation of loan take up. A student receives a government loan if the student is a loan taker ($B_i^L = 1$) and scores above the eligibility threshold for the chosen degree ($r_i \geq \bar{r}_{jt}$). From the Bayes rule, the following identity holds:

$$\frac{\rho^L(\mathbf{x}_i)}{1 - \rho^L(\mathbf{x}_i)} = \frac{\overbrace{P(B_i = 1 | d_i^a, \mathbf{x}_i)}^{\text{observed}}}{\overbrace{P(B_i = 0 | d_i^a, \mathbf{x}_i)}^{\text{observed}}} \cdot \frac{\overbrace{P(d_i^a | B_i = 0, \mathbf{x}_i)}^{\text{selection}}}{\overbrace{P(d_i^a | B_i = 1, \mathbf{x}_i)}^{\text{selection}}} \quad (14)$$

where $\rho^L(\mathbf{x}_i) \equiv P(B_i = 1 | \mathbf{x}_i)$ —as defined in Equation (12)—, and d_i^a indicates whether students are enrolled in degrees for which their scores are high enough to get loans. Formally:

$$d_i^a = \sum_{j \in \mathcal{J}} d_{ij} \cdot \mathbb{1}(r_i \geq \bar{r}_{jt}) .$$

Given the parameters of students' utility function, we can calculate the probability of each student enrolling in each degree, with and without a loan; hence we can compute the selection term in Equation (14). Therefore, ρ_L can be directly recovered from this equation.

We follow a similar procedure to estimate students' propensity to receive discounts (B_i^D). However, there is an important difference relative to the allocation of loans. The allocation of discounts depends on the parameters of the targeting function ($\rho^D(\mathbf{x}_i) = \beta^D \mathbf{x}_i$), but also on degrees' chosen discount threshold (\bar{D}_{jt}). Since \bar{D}_{jt} is not observed, it also must be estimated.

Using the Bayes rule and the logistic functional form of η_i^D , we have that:

$$P(\overbrace{D_{ij} = 1 | d_{ij} = 1}^{\text{observed}}) = \frac{e^{\beta^D \mathbf{x}_i - \bar{D}_{jt}}}{1 + e^{\beta^D \mathbf{x}_i - \bar{D}_{jt}}} \cdot \frac{\overbrace{P(d_{ij} = 1 | D_{ij} = 1)}^{\text{selection}}}{P(d_{ij} = 1)} \quad (15)$$

where $D_{ij} \equiv (B_i^D \geq \bar{D}_{jt})$ indicates whether student i would receive a discount in degree j . Note that D_{ij} is observed for the degree in which the student is enrolled ($d_{ij} = 1$).

Given the parameters of students' utility function, we can calculate the probability of each student enrolling in each degree, with and without a discount; hence we can compute the selection term in Equation (15). We then estimate β^D and \bar{D}_{jt} from (15) by maximum likelihood:

$$\{\beta^{D*}, \{\bar{D}_{jt}^*\}_{jt}\} = \arg \max_{\{\beta^D, \{\bar{D}_{jt}\}_{jt}\}} \prod_{i | j_i \neq 0} P(D_{ij_i} = 1 | d_{ij_i} = 1)^{D_{ij_i}} \cdot P(D_{ij_i} = 0 | d_{ij_i} = 1)^{1 - D_{ij_i}} ,$$

where j_i is the degree where student i is enrolled.

6.3 Estimation: Supply

We now describe how we estimate the supply-side parameters; that is, the parameters of the marginal cost function. To recover colleges' unobserved cost shifters, we follow the standard practice in the literature (Berry et al., 1995; Nevo, 2001; Backus et al., 2021) and invert colleges' first order conditions. More specifically, we estimate c_{jt} , κ_{jt} , and ω_{ft} from the first order conditions with respect to p_{jt}^F , p_{jt}^D , and \bar{D}_{ft} , presented in Equations (7), (8), and (9).

The remaining supply parameter to be estimated is ψ , which measures how costly it is for colleges to deviate from their exogenous propensity to offer discounts. Motivated by the event study discussed in Section 4.2 and shown in Figure 4, we estimate ψ by leveraging the increase in discounts that occurred in response to the reduction in loan availability. Intuitively, degrees that had more students with loans before the policy change faced a larger demand shock and hence offered more discounts afterwards. The magnitude of this response is driven by ψ . If ψ is large, degrees' pricing policy is mostly driven by their exogenous propensity of offering discounts (ω_{ft}) and the response will be muted. Formally, we impose the following moment condition:

$$E[N_{f2012}^L \cdot \Delta\omega_f] = 0 , \quad (16)$$

where N_{f2012}^L is the number of students with loans in college f in 2012. Equation (16) means that colleges that had more students with loans increased their discounts in response to the policy change, and not because of a coincidental change in unobserved shocks ($\Delta\omega_f$). In Section 4.2, we provided evidence supporting this assumption. We implement this identification strategy by finding ψ such that Equation (16) holds.

6.4 Parameter estimates

We now present the results of our structural estimation. The estimated parameters are in Table 4. Below we discuss each of the estimates, beginning with the parameters of the utility function.

Figure 5a shows the implied price elasticities. Low-income students are more sensitive to prices, with a median elasticity without loans of -5.93, compared to -1.40 for high-income students. Notice that these estimates are much larger than the ones obtained by 2SLS (Table 1). The reasons are twofold. First, the 2SLS estimates assume all students pay the same price for a given degree. Hence, the price variable used in the 2SLS estimation has measurement error, leading to a downward bias in the estimated elasticity. Second, the 2SLS does not separate between the students with and without loans. Indeed, our structural estimates show that students become substantially less price elastic when using loans, with a median elasticity of -2.74 and -0.50 for high- and low-income students, respectively.

Some remarks about these estimates. First, there is a vast difference between the price elasticities of high- and low-income students, the low-income ones being almost four times more elastic. Second,

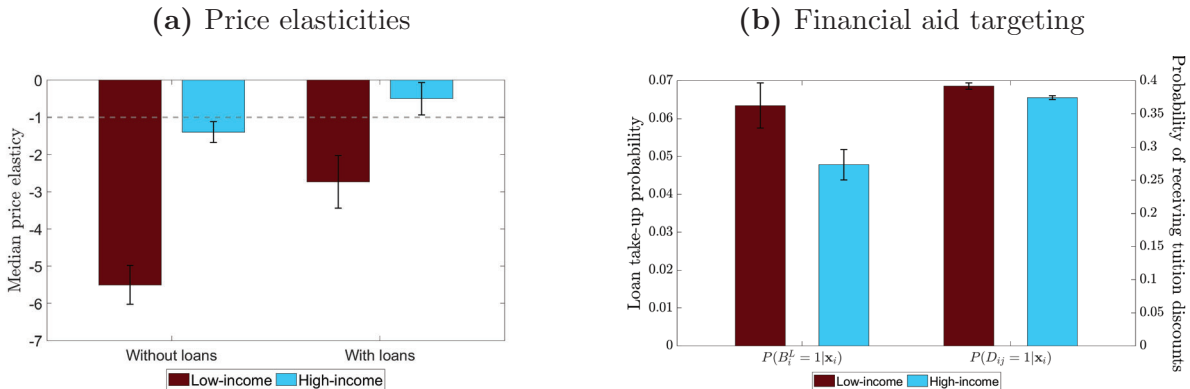
Table 4: Estimated parameters

	Demand		Supply		
	Low income	High income		Pre	Post
α_{w0}	-7.79*** (0.44)	-1.98*** (0.24)	ψ	0.025*** (0.007)	
α_{w1}	-3.94*** (0.60)	-0.72* (0.41)	c_{jt} (average)	0.68 [0.49]	0.70 [0.65]
$\beta_w^r \times 10^4$	1.18*** (0.15)	0.22* (0.12)	κ_{jt} (average)	0.01 [0.11]	0.11 [0.18]

Notes: Prices, c_{jt} , and κ_{jt} are in 1000 BRL. Standard deviations of marginal costs across degrees are presented in brackets. Standard errors of estimated parameters are presented in parentheses and were estimated using bootstrap (details in Appendix H). * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$.

even though loans do make students substantially less elastic, the baseline difference is so large between income groups that low-income students with loans are still more elastic than high-income students without. Third, these elasticities imply an average mark-up—defined as the ratio of price minus marginal cost to price—of 30%.

Figure 5: Parameter estimates: demand



Notes: Panel (a) shows the median price elasticities implied by the estimated demand parameters. Panel (b) shows the probability of a student being a loan taker [$P(B_i^L = 1 | \mathbf{x}_i)$] and receiving a tuition discount [$P(D_{ij} = 1 | \mathbf{x}_i)$]. Discount probabilities are the average across degrees. Error bars are 95% confidence intervals estimated using bootstrap (details in Appendix H).

The estimated score parameters (β^r) implies that differences in scores are responsible for 47% of the variation in degrees' mean utility for low-income students, and 9% for high-income students.

We now move to the parameters governing the allocation of financial aid. That is, the probability that students receive loans and/or discounts given their characteristics. Appendix Table A.1 has the full list of estimated parameters. Figure 5b shows the average estimated probability of being a loan taker and of receiving a discount for both high- and low-income students. Low-income students are more likely to be loan takers: 6.4%, compared to 4.8% for the high-income ones. Low-income students are also more likely to receive discounts: 39%, compared to 37% for the high-income ones.

Note that the estimated differences between high- and low-income is much smaller than the observed differences in financial aid usage: 25% of the low-income students receive loans, compared to 9% of the high-income ones; and 54% of the low-income students receive discounts, compared to 39% of the high-income ones. Students self-select into degrees in which they have financial aid available and low-income students, being more price sensitive, even more so. That is why the observed difference in financial aid usage is larger than the underlying difference in the structural parameters governing the allocation of financial aid.

Finally, we present the estimated supply-side parameters. Appendix Figure A.5 shows the distribution of marginal costs. All the estimated marginal costs are above zero. Moreover, colleges must pay a cost (ψ) to deviate from their exogenous propensity of offering discounts (ω_{ft}). The estimated ψ implies that changing the number of students who receives discounts by 10% increases marginal costs by 6.3%, on average.

7. Counterfactuals

In this section, we use our model to study the equilibrium effects of student loans. In Section 7.1, we assess the fit of our model, leveraging the observed responses to the 2015 policy change. In Section 7.2, we investigate the mechanisms driving the observed effects of the loan program. For this purpose, we decompose the effects into partial and general equilibrium responses and further decompose general equilibrium responses into the direct and composition effects. In Section 7.3, we leverage the findings of Section 7.2 to propose a policy that results in better outcomes. More specifically, we show that an alternative policy that gives loans only to low-income students strengthens the composition effect, leading to lower prices and higher enrollment. In Section 7.4, we investigate the welfare implications of different policy designs. Finally, in Section 7.5, we show how the importance of the composition effects depends on the characteristics of the market, in particular, the extent of price discrimination.

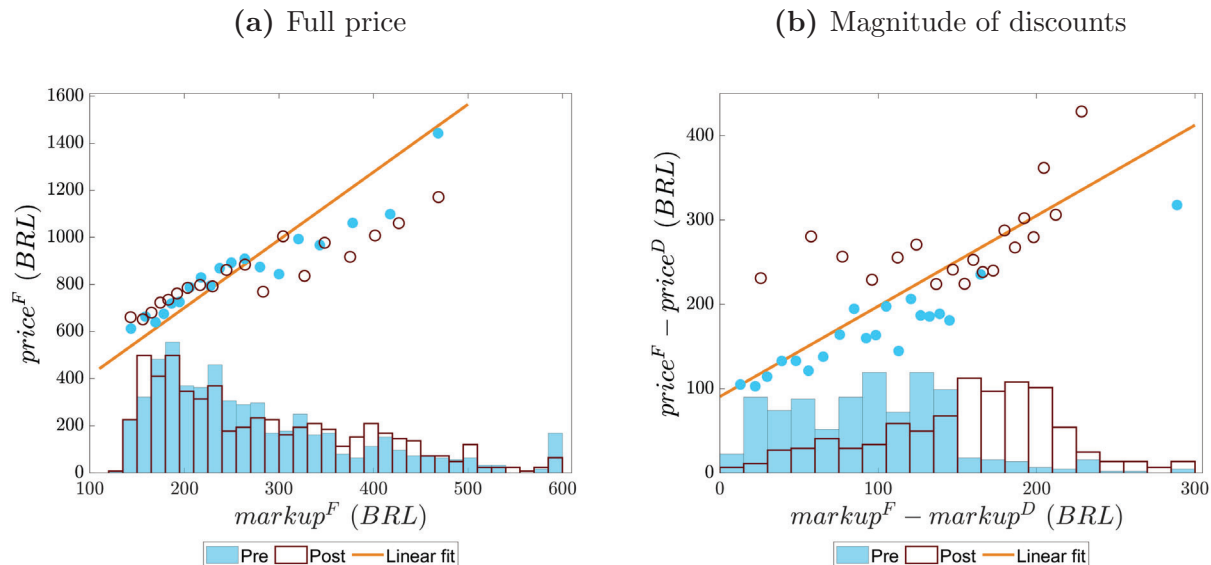
7.1 Model fit: The effects of the 2015 policy change

We now assess the fit of our model by evaluating how well it explains the observed responses to the 2015 policy change. First, we consider the variation of prices and discounts. Second, we discuss quantities: enrollment and financial aid allocation.

