Equilibrium Price Responses to Targeted Student Financial Aid*

Nano Barahona† Cauê Dobbin§ Sebastián Otero‡

April 6, 2026

Abstract

We study how colleges adjust tuition in response to student financial aid, emphasizing the role of targeting. Our framework highlights two forces: a *direct* effect that increases tuition, and a *composition* effect that can reduce tuition when aid targets more price-sensitive students. Using a reform in Brazil's student loan program, we document both effects empirically and estimate an equilibrium model to evaluate counterfactual policies. We find that merit-based targeting *raises* tuition by 2.2%, while need-based targeting *lowers* it by 0.8%, reflecting that lower-income students are more price-sensitive. These price effects have a strong impact on enrollment decisions.

Keywords: Student financial aid, Pass-through, Targeting

JEL Codes: L13, I22, I23, H22

*We are grateful to Caroline Hoxby, Rebecca Diamond, Melanie Morten, Isaac Sorkin, and Constantine Yannelis for their invaluable guidance and support throughout this project. We also thank Claudia Allende, Luis Armona, Michael Best, Sandra Black, Arun Chandrasekhar, Pascaline Dupas, Liran Einav, Marcel Fafchamps, Matthew Gentzkow, Neale Mahoney, Nathan Miller, and Lesley Turner for their insightful comments and suggestions. We benefited greatly from discussions at the NBER Summer Institute, the Barcelona School of Economics Summer Forum, and the Econometric Society Summer and Winter Meetings, as well as from seminar participants at Stanford University, the University of Chicago, Yale University, Georgetown University, the World Bank, Center for Monetary and Financial Studies (CEMFI), Institute for Fiscal Studies, Toulouse School of Economics, Einaudi Institute for Economics and Finance, Pompeu Fabra University, Fundação Getulio Vargas Rio de Janeiro (FGV-RJ) and São Paulo (FGV-SP), Pontifical Catholic University of Rio de Janeiro (PUC-RJ), and Federal University of Paraíba (UFPB). 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, the NAEd/Spencer Dissertation Fellowship Program, and the Georgetown Center for Economic Research (GCER). Any remaining errors are our own. †University of California, Berkeley and NBER. Email: `nanobk@berkeley.edu`. §Georgetown University. Email: `caue.dobbin@georgetown.edu`. ‡Columbia University and NBER. Email: `so2699@columbia.edu`.

1 Introduction

Governments around the world invest substantial resources in student financial aid programs with the aim of expanding access to higher education. Such initiatives are widespread across both Latin America and OECD countries.[1] Despite their prevalence, policymakers have long expressed concern that these programs may incentivize colleges to raise tuition and capture a portion of public funds—an issue known as the Bennett Hypothesis (Bennett, 1987).

Conceptually, student aid programs can generate two opposing forces. On one hand, aid increases students' ability to pay, which may induce colleges to raise tuition. On the other hand, if aid is targeted to lower-income students—who typically have lower willingness to pay—colleges may reduce tuition to attract them. As a result, the net impact of student aid on tuition is theoretically ambiguous and depends on how the aid is targeted—a mechanism overlooked in the literature.

This paper examines how tuition responses to financial aid depend on targeting. We begin by presenting a conceptual framework that formalizes the channels through which aid can affect tuition under imperfect competition. We then provide empirical evidence that the forces identified in the framework are quantitatively relevant, leveraging a major reform of Brazil's federal student loan program. Building on this analysis, we develop and estimate an equilibrium model of supply and demand in higher education. The model allows us to compare the effects of merit-based and need-based financial aid—two commonly used targeting schemes—on tuition and enrollment.

In our framework, a monopolistic college sets tuition to maximize profits, while the government provides targeted financial aid. We show that the effect of aid on tuition has two components: a *direct* effect and a *composition* effect. The direct effect comes from aid raising recipients' willingness to pay (WTP), which shifts the demand curve outward and puts upward pressure on tuition. The composition effect, on the other hand, depends on how the aid is targeted. When aid goes to students with lower baseline WTP, it narrows the gap between high- and low-WTP students, causing the demand curve to rotate and become flatter (more price-elastic). This rotation pushes prices down and can lead to a lower equilibrium tuition. In contrast, when aid is directed toward higher-WTP students, the demand curve becomes steeper (less price-elastic), which amplifies the rise in tuition.