We begin by investigating how much of the variation in prices is explained by the model’s endogenous responses. According to colleges’ first order conditions—Equations (7) and (8)—prices are given by the sum of marginal costs and markups. The former is exogenously determined, whereas the latter is a function of students’ price elasticities and represents the model’s endogenous responses.

Figure 6 shows the relationship between prices and markups. We see an approximately linear pattern, with markups explaining 6% and 22% of the variation in full prices and discounts, respec-

Figure 6: Prices and markups



Notes: This figure shows the relationship between prices and markups. Each observation is a degree-year. The blue dots and bars represents observations from 2014 (pre-policy) and the red ones from 2016 (post-policy). The bars show the distribution of markups. The dots are a binned scatter plot with equally sized dots. The linear fit is calculated with observations from both years.

tively, the remainder being attributed to variation in costs. Hence, the overall variation in prices is mostly driven by exogenous factors.

We now focus on the change in prices before and after the policy change. On the one hand, Figure 6a shows that, after the policy change, there was a small increase in full prices (1.4%, on average), that is fully driven by changes in markups. On the other hand, Figure 6b shows much larger increases in the magnitude of discounts (67%, on average). Most of this increase (53%) is driven by changes in markups. These results indicate that our model effectively describes how equilibrium prices respond to changes in the student loan policy.

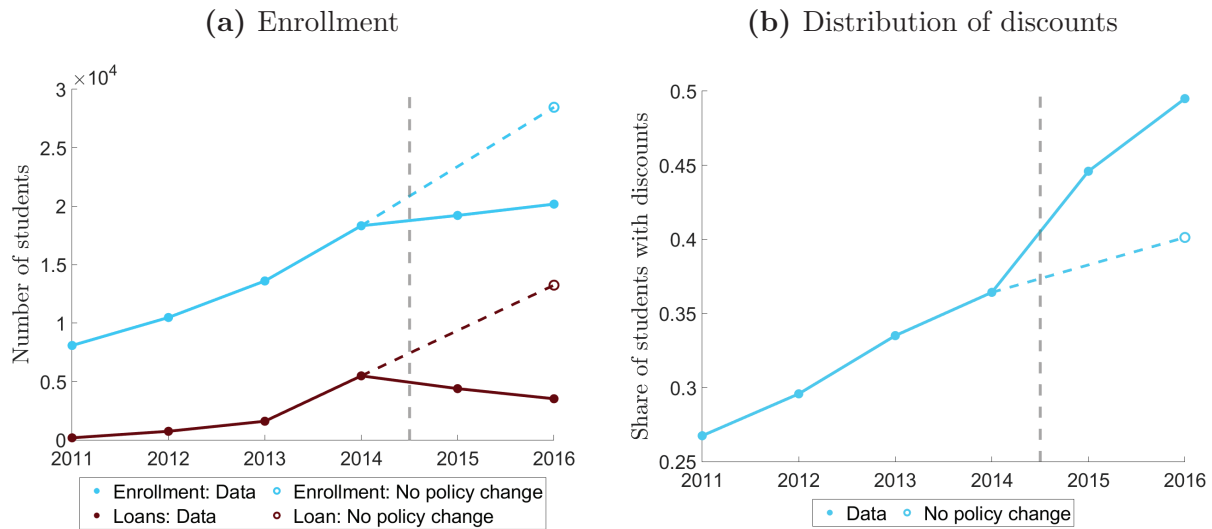
Next, we assess how well our model predicts enrollment and financial aid allocation under different loan policies. For this purpose, we use the model to predict what would have happened in 2016 if the rules of the loan program had not changed and compare results with the observed trends. More precisely, we compute a counterfactual equilibrium with the parameters governing the allocation of loans at their 2014 values, and all the time-varying shocks at their 2016 values:

$$\begin{aligned}
 \{\tilde{\rho}^L, \{\tilde{r}_j\}_{j \in \mathcal{J}}\} &= \{\rho_{2014}^L, \{r_{j2014}\}_{j \in \mathcal{J}}\}, \\
 \{\{\tilde{c}_j, \tilde{\kappa}_j\}_{j \in \mathcal{J}_f}, \tilde{\omega}_f\}_{f \in \mathcal{F}} &= \{\{\bar{c}_{j2016}, \kappa_{j2016}\}_{j \in \mathcal{J}_f}, \omega_{f2016}\}_{f \in \mathcal{F}}, \\
 \{\tilde{\xi}_{ij}\}_{I \in \mathcal{I}_{2016}, j \in \mathcal{J}} &= \{\xi_{ij2016}\}_{I \in \mathcal{I}_{2016}, j \in \mathcal{J}},
 \end{aligned} \tag{17}$$

where tildes (\sim) represent counterfactual parameters.

The results are in Figure 7. In Panel (a), we see that, before the policy change, enrollment was in a strong upward trend, from 8,000 in 2011 to 18,000 in 2014. The number of FIES loans taken by these students followed a similar trend, from 200 in 2011 to 5,400 in 2014. In 2015, the more restrictive rules of the FIES program broke this trend resulting in the number of new loan contracts dropping 35% and total enrollment remaining stable between 2014 and 2016. Panel (b) shows the trends in the share of students with discounts. We see a slight upward trend until 2014 then followed by a sharp increase between 2014 and 2016.

Figure 7: Model fit: The effects of the 2015 policy change



Notes: This figure shows the estimated effects of the reform in the federal loan program that occurred in 2015. The solid lines show the observed trends. The dashed lines show a counterfactual prediction of what would have happened in 2016 if the loan program kept the rules it had in 2014, as defined in Equation (17).

Our counterfactual simulation indicates that if the loan program rules had not changed, enrollment, loans, and discounts would have kept approximately the same trend seen until 2014. Indeed, compared to what we observe in 2016, total enrollment would have been 41% higher, new loans contracts 3.7 times higher, and the share of students with discounts 19% lower. Moreover, the cost of the loan program to the federal government would have been nine times larger.

Note that we estimate the model only with 2014 and 2016 data and use no information about previous trends. Hence, the fact that the counterfactual equilibrium fits these trends indicates that our model effectively captures the mechanisms through which loan availability affects the higher education market. Moreover, these results also suggest that our estimation strategy recovers the underlying shocks governing the trends in this market. For example, the overall increase in the demand for higher education, captured in the model by changes in the demand shock (ξ_{jt}); and the growth of degree-search platforms, which reduced the administrative costs of offering discounts (κ_{jt}).

Taken together, the results in this section indicate that our model is well suited to describe how equilibrium prices and quantities change under different policy designs. Next, we use our framework to compare alternative policy designs in terms of enrollment and student welfare.

7.2 Decomposing the direct and composition effects

This section estimates the equilibrium effects of the current Brazilian student loan policy and decomposes these impacts into three components: partial equilibrium responses, direct price effects, and composition price effects. Decomposing the effects of loans into these three mechanisms allows us to understand better how different policies lead to different results and how to design a policy that achieves the desired outcomes. For example, in Section 7.3, we build upon the results of this section to propose an alternative loan policy that results in higher enrollment.

To implement the decomposition, we compare the outcomes of four different equilibria. First, an equilibrium without loans. Second, we introduce loans but shut down price responses. Third, we allow for the direct effect of loans on prices. Fourth, we also allow for the composition effect of loans on prices, which brings us back to the outcomes observed in the data. We now describe how we compute each of these counterfactual equilibria. In all of them, we keep demand shocks and marginal costs at their 2016 levels:

$$\begin{aligned} \left\{ \{ \tilde{c}_j, \tilde{\kappa}_j \}_{j \in \mathcal{J}_f}, \tilde{\omega}_f \right\}_{f \in \mathcal{F}} &= \left\{ \{ \bar{c}_{j2016}, \kappa_{j2016} \}_{j \in \mathcal{J}_f}, \omega_{f2016} \right\}_{f \in \mathcal{F}} , \\ \{ \tilde{\xi}_{ij} \}_{I \in \mathcal{I}_{2016}, j \in \mathcal{J}} &= \{ \xi_{ij2016} \}_{I \in \mathcal{I}_{2016}, j \in \mathcal{J}} . \end{aligned} \tag{18}$$

In the first scenario, the equilibrium without loans is a counterfactual equilibrium in which there are no loan takers:

$$\tilde{\rho}^L(\mathbf{x}) = 0, \forall \mathbf{x}. \tag{19}$$

In the second scenario, to compute partial equilibrium responses, we allocate loans as in the data ($\tilde{\rho}^L = \rho_{2016}^L$), but hold prices fixed at their level in the equilibrium without loans.

In the third scenario, we allocate loans as in the data ($\tilde{\rho}^L = \rho_{2016}^L$) and allow for price responses, but shut down the composition effect. The decomposition derived in our simplified conceptual framework (Section 2) does not directly apply to the structural model (Section 5). Hence, we now introduce an alternative definition of the composition effect. For clarity, we discuss here the case of a small single-product college without discounts. Appendix G shows the general case and corresponding

proofs. We have that equilibrium markups are equal to:

$$markup_j \approx \frac{1}{\underbrace{P(w_i = Low|d_{ij} = 1)}_{\text{composition effect}} \cdot \alpha_j^{Low\ Income} + \underbrace{P(w_i = High|d_{ij} = 1)}_{\text{composition effect}} \cdot \alpha_j^{High\ Income}}, \quad (20)$$

where

$$\begin{aligned} \alpha_j^{Low\ Income} &= \overbrace{P(L_{ij} = 0|w_i = Low \ \& \ d_{ij} = 1)}^{\text{direct effect}} \cdot \alpha_{Low,0} + \overbrace{P(L_{ij} = 1|w_i = Low \ \& \ d_{ij} = 1)}^{\text{direct effect}} \cdot \alpha_{Low,1}, \\ \alpha_j^{High\ Income} &= \overbrace{P(L_{ij} = 0|w_i = High \ \& \ d_{ij} = 1)}^{\text{direct effect}} \cdot \alpha_{High,0} + \overbrace{P(L_{ij} = 1|w_i = High \ \& \ d_{ij} = 1)}^{\text{direct effect}} \cdot \alpha_{High,1}. \end{aligned}$$

Equation (20) shows that degree j 's markup is a function of the average price sensitivity of its students. The direct effect is that increasing the share of the students in a degree who have loans increases the markup because loans reduce price sensitivity. The composition effect is that when loans increase the share of low-income students in degree j markup goes down since low-income students are more price sensitive.

To allow for the direct effect and shut down the composition effect, we compute prices using Equation (20) and input the share of low-income students $[P(w_i = Low|d_{ij} = 1)]$ from the equilibrium without loan and the share of students with loans $[P(L_{ij} = 0|w_i = Low \ \& \ d_{ij} = 1)$ and $P(L_{ij} = 0|w_i = High \ \& \ d_{ij} = 1)]$ from the data.

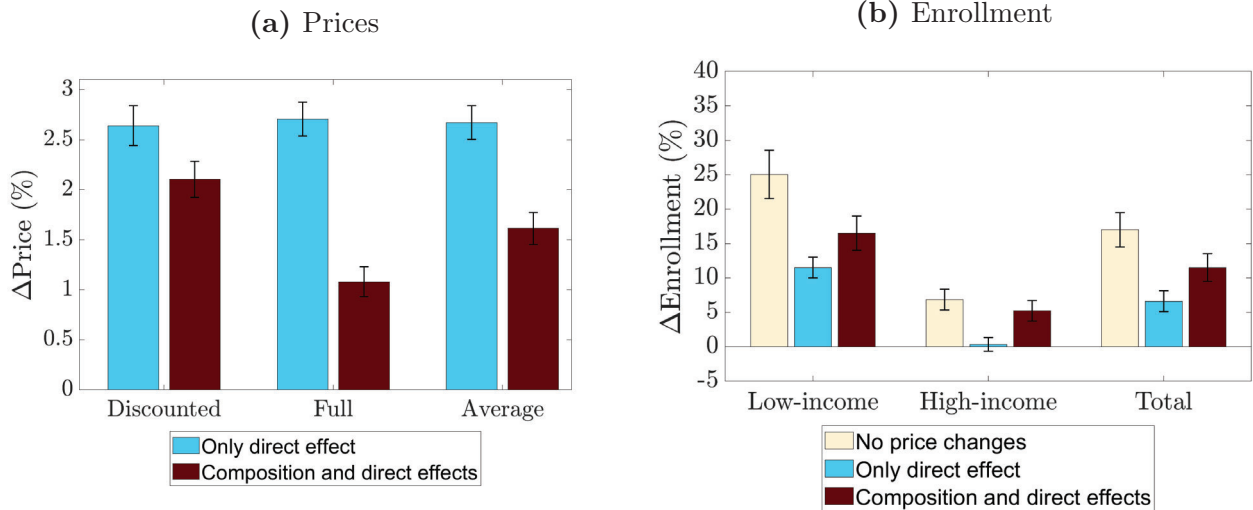
Finally, in the fourth scenario, we allocate loans as in the data ($\tilde{\rho}^L = \rho_{2016}^L$), and allow for both direct and composition price effects. This brings us back to the equilibrium observed in the data.

The results are in Figure 8. It shows equilibrium outcomes in each of the three equilibria with loans described above, compared to the equilibrium without loans. Panel (a) shows the effects on prices. On average, the direct effect raises prices by 2.7%, whereas the composition effect reduces them by 1.1%. The composition effect is stronger for full prices (-1.6%) than discounted prices (-0.5%). The reason is that loans change the composition of students paying full price more than those paying a discounted price.

Panel (b) shows effects on enrollment. In partial equilibrium, the loan program increases total enrollment by 17.0%, relative to the counterfactual without loans. The price changes induced by the direct effect reduce enrollment by 10.4%, and the changes induced by the composition effect increase enrollment by 4.9%. The net effect of these three forces is an 11.5% increase in total enrollment.

In summary, we find that both the direct and the composition effects are responsible for prices responses that lead to substantial changes in total enrollment. In particular, price reductions induced by the composition effect are responsible for 40% of the enrollment gains of the current Brazilian loan program.

Figure 8: Equilibrium effects of the current Brazilian student loan program



Notes: This figure shows the equilibrium effects of the current Brazilian student loan program. In both panels, the y-axis shows differences relative to an equilibrium without loans. In Panel (a), discounted prices are weighted by the number of students with discounts in each degree, full prices by the number of students paying full price, and the average price by the total number of students. All weights are taken from the counterfactual without loans. Error bars are 95% confidence intervals estimated using bootstrap (details in Appendix H).

7.3 Alternative policy designs

This section builds upon our previous findings to propose an alternative allocation of loans: giving loans only to low-income students. The reason for focusing on this specific policy is the following. We have seen that low-income students are more price-sensitive than high-income ones (Figure 5a). Moreover, loans reduce students' price sensitivity, albeit by less than the baseline difference between the groups. That is, low-income students with loans are still more price-sensitive than high-income students without. These patterns suggest that targeting loans to low-income students can strengthen the composition effect and lower prices, leading to higher enrollment. Motivated by this insight, we now discuss the equilibrium outcomes of an alternative policy design that gives loans only to low-income students.

More specifically, we define the counterfactual policy as:

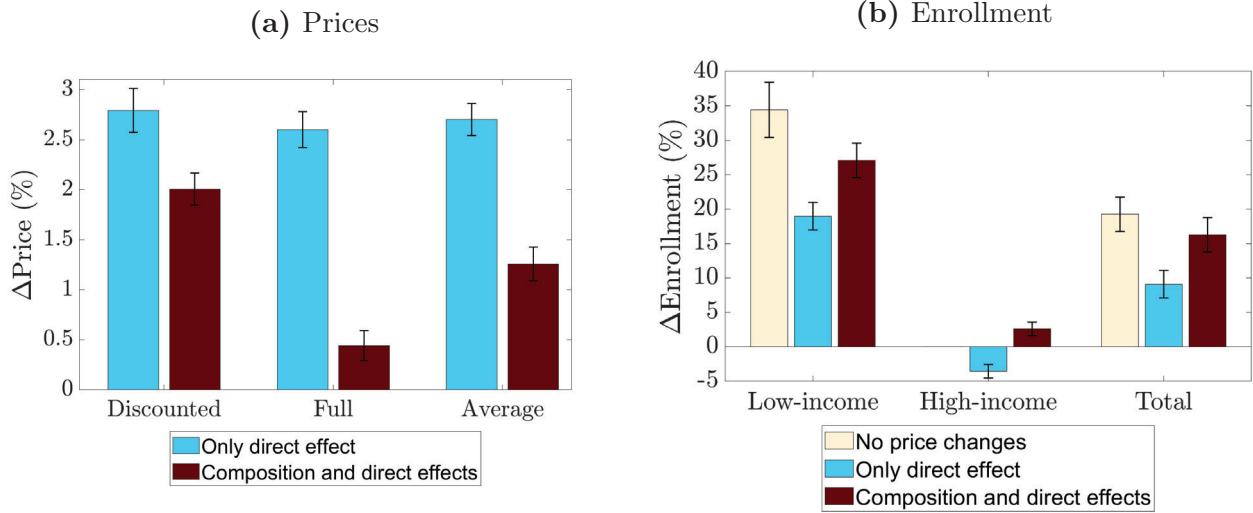
$$\tilde{\rho}^L(\mathbf{x}_i) = \begin{cases} 0 & , \text{ if } w_i = \text{high income} \\ \rho_{2016}^L(\mathbf{x}_i) & , \text{ if } w_i = \text{low income} \end{cases} . \quad (21)$$

Unobserved shocks are kept unchanged at their 2016 levels in the counterfactual, both on the demand (ξ_{ijt}) and on the supply ($\bar{c}_{jt}, \kappa_{jt}, \omega_{ft}$) sides, as defined in Equation (18). Score cutoffs (\tilde{r}_{jt}) are adjusted proportionally for all degrees to keep the total budget of the program the same it was

in 2016.

Figure 9 presents the effects of the alternative policy. All outcomes are shown as changes with respect to the equilibrium without any loans. We decompose direct and composition effects following the same procedure described in Section (7.2).

Figure 9: Alternative policy design: loans only to low-income students



Notes: This figure shows the equilibrium effects of an alternative policy that gives loans only to low-income students. In both panels, the y-axis is the difference between outcomes in an equilibrium with loans only to low-income students and an equilibrium without loans. In Panel (a), discounted prices are weighted by the number of students with discounts in each degree, full prices by the number of students paying full price, and the average price by the total number of students. All weights are taken from the counterfactual without loans. Error bars are 95% confidence intervals estimated using bootstrap (details in Appendix H).

Panel (a) shows the effects of loans on prices under the alternative policy. On the one hand, we see that the direct effect is very similar to the observed in the current policy. Both discounted and full prices increase, on average, by around 2.7%. On the other hand, the composition effect is stronger in the alternative policy and reduces the average price by 1.5%, 40% more than the current policy.

Panel (9b) shows effects on enrollment. We see that the alternative policy increases total enrollment by 16.3%, compared to 11.5% in the current one, a 41% difference. Half of the difference in enrollment between the two policies comes from partial equilibrium responses. Low-income students are more affected by loans, at least in the extensive margin; hence targeting this group is more effective in increasing total enrollment. The composition effect is responsible for the other half. The direct effect is very similar in both policies and has negligible impacts on enrollment.

In addition to the effects on total enrollment, the two policies have important distributional differences. As expected, the alternative policy benefits more low-income students. On the one hand, the current policy increases the enrollment of low-income students by 16%, the alternative by

27%. On the other hand, the current policy increases the enrollment of high-income students by 5.2%, the alternative by 2.6%.

Notice that the alternative policy still increases the total enrollment of high-income students, even though this group does not receive any loans under this design. The reason is that loans increase the share of low-income students more in more expensive degrees, initially dominated by high-income students. Hence, the composition effect is stronger in these degrees, which leads to lower prices. In Section (7.4), we will show that these lower prices lead to large welfare gains to inframarginal high-income students.

Despite the considerable gains of giving loans only to low-income students, perfectly targeting them might be challenging in practice due to misreporting and fraud. A feasible alternative is to restrict loans to public high school students, who are on average poorer. In our sample, 82% of public high school students are low-income compared to 42% in private high schools. Moreover, implementation would be simple since the government already knows where each student was enrolled in high school. Indeed, the Ministry of Education already reserves spots in public universities for students coming from public high schools (Mello, 2021; Otero et al., 2021), and the same system could be used in the allocation of loans.

Appendix Figure A.6 shows the counterfactual outcomes of an alternative policy that gives loans only to public high school students. Indeed, the patterns are very similar to those obtained by targeting low-income students, but the magnitudes are smaller. For example, targeting public high schools obtains 28% of the enrollment gains of perfectly targeting low-income students.

In summary, we take the following lessons from the results in this section. First, restricting loans to low-income students would lead to substantial gains in terms of total enrollment. Second, half of these gains are due to supply-side responses from a stronger composition effect in the alternative policy.

7.4 Welfare implications

The outcomes of alternative policies differ not only in terms of total enrollment, but also in the allocation of students across degrees. Hence, we need a framework for evaluating outcomes that considers the distribution of students across degrees.

Our demand model also gives us such a framework: comparing total student welfare in each allocation. We define students' *experience utility* U_{ij}^* (Kahneman, 1994; Allcott, 2013), which represents welfare, as:

$$U_{ij}^* = \beta_i^h \mathbf{h}_j + \alpha_i^* p_{ijt} + \xi_{ijt} ,$$

where α_i^* is the welfare-relevant price sensitivity of student i . We measure student welfare as total

consumer surplus, which is given by:

$$CS = \int_{i \in \mathcal{I}} \frac{1}{\alpha_i^*} \sum_{j \in \mathcal{J}} d_{ij} U_{ij}^* di . \quad (22)$$

Under the parameterization defined in Equation 11, consumer surplus can be computed as (Train, 2015):

$$CS = \int_{i \in \mathcal{I}} \frac{1}{\alpha_i^*} \left[\underbrace{\log \left(\sum_{j \in \mathcal{J}} e^{V_{ij}} \right)}_{\text{behavior welfare}} + \overbrace{\sum_{j \in \mathcal{J}} s_{ij} (\alpha_i^* - \alpha_{iL_{ij}})}^{\text{difference between experience and behavior welfare}} \right] di , \quad (23)$$

where

$$V_{ij} \equiv \beta_i^h \mathbf{h}_j + \alpha_{iL_{ij}} p_{ijt} + \gamma_{w_{ij}} + \gamma_{w_{it}} + \xi_{w_{ijt}} .$$

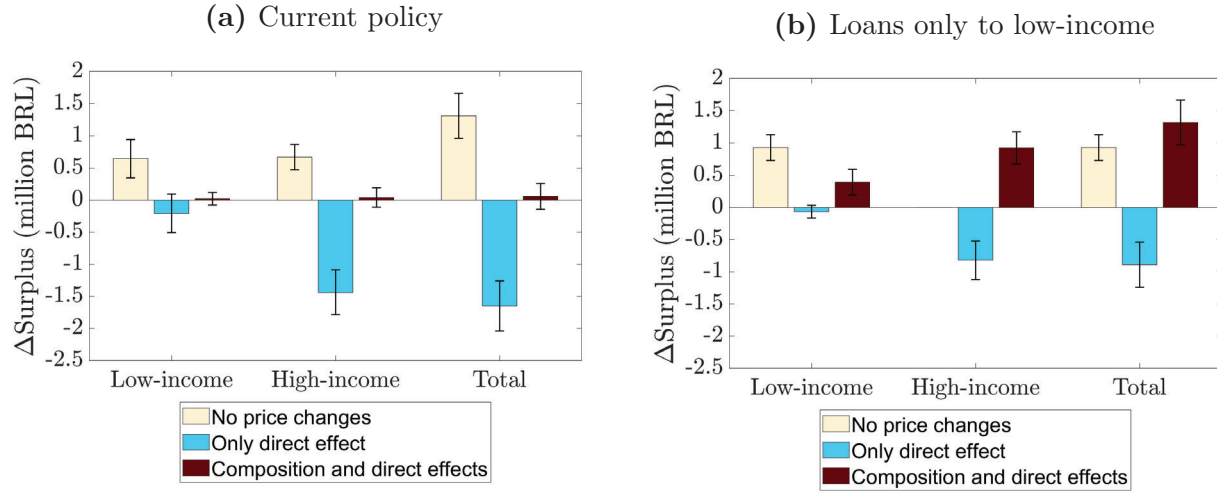
Note that students' *experience utility* might be different from their *behavior utility*, described in Equation (4). Our specification of students' *behavior utility* allows their price sensitivity to depend on whether they have a loan, raising the question of which sensitivity parameter should be used in welfare calculations. Students' behavior reveals their underlying *experience utility* if there are no market imperfections. Two imperfections are particularly relevant in our context. First, some students are credit constrained. Second, students might have biased expectations regarding the returns of different degrees, resulting in price elasticities that are too low or too high.

Our baseline specification takes the extreme interpretation that student loans solve all the existing market failures. Under this assumption, the choices of students with loans reveal the parameters of their underlying *experience utility*. Note that this assumption means that the current subsidy in the loans' interest rate perfectly corrects any differences between students' baseline price elasticities and the one that reflects the social value of higher education. The logic behind this assumption is that elected policymakers chose the subsidy rate to reflect their view regarding the social returns of higher education. However, we also discuss how our results change under alternative assumptions.

Figure 10 shows the effects of different policy designs on consumer surplus. Panel (a) presents the current policy, and Panel (b) the low-income only alternative policy, defined in Equation (21). The figure shows changes in consumer surplus, relative to the equilibrium without loans, defined in Equation (19). The effects are decomposed into partial equilibrium, direct price effects, and composition price effects following the procedure developed in Section 7.2.