Given a targeting scheme, the magnitudes of the direct and composition effects depend on three key parameters. First, the extent to which aid increases WTP determines the strength of the direct effect, since a greater increase in WTP leads to a larger upward shift in the demand curve. Second, the degree of heterogeneity in WTP across students: the greater the disparity in WTP between targeted and non-targeted students, the stronger the composition effect. Third, the degree of price discrimination: The more effectively colleges engage in price discrimination, the more they can offer lower prices to low-WTP students while maintaining higher prices for others, thereby weakening the composition effect.

We explore the empirical relevance of these forces in Brazil's private higher education market, which

[1] For detailed overviews, see Marta Ferreyra et al. (2017) for Latin America and OECD (2014) for OECD countries.

accounts for 75% of all incoming college students. This context offers several key advantages for studying the equilibrium effects of financial aid. First, for-profit institutions play a prominent role, making tuition and enrollment outcomes highly responsive to market forces. Second, the federal government allocates subsidized student loans through a centralized system with clear eligibility thresholds, allowing us to estimate how loan access affects student demand. Third, a major policy reform in 2015 sharply reduced the availability of loans, providing a natural experiment to assess the broader market impact of financial aid. Finally, unlike in the U.S., Brazilian colleges have limited access to students' financial information, which constrains their ability to price discriminate.[2]

We combine multiple data sources for our analysis. First, we merge several administrative records to build a comprehensive individual-level dataset covering the full population of students who took Brazil's college entrance exam. This dataset includes students' test scores, detailed demographic information—such as gender, race, and household income—and their college enrollment status. For enrolled students, we observe their chosen college and degree program, as well as whether they received a government loan and the amount borrowed. Second, we construct a novel dataset on college tuition fees by merging government records with data obtained through partnerships with private companies. This allows us to recover tuition information for degree programs accounting for 98% of total enrollment. Our analysis covers the period from 2012 to 2017 and includes over 20 million exam takers and 13,567 unique degree programs offered across 695 colleges.

The empirical analysis begins by examining trends in the Brazilian higher education market before and after the 2015 reform of the federal student loan program. The reform triggered a sharp reduction in new loan contracts, with their number falling nearly fourfold between 2014 and 2017. Real tuition fell by 5.2% following the reform, while the stock prices of the four largest higher education firms declined by 20% to 40%, reflecting both the unexpected nature of the policy and its impact on future profitability expectations.

We leverage the reform to examine how financial aid availability impacts tuition. We find that degree programs with a higher pre-reform share of loan recipients experienced larger reductions in tuition and enrollment compared with less exposed programs. We then explore heterogeneity in tuition responses based on our theoretical framework. In markets in which loan recipients were relatively higher-income, reduced loan availability led to lower tuition, consistent with the Bennett Hypothesis. However, in markets in which loan recipients were mostly lower-income, reduced loan availability caused tuition to increase, which supports our prediction that price effects depend on aid allocation.

Motivated by these findings, we develop an equilibrium model of higher education to evaluate how different targeting designs for government loans affect tuition fees and college enrollment. On the supply side, colleges are modeled as multi-product firms offering a portfolio of degree programs. They set tuition prices for each program to maximize profits while accounting for the allocation of government

[2] In the United States, colleges use financial aid application data to price discriminate (Fillmore, 2023), allowing them to better target marginal students and potentially reducing the importance of composition effects.

loans. Also, colleges engage in price discrimination by offering tuition discounts. On the demand side, students decide whether to enroll in college and select a degree program while considering tuition prices and loan availability. Colleges are assumed to be non-selective and to allow students to enroll in any degree of their choice.[3]

To estimate the model, we employ the generalized method of moments (GMM) approach developed by Berry et al. (1995) and extended by Petrin (2002), which integrates instrumental variables with micro-moments to identify key parameters. We estimate price elasticities using exposure to the 2015 loan reform, proxied by the pre-reform share of students with loans, as an instrument. This approach assumes that observed price changes following the reform primarily reflect endogenous responses to reduced loan availability, rather than unobserved shocks correlated with policy exposure. To identify the impact of loans on student demand, we exploit the discontinuity in loan access at eligibility score thresholds. This strategy relies on the assumption that scores are not systematically manipulated around these cutoffs, a plausible assumption given that eligibility thresholds were ex ante unknown.