In partial equilibrium, low-income students gain more from the alternative policy and high-income students from the current one. In aggregate, the alternative policy results in a more modest increase in welfare than the current one. The picture is entirely different in general equilibrium. Supply-side

Figure 10: Consumer surplus under different policy designs



Notes: This figure shows total consumer surplus under two alternative loan policies. Panel (a) shows consumer surplus in the current policy, and Panel (b) in an alternative policy that gives loans only to low-income students. In both panels, the y-axis shows changes relative to an equilibrium without loans, defined in Equation (19). Error bars are 95% confidence intervals estimated using bootstrap (details in Appendix H).

responses (lower prices) lead to large welfare gains in the alternative policy, and both income groups are better off than in the current policy. In the current policy, supply-side responses nearly crowd out all the partial equilibrium welfare gains, whereas supply-side responses amplify the gains in the alternative policy. Consequently, consumer surplus increases 23 times more in the alternative policy. To benchmark the magnitudes of these effects, let us compare them with the program’s total budget. In the current policy, for each \$1.00 invested in the program, consumer surplus increases \$0.31 in partial equilibrium. However, supply-side responses reduce it by \$0.30, resulting in modest and not statistically significant \$0.01 net gain. In contrast, consumer surplus increases \$0.22 in partial equilibrium in the alternative policy, and supply-side responses further increase it by \$0.09, resulting in a \$0.31 net gain.

The alternative policy’s price responses produce better welfare consequences for two reasons. First, the composition effect is stronger under the alternative policy, which leads to lower prices. Second, the price increases associated with the direct effect reduce welfare less in the alternative policy, even though the average direct effect on prices is the same in both policies, as discussed in Section 7.3. The reason is that the alternative policy’s price increases are concentrated in degrees with more low-income students. Since these degrees are also cheaper, the average price increase, in dollars, is 35% larger in current policy, even though the average *percent* increase is similar in the two policies.

As discussed above, our baseline results assume that the loan program corrects existing market failures and, hence, the choices of students with loans reveal the parameters of their underlying

experience utility. An alternative assumption is that there are no market failures and, hence, students' choices without loans reveal the parameters of their underlying experience utility. Appendix I shows how our results change for a series of alternative assumptions between these two extremes. We find that the low-income-only alternative policy results in higher welfare under any considered assumption. The reason is that the alternative policy results in lower prices, which always benefits students.

The main takeaway of this section is that different policy designs might have substantial welfare consequences beyond their effects on total enrollment. In particular, policies that better target low-income students lead to lower prices due to a stronger composition effect, resulting in large gains in consumer surplus.

7.5 The role of price discrimination

In Sections 7.3 and 7.4, we showed how an alternative policy that gives loans only to low-income students reduces prices due to a stronger composition effect, leading to higher enrollment. This section discusses how these results depend on the extent of price discrimination. For this purpose, we show how the current and the alternative policies compare under perfect price discrimination. There is no composition effect in this scenario, and all price changes come from the direct effect.

More specifically, we simulate counterfactual equilibria where only high-income students pay full prices, and only low-income students pay discounted prices. Hence, colleges consider only the price elasticity of low-income students when setting their discounted price and high-income students when setting their full price. This implies that there is no composition effect under perfect price discrimination.

Formally, we model perfect price discrimination the following way:

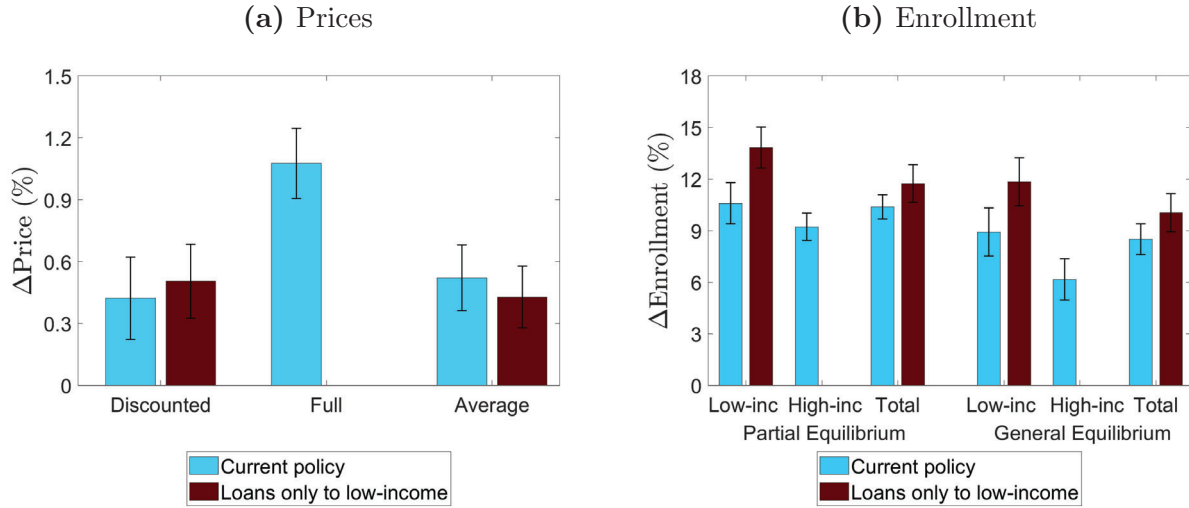
$$\tilde{\rho}^D(\mathbf{x}_i) = \begin{cases} \infty & , \text{ if } w_i = \text{low income} \\ -\infty & , \text{ if } w_i = \text{high income} \end{cases} , \quad (24)$$

where ρ^D governs students' propensity to receive discounts, as defined in Section 6.1. Unobserved shocks are kept unchanged at their 2016 levels in the counterfactual, both on the demand and the supply sides, as defined in Equation (18).

We simulate the outcomes of three different policy counterfactuals under perfect price discrimination. First, without any loans: $\tilde{\rho}^L(\mathbf{x}_i) = 0$. Second, following the current allocation of loans: $\tilde{\rho}^L(\mathbf{x}_i) = \rho_{2016}^L(\mathbf{x}_i)$. Third, an alternative policy that gives loans only to low-income students, as defined in Equation (21).

Figure 11 compares the outcomes in these three scenarios. It shows the changes in prices and enrollment in the current and the alternative loan policies compared to the equilibrium without loans. We see that under perfect price discrimination (without the composition effect), giving loans only to low-income students has much more modest impacts on equilibrium outcomes.

Figure 11: Comparing alternative policy designs under perfect price discrimination



Notes: This figure shows the equilibrium outcomes of counterfactuals with perfect price discrimination. It compares the equilibrium outcomes of two different student loan programs: the observed 2016 policy (in blue) and a counterfactual policy that gives loans only to low-income students (in red). Changes are relative to an equilibrium without any loans. The counterfactual policy is defined in Equation (21). All unobserved shocks are kept unchanged, both on the demand and on the supply sides, as defined in Equation (?). Score cutoffs (\tilde{r}_j) are adjusted proportionally for all degrees to keep the total budget of the program the same it was in 2016. In partial equilibrium, pricing policies are kept the same as in the equilibrium without loans and, in general equilibrium, they follow colleges' first order conditions. In Panel (a) prices are the average across all degrees, weighted by market shares in the equilibrium without loans. Error bars are 95% confidence intervals estimated using bootstrap (details in Appendix H).

Let us begin discussing the effects on prices. On the one hand, in Section 7.3 we have seen that the current policy raises the average price 0.4 percentage points more than the alternative policy under imperfect discrimination. Moreover, the composition effect was entirely responsible for this difference. On the other hand, there is no composition effect under perfect discrimination. As a result, Figure 11a shows that the difference between the average prices of the two policies is less than 0.1 percentage points.

Panel (b) shows how the policies compare in terms of enrollment. The gains obtained by giving loans only to low-income students are much more modest without the composition effect. On the one hand, under imperfect discrimination, the alternative policy raises total enrollment 41% more than the current policy (Section 7.3). On the other hand, under perfect discrimination, the alternative policy raises total enrollment only 18% more than the current policy (Figure 11). The difference is almost entirely due to supply-side responses. Indeed, in partial equilibrium, the enrollment gains of switching from the baseline to the alternative policy are around 13% under either perfect or imperfect discrimination.

In summary, the outcomes of different policy designs depend on the market structure. For

example, when there is a large difference in price elasticity between high- and low-income students, and colleges have limited ability to discriminate between these groups, policies that increase the market-shares of low-income students reduce prices, leading to large enrollment gains. However, this is not the case under perfect price discrimination because there is no composition effect.

8. Conclusion

In this paper, we study the equilibrium effects of subsidized student loan programs. Despite the popularity of these programs, policymakers have long been worried that they enable colleges to raise their tuition and capture a large share of the invested public funds, undermining the policy's effectiveness.

We show that student loans affect equilibrium prices in two ways. First, loans reduce students' price elasticity, leading to higher markups and prices. We refer to this mechanism as the direct effect. Second, loans increase the market share of low-income students. Since they are more price elastic, the average elasticity of the market increases, reducing markups and prices. We refer to this mechanism as the composition effect. These opposing forces imply that the net impact on prices is ambiguous and depends on how the government targets loans, a fact the previous literature has not discussed.

The magnitude of the direct and composition effects depends on three key parameters. First, the effect of loans on students' price elasticities: if students who receive loans become much less sensitive to prices, the direct effect will be strong. Second, the heterogeneity in price elasticity across students: the larger the difference in price elasticity between high- and low-income students, the stronger the composition effect. Third, the accuracy of price discrimination: in the extreme, if price discrimination is perfect, colleges charge one individualized price to each student, with no composition effect.

We investigate the empirical relevance of these forces in the context of the Brazilian higher education market. We exploit a policy change that resulted in a drastic reduction in the availability of loans and was followed by a large expansion of tuition discounts. We find that loans substantially reduce students' price elasticity, suggesting a strong direct effect. Moreover, low-income students are much more price elastic, and colleges have limited ability to price discriminate, suggesting the composition effect is also strong.

Guided by these patterns, we develop a model of the supply and demand for higher education. The goals of the model are twofold. First, to put together the different pieces of empirical evidence and estimate the net effect of student loans on prices, considering both the direct and composition effects. Second, to compare the outcomes of alternative policy designs.

We estimate our model with data from the Brazilian higher education market and use it to simulate the outcomes of alternative student loan programs. We find that increasing loan availability

to low-income students reduces prices. Indeed, an alternative design that keeps the same budget, but gives loans only to low-income students, reduces tuition costs and leads to large gains in terms of enrollment and consumer surplus. All the gains in surplus and half of the rise in enrollment come from supply-side responses.

In summary, our results show that, when there is vast heterogeneity in price sensitivity across students, financial aid policies change the composition of the market and consequently equilibrium prices. Moreover, the direction of these effects is ambiguous and depends on how the aid is targeted. The framework developed here can be used to study other markets in which targeted public subsidies coexist with private provision of goods and services, such as the health care and food industries.

There are important issues left for future research. Most importantly, we take the market structure as given in our analysis. Even though this is a reasonable approximation in the short-term window we consider, in the long run, colleges will likely respond to major reforms in student loan programs by changing their quality, offering new degrees, and even entering or exiting the market. Better understanding these long-run implications will be crucial for designing more efficient student aid programs.

References

- Allcott, H. (2013, August). The Welfare Effects of Misperceived Product Costs: Data and Calibrations from the Automobile Market. *American Economic Journal: Economic Policy* 5(3), 30–66.
- Allende, C. (2020). Competition Under Social Interactions and the Design of Education Policies. pp. 104.
- Angrist, J., D. Autor, S. Hudson, and A. Pallais (2014, December). Leveling Up: Early Results from a Randomized Evaluation of Post-Secondary Aid. Technical Report w20800, National Bureau of Economic Research, Cambridge, MA.
- Armona, L. and S. Cao (2021). Redesigning Federal Student Aid in Higher Education. *Working Paper*, 104.
- Backus, M., C. Conlon, and M. Sinkinson (2021). Common Ownership and Competition in the Ready-to-Eat Cereal Industry. *American Economic Journal: Microeconomics* 13(3).
- Bau, N. (2019). Estimating an Equilibrium Model of Horizontal Competition in Education. *Working Paper*, 65.
- Bennett, W. J. (1987, February). Our greedy colleges. *The New York Times*, 31–32.
- Berry, S., J. Levinsohn, and A. Pakes (1995, July). Automobile Prices in Market Equilibrium. *Econometrica* 63(4), 841.
- Bo, I. and R. Hakimov (2019, July). Iterative Versus Standard Deferred Acceptance: Experimental Evidence. *The Economic Journal*, uez036.
- Cellini, S. R. and C. Goldin (2014, November). Does Federal Student Aid Raise Tuition? New Evidence on For-Profit Colleges. *American Economic Journal: Economic Policy* 6(4), 174–206.
- Coey, D., B. J. Larsen, and B. C. Platt (2020, December). Discounts and Deadlines in Consumer Search. *American Economic Review* 110(12), 3748–3785.
- DellaVigna, S. and M. Gentzkow (2019, November). Uniform Pricing in U.S. Retail Chains*. *The Quarterly Journal of Economics* 134(4), 2011–2084.
- Dinerstein, M. and T. D. Smith (2021). Quantifying the Supply Response of Private Schools to Public Policies. *American Economic Review*, 80.
- Dobbin, C., C. Yannelis, N. Barahona, H. Hanson, and S. Otero (2021). Skin in the Game: College’s Financial Incentives and Student Outcomes. *Working Paper*, 37.


- Eaton, C., S. T. Howell, and C. Yannelis (2020, September). When Investor Incentives and Consumer Interests Diverge: Private Equity in Higher Education. *The Review of Financial Studies* 33(9), 4024–4060.
- Epple, D., R. Romano, S. Sarpca, H. Sieg, and M. Zaber (2019, March). Market power and price discrimination in the US market for higher education. *The RAND Journal of Economics* 50(1), 201–225.
- Epple, D., R. Romano, and H. Sieg (2006). Admission, Tuition, and Financial Aid Policies in the Market for Higher Education. *Econometrica* 74(4), 885–928.
- Fack, G. and J. Grenet (2015, April). Improving College Access and Success for Low-Income Students: Evidence from a Large Need-Based Grant Program. *American Economic Journal: Applied Economics* 7(2), 1–34.
- Ferreira, M. M. and G. Kosenok (2018, March). Charter school entry and school choice: The case of Washington, D.C. *Journal of Public Economics* 159, 160–182.
- Fillmore, I. (2016, April). Price Discrimination and Public Policy in the U.S. College Market. *Employment Research* 23(2), 5–6.
- Hausman, J. (1994, December). Valuation of New Goods under Perfect and Imperfect Competition. Technical Report w4970, National Bureau of Economic Research, Cambridge, MA.
- Hoxby, C. M. (Ed.) (2004). *College choices: the economics of where to go, when to go, and how to pay for it*. A National Bureau of Economic Research conference report. Chicago: University of Chicago Press.
- Kahneman, D. (1994). New Challenges to the Rationality Assumption. *Journal of Institutional and Theoretical Economics* 150, 20.
- Kelchen, R. (2019, December). An empirical examination of the Bennett hypothesis in law school prices. *Economics of Education Review* 73, 101915.
- Londono-Velez, J. and C. Rodriguez (2020). Upstream and Downstream Impacts of College Merit-Based Financial Aid for Low-Income Students: Ser Pilo Paga in Colombia. *American Economic Journal: Economic Policy*, 40.
- Long, B. T. (2004). How Do Financial Aid Policies Affect Colleges? The Institutional Impact of the Georgia HOPE Scholarship. *The Journal of Human Resources* 39(4), 1045.
- Lucca, D. O., T. Nadauld, and K. Shen (2018, June). Credit Supply and the Rise in College Tuition: Evidence from the Expansion in Federal Student Aid Programs. *The Review of Financial Studies*.
- Marta Ferreira, M., C. Avitabile, J. Botero Alvarez, F. Haimovich Paz, and S. Urzua (2017, May). *At a Crossroads: Higher Education in Latin America and the Caribbean*. The World Bank.

- Mello, U. (2021). Centralized Admissions, Affirmative Action and Access of Low-income Students to Higher Education. *American Economic Journal: Economic Policy*, 70.
- Neilson, C. A. (2019). Approximating the Equilibrium Effects of Informed School Choice. *Working Paper 628*, 52.
- Neilson, C. A. (2021). Targeted Vouchers, Competition Among Schools, and the Academic Achievement of Poor Students. *Working Paper*, 55.
- Nevo, A. (2001). Measuring Market Power in the Ready-to-Eat Cereal Industry. *Econometrica*.
- OECD (2014). Indicator B5: How much do tertiary students pay and what public support do they receive? *Education at a Glance 2014: OECD Indicators*.
- OECD (2019, September). *Education at a Glance 2019: OECD Indicators*. Education at a Glance. OECD.
- Otero, S., N. Barahona, and C. Dobbin (2021). Affirmative Action in Centralized College Admission Systems. *Working Paper*.
- Singell, L. D. and J. A. Stone (2007, June). For whom the Pell tolls: The response of university tuition to federal grants-in-aid. *Economics of Education Review* 26(3), 285–295.
- Singleton, J. D. (2019, July). Incentives and the Supply of Effective Charter Schools. *American Economic Review* 109(7), 2568–2612.
- Solis, A. (2017, April). Credit Access and College Enrollment. *Journal of Political Economy* 125(2), 562–622.
- Tiffany, F. G. (1998). The Competitive Use of Price Discrimination by Colleges. *Eastern Economic Journal* 24(1), 13.
- Train, K. (2015, September). Welfare calculations in discrete choice models when anticipated and experienced attributes differ: A guide with examples. *Journal of Choice Modelling* 16, 15–22.
- Turner, L. J. (2017). The Economic Incidence of Federal Student Grant Aid. pp. 41.
- Turner, N. (2012, August). Who benefits from student aid? The economic incidence of tax-based federal student aid. *Economics of Education Review* 31(4), 463–481.
- van der Klaauw, W. (2002, November). Estimating the Effect of Financial Aid Offers on College Enrollment: A Regression-Discontinuity Approach. *International Economic Review* 43(4), 1249–1287.
- Weyl, E. G. and M. Fabinger (2013, June). Pass-Through as an Economic Tool: Principles of Incidence under Imperfect Competition. *Journal of Political Economy* 121(3), 528–583.
- Zerkle, J. H. (2006). Price Discrimination. *Journal of Economics and Finance Education*, 7.

Online Appendix for:
The Equilibrium Effects of Subsidized Student Loans

A. Additional Figures and Tables

Figure A.1: The expansion of private higher education



Notes: This figure shows the number of incoming higher-education students in each year in Brazil, by type of institution (private, state, or federal). The vertical line marks a reform that allowed the entry of for-profit institutions. Source: Census of Higher Education.

Table A.1: Estimated parameters: Financial aid targeting


Income	Parents have higher education	Plan to apply to a loan	Public high school	ρ_{2014}^L	ρ_{2016}^L
4	1	0	0	0.01 (0.01)	0.00 (0.00)
2	1	0	1	0.02** (0.01)	0.02* (0.01)
2	1	0	0	0.02** (0.01)	0.02** (0.01)
4	0	0	0	0.03*** (0.01)	0.01 (0.01)
4	1	0	1	0.03*** (0.01)	0.01* (0.00)
1	1	0	0	0.03** (0.02)	0.02* (0.01)
2	0	0	0	0.03*** (0.01)	0.02** (0.01)
2	1	1	0	0.04*** (0.02)	0.03** (0.02)
3	1	0	0	0.04*** (0.01)	0.02* (0.01)
1	1	0	1	0.04*** (0.02)	0.02* (0.01)
1	1	1	0	0.04*** (0.01)	0.03** (0.02)
1	0	0	0	0.05*** (0.02)	0.03*** (0.01)
2	0	0	1	0.05*** (0.03)	0.04*** (0.02)
4	1	1	0	0.06*** (0.02)	0.02** (0.01)
3	1	0	1	0.06*** (0.02)	0.03*** (0.01)
3	0	0	0	0.07*** (0.03)	0.03** (0.02)
4	0	0	1	0.07*** (0.02)	0.02* (0.01)
2	0	1	0	0.07*** (0.04)	0.04*** (0.02)
3	1	1	0	0.07*** (0.03)	0.04*** (0.02)
1	0	1	0	0.07*** (0.04)	0.06*** (0.03)
2	1	1	1	0.08*** (0.03)	0.02* (0.01)
4	1	1	1	0.08*** (0.03)	0.03* (0.02)
1	0	0	1	0.08*** (0.04)	0.05*** (0.02)
4	0	1	0	0.09*** (0.03)	0.03*** (0.01)
2	0	1	1	0.10*** (0.05)	0.06*** (0.02)
3	0	0	1	0.10*** (0.04)	0.05*** (0.02)
3	1	1	1	0.11*** (0.04)	0.04*** (0.02)
1	1	1	1	0.11*** (0.05)	0.11*** (0.04)
1	0	1	1	0.12*** (0.05)	0.07*** (0.03)
4	0	1	1	0.12*** (0.06)	0.04*** (0.02)
3	0	1	0	0.13*** (0.06)	0.06*** (0.03)
3	0	1	1	0.21*** (0.09)	0.10*** (0.04)
				ρ^D	
1	-	-	-	0.00 (·)	
2	-	-	-	0.06 (0.08)	
3	-	-	-	-0.21*** (0.07)	
4	-	-	-	-1.33*** (0.22)	

Notes: This table presents the estimates of the parameters governing the allocation of financial aid, ρ^L and ρ^D . “Income” is the students family income bin, 4 being the highest and 1 the lower. “Parents have higher education” indicates whether at least one of the students parents has higher education. “Plan to apply to a loan” is the answer to a survey question asking whether the student plans to apply to a federal loan. “Public high school” indicates whether the student attended a public high school. Standard errors are presented in parentheses and were estimated using bootstrap (details in Appendix H). * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$.

Figure A.2: The distribution of FIES loans, by income

(a) 2014


(b) 2016



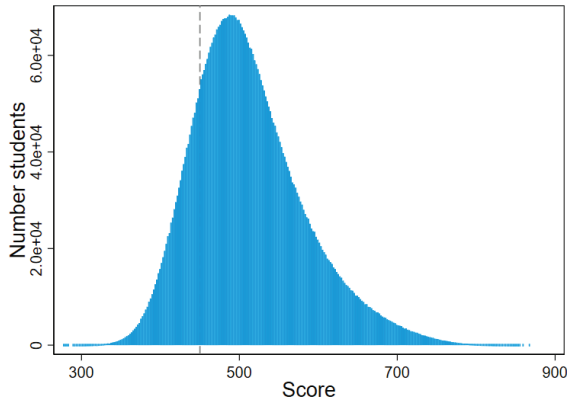
Notes: This figure shows the distribution of income among all incoming students in private sector universities (blue) and among the ones receiving federal loans (red). The vertical dashed line marks the maximum income for participating in the loan program, imposed since 2015. Sources: Census of Higher Education and ENEM Administrative Records.

Figure A.3: Distribution of scores and loan eligibility cutoffs


(a) Degree-specific FIES cutoffs



(b) Distribution of scores

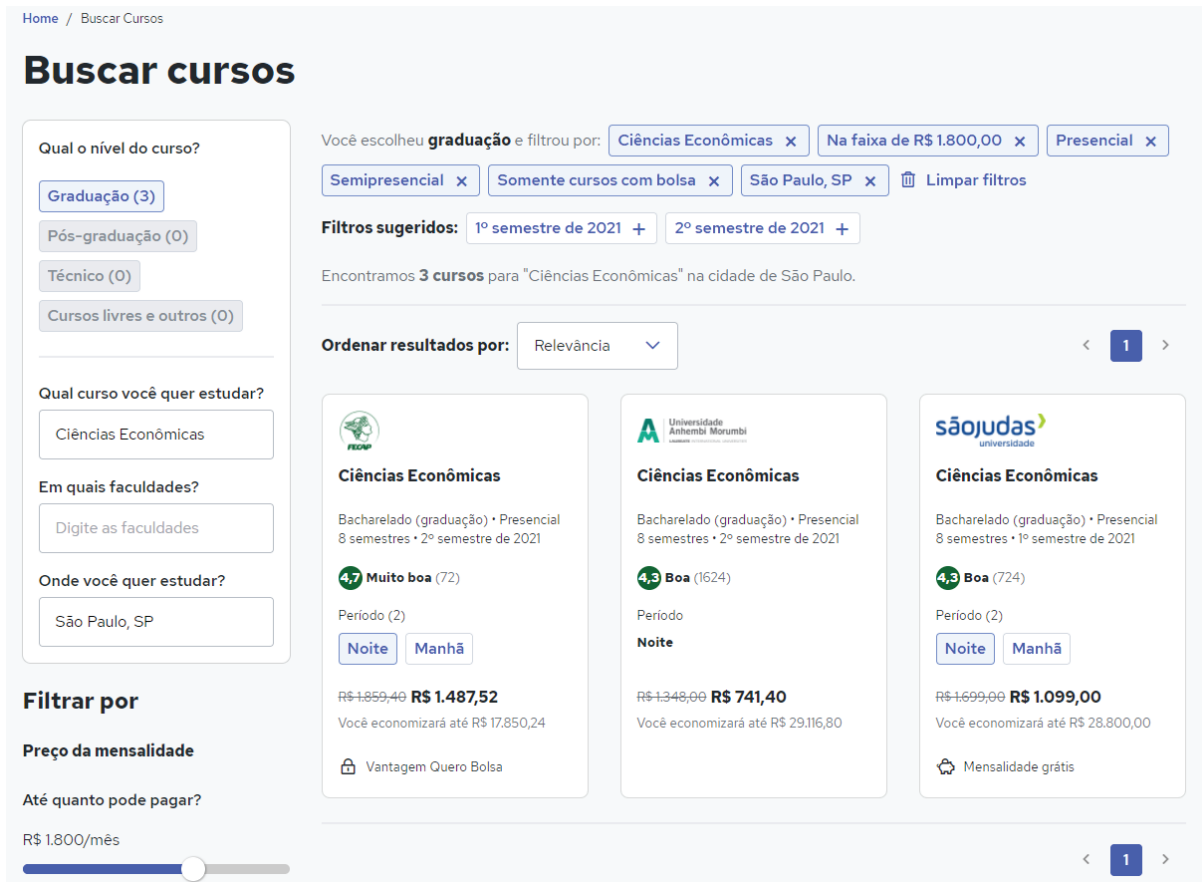


(c) Distribution of scores (around 450 points)