We estimate a median price elasticity of -3.46, consistent with prior findings (Armona and Cao, 2024; Barahona et al., 2025). This estimate confirms that colleges possess a moderate degree of market power and implies that changes in financial aid can trigger meaningful tuition adjustments. Importantly, price sensitivity varies systematically with both household income and loan status, which plays a central role in shaping the magnitude of the direct and composition effects. The median price elasticity is -5.07 for students with below-average income and -2.96 for those with above-average income. Moreover, receiving a loan reduces price sensitivity by 11.1% on average. We also find that tuition discounts are only weakly correlated with student income, which suggests that price discrimination plays a limited role in our setting.

We use the model to estimate the equilibrium effects of two common loan targeting schemes: need-based and merit-based. In both cases, loans are allocated under a fixed budget equal to that of the pre-reform policy. We evaluate the effects of each scheme by comparing them to a counterfactual with no loans.

Under the need-based scheme, loans are directed to low-income students. Since these students are more price sensitive, the composition effect exerts downward pressure on tuition, offsetting the upward direct effect. The net effect is a 0.8% decrease in prices, relative to the no-loan benchmark. In contrast, the merit-based scheme allocates loans to students with high test scores, which are strongly correlated with higher income and lower price sensitivity. As a consequence, the composition effect reinforces the direct effect, leading to a tuition increase of 2.2%.

Next, we use the model to assess whether the price effects of loan programs are large enough to meaningfully affect enrollment outcomes. To disentangle supply- and demand-side responses, we

[3]The private sector is largely non-selective. According to the 2014 Census of Higher Education, private college degree programs operate at just 48% of their reported capacity, on average. Moreover, about 90% of these programs enroll fewer than 80% of their available seats.

simulate two counterfactuals for each targeting scheme. The first, labeled *demand-only*, isolates demand-side effects by allowing students to respond to loan availability while holding tuition prices fixed. The second, labeled *equilibrium*, incorporates supply-side responses by allowing tuition prices to adjust in response to changes in demand. Comparing enrollment outcomes across these two counterfactuals reveals the role of price adjustments in shaping enrollment decisions.

Under need-based targeting, the associated price decline raises the enrollment gain from 35.8% in the demand-only counterfactual to 39.6% in equilibrium. By contrast, under merit-based targeting, price increases offset much of the enrollment gain observed in the demand-only counterfactual, reducing the effect from 24.2% to 15.8%. Similar patterns emerge when focusing on enrollment in high-quality degrees. These results underscore how endogenous price responses shape the aggregate impact of loan programs. Moreover, whether price adjustments amplify or dampen enrollment gains depends crucially on how loans are targeted, highlighting the central role of the composition effect.

It is important to emphasize that our analysis focuses on the price effects of targeting and abstracts from other dimensions of financial aid design. For example, high-score students may have lower dropout rates, generate positive peer effects, and exhibit higher loan repayment rates, thereby enhancing the financial sustainability of aid programs. The central policy implication of our results is that tuition responses are highly sensitive to how aid is targeted and must be weighed against these broader considerations when designing financial aid policies.

Our work contributes to the literature on tax incidence under imperfect competition (Delipalla and Keen, 1992; Anderson et al., 2001; Weyl and Fabinger, 2013; Miravete et al., 2018; Kroft et al., 2024a,b). A central result in this literature is that pass-through depends on demand elasticity and curvature, but these insights have been developed primarily for uniform policies—common in product taxation—where all consumers face the same tax.[4] We show that when subsidies are targeted—as is common in student aid and other transfer programs—an additional force emerges: targeting reallocates demand across consumers with heterogeneous price sensitivities, endogenously changing the elasticity of residual demand faced by firms. Our decomposition of subsidy incidence into direct and composition effects formalizes this channel and shows that it can either amplify or offset standard pass-through, depending on whether aid is directed toward more or less price-sensitive consumers.

The composition effect we identify is closely related to the "demographic externality" concept introduced by Polyakova and Ryan (2022) in contemporaneous work on health insurance markets.[5] Both papers share the core idea that targeted subsidies shift the composition of demand, thereby altering

[4]A related literature studies optimal taxation with heterogeneous agents, including targeted and nonlinear tax schedules (Ramsey, 1927; Mirrlees, 1971; Diamond, 1975; Saez, 2001; Piketty and Saez, 2013). These frameworks typically assume competitive markets or a social planner, so the mechanism we emphasize—whereby targeting alters the elasticity of residual demand faced by firms—does not emerge.