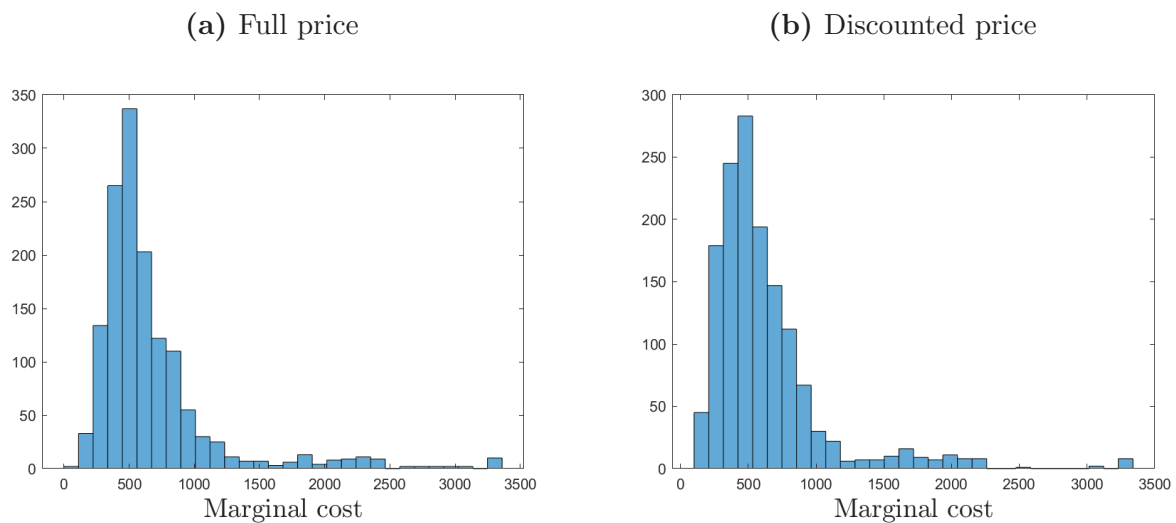
Notes: Panels (b) and (c) show the distribution of ENEM scores. The vertical line marks the minimal score to receive a federal loan (450 points). Panel (b) includes all students who took ENEM in 2015 and Panel (c) the ones with score between 400 and 500. Source: ENEM administrative records. Panel (a) of this figure shows the distribution of degree-specific cutoffs to receive a loan in the FIES program. The vertical line marks the minimal cutoff to participate in the program, imposed by the federal government (450 points). Over 99% of the degrees have cutoffs above 450 points, so this restriction is not binding. Sample: all degrees participating in FIES in 2016. Source: Census of Higher Education, ENEM administrative records, and FIES administrative records.

Figure A.4: Screenshot of an online education marketplace



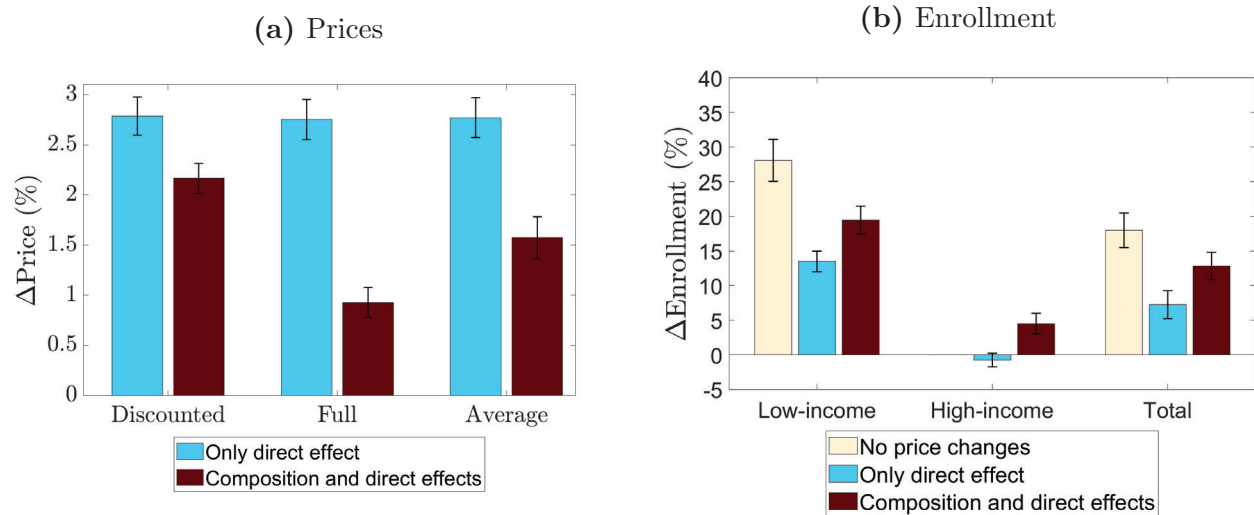
Notes: Screenshot of *QueroBolsa*, the largest online education marketplace in Brazil, responsible for 15% of the tuition discounts in the country.

Figure A.5: Distribution of marginal costs



Notes: This figure shows the distribution of total marginal costs, at the degree level, for students paying the full and discounted prices, as defined in Equation E.1. Marginal costs are estimated inverting colleges' first order conditions, presented in Appendix E. Marginal costs are in BRL.

Figure A.6: Alternative policy: loans only to public high school students



Notes: This figure shows the equilibrium effects of an alternative policy that gives loans only to public high school students. In both panels, the y-axis is the difference between outcomes in an equilibrium with loans only to public high school and an equilibrium without loans. In Panel (a), discounted prices are weighted by the number of students with discounts in each degree, full prices by the number of students paying full price, and the average price by the total number of students. All weights are taken from the counterfactual without loans. Error bars are 95% confidence intervals estimated using bootstrap (details in Appendix H).

B. Conceptual framework – Details

B.1 Proof

In this appendix, we derive Equation (1). To simplify the notation, define $S(p, \{a_x\}_{x \in \mathcal{X}}) \equiv \sum_{x \in \mathcal{X}} N_x s_x(p, a_x)$. The college's problem becomes:

$$\max_p \left\{ S(p, \{a_x\}_{x \in \mathcal{X}}) \cdot (p - c) \right\}.$$

The first-order condition is:

$$(p - c) \cdot \frac{\partial S}{\partial p} + S = 0.$$

Deriving with respect to a_x :

$$\frac{dp}{da_x} \cdot \frac{\partial S}{\partial p} + (p - c) \cdot \frac{d}{da_x} \frac{\partial S}{\partial p} + \frac{dS}{da_x} = 0.$$

Applying the chain rule:

$$\frac{dp}{da_x} \cdot \frac{\partial S}{\partial p} + (p - c) \cdot \left[\frac{\partial^2 S}{\partial p \partial a_x} + \frac{\partial^2 S}{\partial p^2} \frac{dp}{da_x} \right] + \left[\frac{\partial S}{\partial a_x} + \frac{\partial S}{\partial p} \frac{dp}{da_x} \right] = 0.$$

Replacing the first-order condition:

$$\frac{dp}{da_x} \cdot \frac{\partial S}{\partial p} - \frac{S}{\frac{\partial S}{\partial p}} \cdot \left[\frac{\partial^2 S}{\partial p \partial a_x} + \frac{\partial^2 S}{\partial p^2} \frac{dp}{da_x} \right] + \left[\frac{\partial S}{\partial a_x} + \frac{\partial S}{\partial p} \frac{dp}{da_x} \right] = 0. \quad (\text{B.1})$$

Now, note the following identities:

$$\begin{aligned} \frac{\partial S}{\partial a_x} &= \sum_{\tilde{x} \in \mathcal{X}} N_{\tilde{x}} \frac{\partial s_{\tilde{x}}}{\partial a_x} \\ &= N_x \frac{\partial s_x}{\partial a_x}, \end{aligned} \quad (\text{B.2})$$

and:

$$\begin{aligned}
\frac{\partial^2 S}{\partial p \partial a_x} &= \frac{\partial}{\partial p} \frac{\partial S}{\partial a_x} \\
&= \frac{\partial}{\partial p} N_x \frac{\partial s_x}{\partial a_x} \\
&= N_x \frac{\partial}{\partial a_x} \frac{\partial s_x}{\partial p} \\
&= N_x \frac{\partial}{\partial a_x} \left[\frac{s_x}{p} \cdot \frac{p}{s_x} \frac{\partial s_x}{\partial p} \right] \\
&= -\frac{N_x}{p} \frac{\partial}{\partial a_x} [s_x \cdot \eta_x] \\
&= -\frac{N_x}{p} \left[s_x \cdot \frac{\partial \eta_x}{\partial a_x} + \eta_x \cdot \frac{\partial s_x}{\partial a_x} \right] \\
&= -N_x \frac{s_x}{p} \cdot \frac{\partial \eta_x}{\partial a_x} - N_x \frac{\eta_x}{p} \cdot \frac{\partial s_x}{\partial a_x}
\end{aligned} \tag{B.3}$$

Replacing Equations (B.2) and (B.3), into (B.1), and isolating $\frac{1}{p^*} \frac{dp^*}{da_x}$, we get Equation (1):

$$\frac{1}{p^*} \frac{dp^*}{da_x} = \left(\frac{N_x s_x}{\sum_{\tilde{x} \in \mathcal{X}} N_{\tilde{x}} s_{\tilde{x}}} \right) \cdot \frac{1}{\eta^2} \cdot \frac{1}{2 - \lambda} \cdot \left[\left(-\frac{\partial \eta_x}{\partial a_x} \right) + (\eta - \eta_x) \cdot \frac{1}{s_h} \frac{\partial s_x}{\partial a_x} \right]. \quad \square$$

B.2 Price discrimination

We now introduce price discrimination to the model. Consider a market with only one college, a single-product monopolist with marginal cost c . The college price discriminates by charging different prices $\mathbf{p} = \{p_q\}_{q \in \mathcal{Q}}$. Students are divided into groups of consumer types x , and \mathcal{X} is the set of consumer types. Let N_x be the size of group x and $s_x(p, a_x)$ the probability that a student of group x enrolls in college. The probability that a student of group x pays price q is ρ_{xq} . Finally, a_x is a continuous variable denoting the generosity of financial aid for group x .

The problem of the college is given by:

$$\mathbf{p}^* = \arg \max_{\mathbf{p}} \sum_{x \in \mathcal{X}} \sum_{q \in \mathcal{Q}} \rho_{xq} N_x \cdot s_x(p_q, a_x) \cdot (p_q - c).$$

Let η_q be the elasticity of total demand with respect to price q , η_{xq} the elasticity of demand coming from group x with respect to price q , and λ_q a measure of the curvature of the demand

curve.⁹

The effect of a marginal increase of financial aid on the equilibrium price p_q^* is:

$$\frac{1}{p_q^*} \frac{dp_q^*}{da_x} = \underbrace{\left(\frac{\rho_{xq} N_x s_x}{\sum_{\tilde{x} \in \mathcal{X}} \rho_{\tilde{x}q} N_{\tilde{x}} s_{\tilde{x}}} \right)}_{\text{scale}} \cdot \underbrace{\frac{1}{\eta_q^2}}_{\text{market power}} \cdot \underbrace{\frac{1}{2 - \lambda_q}}_{\text{curvature}} \cdot \left[\underbrace{\left(-\frac{\partial \eta_{xq}}{\partial a_x} \right)}_{\text{direct effect}} + \underbrace{(\eta_q - \eta_{xq}) \cdot \frac{1}{s_h} \frac{\partial s_x}{\partial a_x}}_{\text{composition effect}} \right]. \quad (\text{B.4})$$

There are two important differences between Equations (1), without price discrimination, and (B.4), with price discrimination.

First, the scale component in Equation (B.4) depends not only on the size of group x , but on the number of individuals from the group that pay price p_q . Consider an extreme example in which individuals of group x never pay p_q , that is, $\rho_{xq} = 0$. In this case, increasing financial aid for group x has no effect on p_q .

Second, with or without price discrimination, the magnitude and the direction of the composition effect depend on how the price elasticity of the targeted group compares with the overall price elasticity. However, with price discrimination, the *relevant elasticity is with respect to p_q* . Consider, for example, an extreme case with perfect price discrimination, in which only group x pays p_q . Formally, $\rho_{xq} = 1$ and $\rho_{x'q} = 0, \forall x' \neq x$. In this case, $\eta_q = \eta_{xq}$, and there is no composition effect.

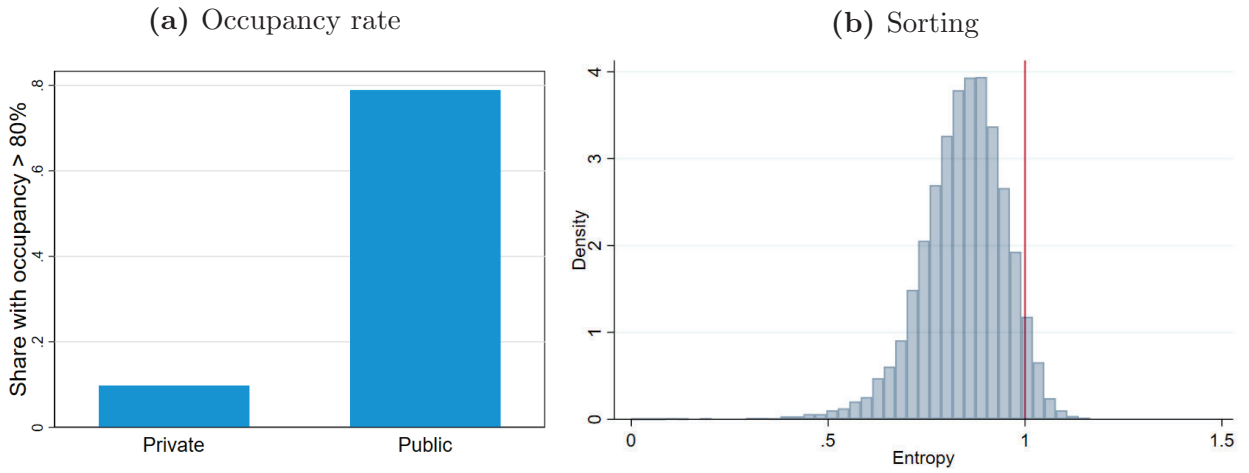
⁹ $\eta_{xq} \equiv -\frac{p_q}{\rho_{xq} N_x s_x} \frac{\partial(\rho_{xq} N_x s_x)}{\partial p_q}$; $\eta_q \equiv -\frac{p_q}{\sum_{x \in \mathcal{X}} \rho_{xq} N_x s_x} \frac{\partial \left(\sum_{x \in \mathcal{X}} \rho_{xq} N_x s_x \right)}{\partial p_q}$; $\lambda_q \equiv \frac{\left(\sum_{x \in \mathcal{X}} \rho_{xq} N_x s_x \right) \frac{\partial^2 \left(\sum_{x \in \mathcal{X}} \rho_{xq} N_x s_x \right)}{\partial p_q^2}}{\left[\frac{\partial \left(\sum_{x \in \mathcal{X}} \rho_{xq} N_x s_x \right)}{\partial p_q} \right]^2}$.

C. Selectivity in private colleges

This appendix discuss the selectivity of private colleges and show empirical patterns suggesting that they are not selective.

Figure C.1a shows that only 10% of the degrees in private colleges fill over 80% of their available spots. In contrast, almost 80% of the degrees in public colleges fill over 80% of their spots. These patterns suggests that admission to public degrees is selective, whereas to private degrees is not. This is expected, since public institutions are free and more prestigious, whereas the private ones charge a tuition and are perceived as lower quality.

Figure C.1: Selectivity in private colleges



Notes: Panel a shows occupancy rates of posted spots in public and private degrees. Panel b shows the distribution of entropy (based on student scores) in private degrees, as defined in Equation C.1. Entropy is normalized by the entropy of the microregion where the degree is located.

We now investigate if there is strong evidence of sorting based scores in private degrees. For this we calculate the *entropy* of each degree. A low entropy means strong sorting. More specifically, we divide students in seven bins based on their ENEM score and calculate the entropy EN_j of each degree j as:

$$EN_j = - \sum_{b=1}^7 p_j(b) \log p_j(b) , \quad (\text{C.1})$$

where $p_j(b)$ is the share of students in degree j who are in the score bin b .

Figure C.1b shows the distribution of entropy of private degrees, normalized by the entropy of the region (commuting zone) where the degree is. An entropy of zero means that all students in the degree are in the same score bin. An entropy of one means that the distribution of student score in the degree is the same as in the degree's region. That is, zero entropy means perfect sorting and one means no sorting. We see there is limited evidence of sorting based on scores, providing further

evidence that private degrees are not selective.

D. Subsidized loans and market structure

In this appendix, we provide evidence that the large reduction in loan availability that occurred in 2015 did not substantially change the structure of the market, at least in the short run. In particular, that is did not lead to the closure of degrees or a reduction in quality.

First, we look at aggregate numbers. Churn is low at the college level: weighed by size, less than 0.5% of colleges exit each year, including in the years following the policy change. At degree level, churn is high, but not higher than usual after the policy change: 14% of degrees that were enrolling new students in 2012 were not anymore in 2014, compared to 10% between 2014 and 2016. Moreover, the degrees that close are the smaller ones: weighted by size, less than 3% of the degrees close every year.

Now we discuss whether degrees that were more affected by the policy were more likely to close or reduce quality. We use the share of students with loans in 2012 as our measure of exposure. More specifically, we estimate the following event study regression:

$$Y_{jt} = \sum_{\tau} \beta_{\tau} \cdot S_{j2012}^L \cdot \mathbb{1}\{t = \tau\} + \gamma_j + \gamma_{ht} + \epsilon_{jt} , \quad (\text{D.1})$$


where j is a degree and t is a year, Y_{jt} is the outcome of interest, S_{j2012}^L is the share of students with loans in 2012, γ_j are degree fixed effects, and ϵ_{jt} are residuals. Mirroring the rule what allocate loans to degrees, we also include quality-region-field-year fixed effects (γ_{ht}).

In the main text, Figure (4) shows the estimates of Equation (D.1) with the share of students with discounts as the outcome. We find that more exposed degrees offered more discounts after the policy change.

Figure D.1, Panel (a) shows the estimates of Equation (D.1) with a dummy for having zero incoming students as an outcome. We see no evidence that more affected degrees were more likely to close. More and less exposed degrees followed the same trend until two years after the policy change. In 2017, the more exposed degrees were *less* likely to close, but the magnitude is small.

Figure D.1, Panel (b) shows the estimates of Equation (D.1) with log faculty-to-student ratio as an outcome, as a proxy for quality. The number of students is kept fixed at its 2012 level. We begin the sample in 2013 because this is the first year for which we have access to degree-level faculty data. Moreover, we extend the sample until 2019 to compare the effects of the 2015 policy with another reform implemented in 2018 that explicitly incentivized colleges to improve quality. More specifically, the 2018 reform imposes a penalty to colleges when students using federal loans dropout. In a follow-up paper (Dobbin et al., 2021), we discuss this policy in detail. We see no evidence that more exposed degrees reduced quality after the 2015 reform. Moreover, we do see strong evidence that more exposed degrees increased quality after the 2018 reform, showing that changes in the faculty-to-student ratio is a good proxy for short term investments to improve quality.

Figure D.1: Student loans and market structure



Notes: This figure shows the OLS estimates of β_τ in equation (D.1). The vertical line marks the policy change that reduced the availability of subsidized student loans (2015) and another reform that incentivized colleges to increase quality (2018). Each observation is a degree-year, and a degree is a major-institution pair. Each panel shows results for a different outcome, specified in its title. The brackets represent 5% confidence intervals. Source: Census of Higher Education and FIES administrative records.

In summary, the results in this appendix show that there is no evidence that degrees exited the market or reduced quality in response to the 2015 policy change.

E. First order conditions of colleges' problem

This appendix presents the first order condition of colleges' problems with multi-product firms. These are the conditions used in our structural estimation and counterfactual exercises. The conditions presented in the main text—in Equations (7), (8), and (9)—are simplified versions, used to discuss intuition.

Define:

$$\begin{aligned} c_{kt}^F &= c_{kt} + \psi \cdot (\bar{D}_f - \omega_f)^2 \\ c_{kt}^D &= c_{kt} + \psi \cdot (\bar{D}_f - \omega_f)^2 - \kappa_{jt} \end{aligned} \tag{E.1}$$

The general first order conditions are:

$$\begin{aligned} p_{kt}^F &= c_{kt}^F + \frac{E_i [s_{ik} | B_i^D < \bar{D}_f, i \in \mathcal{I}_t]}{E_i \left[-\frac{\partial s_{ik}}{\partial p_k^F} | B_i^D < \bar{D}_f, i \in \mathcal{I}_t \right]} + \sum_{j \in \mathcal{J}_f \setminus k} \frac{E_i \left[\frac{\partial s_{ij}}{\partial p_k^F} \cdot (p_{jt}^F - c_{jt}^F) | B_i^D < \bar{D}_f, i \in \mathcal{I}_t \right]}{E_i \left[-\frac{\partial s_{ik}}{\partial p_k^F} | B_i^D < \bar{D}_f, i \in \mathcal{I}_t \right]}, \\ p_{kt}^D &= c_{kt}^D + \frac{E_i [s_{ik} | B_i^D \geq \bar{D}_f, i \in \mathcal{I}_t]}{E_i \left[-\frac{\partial s_{ik}}{\partial p_k^D} | B_i^D \geq \bar{D}_f, i \in \mathcal{I}_t \right]} + \sum_{j \in \mathcal{J}_f \setminus k} \frac{E_i \left[\frac{\partial s_{ij}}{\partial p_k^D} \cdot (p_{jt}^D - c_{jt}^D) | B_i^D \geq \bar{D}_f, i \in \mathcal{I}_t \right]}{E_i \left[-\frac{\partial s_{ik}}{\partial p_k^D} | B_i^D \geq \bar{D}_f, i \in \mathcal{I}_t \right]}, \\ \bar{D}_{ft} &= \frac{1}{2\psi} \cdot E[\pi_{if}^F - \pi_{if}^D | B_i^D = \bar{D}_{ft}, i \in \mathcal{I}_t] + \omega_{ft}. \end{aligned}$$

F. The effects of loans on demand

In this appendix, we derive Equation (13). This equation is a close-form estimator of the causal effect of loans on price sensitivity in a simplified model without discounts and in which all students are loan takers. Under these simplifying assumptions, students' utility function becomes:

$$U_{ij} = r_i \beta^r \cdot R_j + \alpha_1 p_j \mathbb{1}\{r_i \geq \bar{r}_j\} + \alpha_0 p_j \mathbb{1}\{r_i < \bar{r}_j\} + \gamma_j + \epsilon_{ij} .$$

The estimation is done separately for high- and low-income students, and only in the post-policy-change period (there are no score cutoffs in the pre period), so we omit the w and t subscripts, for clarity.

The market share s_{jr} of degree j among students with score r is given by:

$$\begin{aligned} s_{jr} &= \int_{r_i=r} \frac{e^{U_{ij}}}{\sum_{k \in \mathcal{J}} e^{U_{ik}}} di \\ &= \int_{r_i=r} \frac{e^{r_i \beta^r \cdot R_j + \alpha_1 p_j \mathbb{1}\{r_i \geq \bar{r}_j\} + \alpha_0 p_j \mathbb{1}\{r_i < \bar{r}_j\} + \gamma_j}}{\sum_{k \in \mathcal{J}} e^{U_{ik}}} di \\ &= e^{r \beta^r \cdot R_j + \alpha_1 p_j \mathbb{1}\{r \geq \bar{r}_j\} + \alpha_0 p_j \mathbb{1}\{r < \bar{r}_j\} + \gamma_j} \int_{r_i=r} \frac{1}{\sum_{k \in \mathcal{J}} e^{U_{ik}}} di . \end{aligned}$$

Therefore:

$$\log s_{jr} = r \beta^r \cdot R_j + \alpha_1 p_j \mathbb{1}\{r \geq \bar{r}_j\} + \alpha_0 p_j \mathbb{1}\{r < \bar{r}_j\} + \gamma_j + \log \int_{r_i=r} \frac{1}{\sum_{k \in \mathcal{J}} e^{U_{ik}}} di .$$

Analogously, for the outside option:

$$\log s_{0r} = r \beta^r \cdot R_0 + \log \int_{r_i=r} \frac{1}{\sum_{k \in \mathcal{J}} e^{U_{ik}}} di .$$

Therefore:

$$\log \frac{s_{jr}}{s_{0r}} = r \beta^r \cdot (R_j - R_0) + \alpha_1 p_j \mathbb{1}\{r \geq \bar{r}_j\} + \alpha_0 p_j \mathbb{1}\{r < \bar{r}_j\} + \gamma_j .$$

Therefore:

$$\begin{aligned}\lim_{r \rightarrow \bar{r}_j^-} \log \frac{s_{jr}}{s_{0r}} &= \bar{r}_j \beta^r \cdot (R_j - R_0) + \alpha_0 p_j + \gamma_j , \\ \lim_{r \rightarrow \bar{r}_j^+} \log \frac{s_{jr}}{s_{0r}} &= \bar{r}_j \beta^r \cdot (R_j - R_0) + \alpha_1 p_j + \gamma_j .\end{aligned}$$

Therefore:

$$\lim_{r \rightarrow \bar{r}_j^+} \log \frac{s_{jr}}{s_{0r}} - \lim_{r \rightarrow \bar{r}_j^-} \log \frac{s_{jr}}{s_{0r}} = p_j \cdot (\alpha_1 - \alpha_0) .$$

Finally, rearranging and summing over all degrees:

$$\alpha_1 - \alpha_0 = \frac{1}{J} \sum_{j \in \mathcal{J}} \frac{\lim_{r \rightarrow \bar{r}_j^+} \log \frac{s_{jr}}{s_{0r}} - \lim_{r \rightarrow \bar{r}_j^-} \log \frac{s_{jr}}{s_{0r}}}{p_j} . \quad \square$$

G. Decomposing the direct and composition effects

G.1 General case

From the first order conditions shown in Appendix E, we have that equilibrium prices are equal to:

$$p_k^F = c_k^F + \frac{1}{P(w_i = Low|d_{ik} = 1) \cdot \alpha_{k,F}^{Low\ Income} + P(w_i = High|d_{ik} = 1) \cdot \alpha_{k,F}^{High\ Income}} + \sum_{k \in \mathcal{J}_f \setminus k} \frac{E_i [s_{ik}s_{ij} \cdot (p_{ik} - c_k(p_{ik}, \mathbf{p}_k)) | B_i^D < \bar{D}_f]}{P(w_i = Low|d_{ik} = 1) \cdot \alpha_{k,F}^{Low\ Income} + P(w_i = High|d_{ik} = 1) \cdot \alpha_{k,F}^{High\ Income}},$$

$$p_k^D = c_k^D + \frac{1}{P(w_i = Low|d_{ik} = 1) \cdot \alpha_{k,D}^{Low\ Income} + P(w_i = High|d_{ik} = 1) \cdot \alpha_{k,D}^{High\ Income}} + \sum_{k \in \mathcal{J}_f \setminus k} \frac{E_i [s_{ik}s_{ij} \cdot (p_{ik} - c_k(p_{ik}, \mathbf{p}_k)) | B_i^D \geq \bar{D}_f]}{P(w_i = Low|d_{ik} = 1) \cdot \alpha_{k,D}^{Low\ Income} + P(w_i = High|d_{ik} = 1) \cdot \alpha_{k,D}^{High\ Income}},$$

where

$$\begin{aligned} \alpha_{k,F}^{w_i} &= P(L_{ik} = 0 | w_i \ \& \ d_{ik} = 1 \ \& \ B_i^D < \bar{D}_f) \cdot (1 - s_{k,F,0}^{w_i}) \cdot \alpha_{w_i,0} \\ &\quad + P(L_{ik} = 1 | w_i \ \& \ d_{ik} = 1 \ \& \ B_i^D < \bar{D}_f) \cdot (1 - s_{k,F,1}^{w_i}) \cdot \alpha_{w_i,1} \\ \alpha_{k,D}^{w_i} &= P(L_{ik} = 0 | w_i \ \& \ d_{ik} = 1 \ \& \ B_i^D \geq \bar{D}_f) \cdot (1 - s_{k,D,0}^{w_i}) \cdot \alpha_{w_i,0} \\ &\quad + P(L_{ik} = 1 | w_i \ \& \ d_{ik} = 1 \ \& \ B_i^D \geq \bar{D}_f) \cdot (1 - s_{k,D,1}^{w_i}) \cdot \alpha_{w_i,1} \\ s_{k,F,L}^{w_i} &= P(d_{ik} = 1 | w_i \ \& \ L_{ik} = L \ \& \ B_i^D < \bar{D}_f) \\ s_{k,D,L}^{w_i} &= P(d_{ik} = 1 | w_i \ \& \ L_{ik} = L \ \& \ B_i^D \geq \bar{D}_f) \end{aligned}$$

To allow for the direct effect and shut down the composition effect, we compute prices using the above equations and input the share of low-income students— $P(w_i = Low|d_{ij} = 1)$ —from the equilibrium without loans—defined in Equation (19)—and the share students with loans— $P(L_{ij} = 0 | w_i = Low \ \& \ d_{ij} = 1)$ and $P(L_{ij} = 0 | w_i = High \ \& \ d_{ij} = 1)$ —from the data.

Moreover, in the baseline results, we compute prices taking s_{ij} , $s_{k,F,L}^{w_i}$, and $s_{k,D,L}^{w_i}$ from the data. However, since there are many degrees, each degree has a very small market shares ($< 0.1\%$). Hence, the change in the results taking these market shares from the data or from the counterfactual without loans is negligible.

G.2 Small single-price single-degree college

The price of a single-price single-degree college is:

$$p_k^F = c_k^F + \frac{1}{P(w_i = Low | d_{ik} = 1) \cdot \alpha_k^{Low\ Income} + P(w_i = High | d_{ik} = 1) \cdot \alpha_k^{High\ Income}} .$$

Moreover, for a small college $s_{k,F,L}^{w_i} \approx 0$, then:

$$\alpha_j^{Low\ Income} \approx P(L_{ij} = 0 | w_i = Low \ \& \ d_{ij} = 1) \cdot \alpha_{Low,0} + P(L_{ij} = 1 | w_i = Low \ \& \ d_{ij} = 1) \cdot \alpha_{Low,1} ,$$

$$\alpha_j^{High\ Income} \approx P(L_{ij} = 0 | w_i = High \ \& \ d_{ij} = 1) \cdot \alpha_{High,0} + P(L_{ij} = 1 | w_i = High \ \& \ d_{ij} = 1) \cdot \alpha_{High,1} ,$$

which gives us Equation (20).

H. Estimating standard errors for the structural parameters

In this appendix, we describe how we calculate confidence intervals for our structural estimates. First, we follow the following steps to build a bootstrap sample:

1. We divide students into 128 for bins, based on their period (pre or post) and the following covariates: the student’s family income, her score, whether she attended a public high school, whether at least one of the student’s parents attended college, and her answer to a survey question asking whether she plans to apply to a government student loan. We collapse family income into four bins divided at two, three, and seven times the federal minimum wage. We collapse score into equally sized bins.
2. For each bin, we build a bootstrap bin by drawing, with replacement, a number of students equal to the size of the bin.
3. We build the bootstrap sample combining all bootstrap bins.

We then estimate our structural model, following the procedure described in Section 6. We then compute all counterfactuals described in Section 7 with the new estimated parameters.

We repeat this procedure 100 times, leaving us with a sample of 100 different parameter estimates and counterfactual simulations. Finally, we build confidence intervals by taking the percentiles 2.5 and 97.5 of each parameter and counterfactual outcome across these 100 simulations.

I. Alternative assumptions in welfare calculations

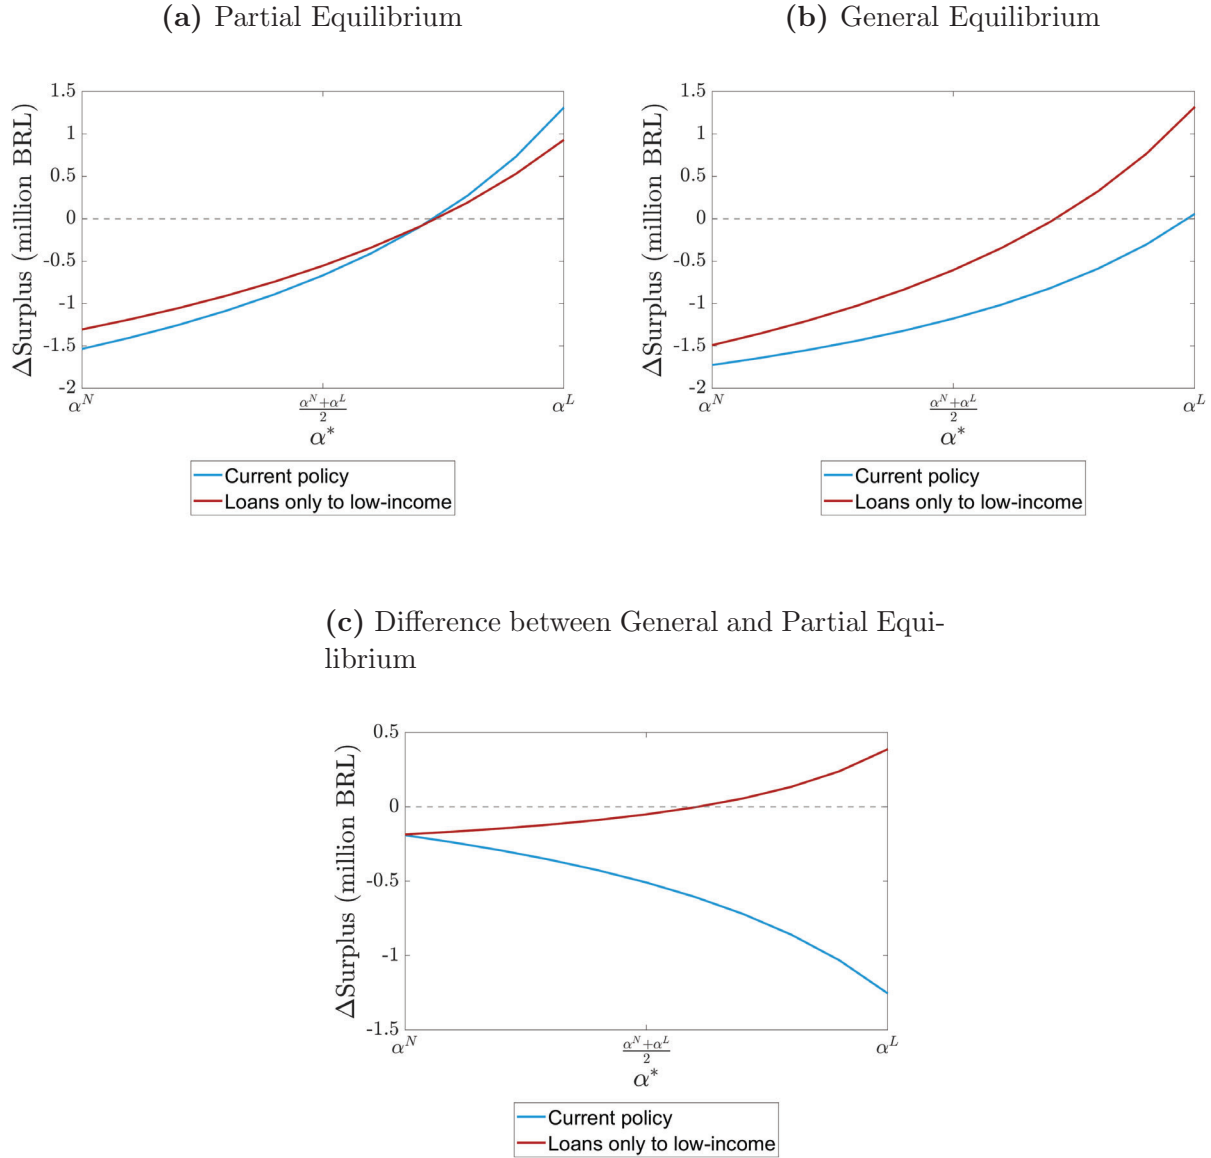
As discussed in Section 7.4, calculating welfare requires additional assumptions. Our baseline results assume that the loan program corrects existing market failures and, hence, the choices of students with loans reveal the parameters of their underlying *experience utility*. An alternative assumption is that there are no market failures and, hence, the choices of students without loans reveal the parameters of their underlying *experience utility*.

This appendix shows how our results change for a series of alternative assumptions between these two extremes. The results are in Figure I.1. Our findings are twofold.

First, Panels (a) and (b) show that whether the overall effects of loans are positive or negative depends on our assumption regarding students' *experience utility*. This result is expected: under the assumption that loans correct existing market failures, loans increase welfare; under the assumption that loans distort an otherwise perfect market, loans decrease welfare. Determining which assumption is more realistic requires estimating the returns of different degrees and is outside of the scope of this paper.

Second, and more related to our object of interest, Panel (c) shows how the welfare consequences of supply-side responses change under different assumptions. We find that supply-side responses increase welfare in our proposed alternative policy relative to the current one under any considered assumption.

Figure I.1: Calculating welfare under alternative assumptions



Notes: This figure shows the results of our welfare estimates, as defined in Equation (23), under different assumptions. The x-axis is the assumed value of α^* , ranging between the sensitivity of students with (α^L) and without (α^N) loans. In Panels (a) and (b), the y-axis is the difference in total consumer surplus in an equilibrium with loans compared to one without. In partial equilibrium—Panel (a)—, pricing policies are kept the same as in the equilibrium without loans and, in general equilibrium—Panel (b)—, they follow colleges’ first order conditions. In Panel (c), the y-axis is the difference in total consumer surplus between partial and general equilibrium.