[5]Earlier work in health and insurance markets studies how prices depend on consumer composition, but emphasizes cost heterogeneity—arising from adverse selection—rather than demand-side externalities (Rothschild and Stiglitz, 1976; Mahoney and Weyl, 2017; Tebaldi, 2025). In contrast, our paper highlights heterogeneity in price sensitivity as the key channel through which consumer composition affects prices.

the residual demand elasticity faced by firms and, in turn, equilibrium prices. The two papers differ along several dimensions. First, our conceptual framework isolates what determines the direction of the composition effect: whether aid is targeted toward more or less price-sensitive consumers. Second, building on this insight, we use our empirical model to compare alternative targeting designs and show that the composition effect can either increase or decrease prices depending on whether aid is merit-based (targeting less price-sensitive students) or need-based (targeting more price-sensitive students). In contrast, Polyakova and Ryan (2022) focus on a single targeting scheme.[6] Third, education markets exhibit tuition discounts, which can attenuate the composition effect, whereas price discrimination is more constrained under the regulated pricing typical of insurance markets.

We also contribute to the empirical literature on the relationship between student financial aid and tuition (Long, 2004; Singell and Stone, 2007; Cellini and Goldin, 2014; Lucca et al., 2019; Baird et al., 2022; Black et al., 2023; De Mello and Duarte, 2020). This literature has generally found that expansions in student aid lead to higher tuition, consistent with the Bennett Hypothesis. We offer a reinterpretation of this evidence: observed price responses to student aid need not reflect pure pass-through; they can also arise as equilibrium responses to changes in the composition of enrolled students. Because most aid programs studied in this literature are targeted, the estimated effects reflect both the direct and composition channels we identify. Our framework thus suggests that the Bennett Hypothesis is not simply true or false—whether aid raises or lowers tuition depends on how it is targeted across students with different price sensitivities.

A related literature examines the effects of student loans on higher-education demand. Several papers exploit discontinuities in loan availability to estimate impacts on college enrollment and student choices (Solís, 2017; Londono-Velez and Rodriguez, 2020; Aguirre, 2021; Bucarey et al., 2020; Hampole, 2024), while others use discrete-choice frameworks to study how financial aid shapes student decisions (Avery and Hoxby, 2004; Joensen and Mattana, 2021). We combine elements of both approaches: we exploit discontinuities in loan eligibility to form micro-moments that discipline a structural demand model. This allows us to recover how loan access and price sensitivity vary across students—heterogeneity that is central to quantifying the composition effect and evaluating how different targeting designs affect equilibrium outcomes.

Finally, our empirical model builds on the literature on equilibrium models of education markets. Foundational work develops frameworks featuring student-level price discrimination (Epple and Romano, 1998; Epple et al., 2006; Fillmore, 2023), as well as endogenous admission policies with granular screening (Fu, 2014). While these features are central in the context of selective U.S. schools, they are less relevant in our setting, where colleges are non-selective and price discrimination relies on limited information about students. We instead follow a strand of the literature that models colleges as differentiated-product firms and estimates demand flexibly, often leveraging observed policy variation

[6]Polyakova and Ryan (2022) compare a policy targeted to low-income individuals to an untargeted policy, but do not consider alternative targeting schemes.

for identification (Bucarey, 2018; Neilson, 2021; Allende, 2021; Bodere, 2023; Sanchez, 2023; Armona and Cao, 2024; Cook, 2025). Several of these papers study the equilibrium effects of financial aid, but they typically model aid as a uniform demand shifter or evaluate a single policy design, so the composition channel—which depends on *who* receives aid—does not emerge. We show that targeting is a first-order determinant of equilibrium price responses: tuition under merit-based targeting is up to 3% higher than under need-based targeting, driven by differences in the composition effect.

The remainder of the paper is structured as follows. Section 2 presents a simplified version of our model and the main theoretical result, which demonstrates that the incidence of a targeted subsidy can be decomposed into direct and composition effects. Section 3 introduces the empirical setting and describes the data, and Section 4 empirically tests our theoretical predictions. Section 5 presents the full model and Section 6 explains the estimation procedure. In Section 7, we perform counterfactual simulations to assess the equilibrium effects of aid programs under different targeting schemes. Finally, Section 8 concludes.

2 Conceptual framework: Price effects of targeted student aid

In this section, we analyze how targeting shapes the price effects of financial aid from a theoretical perspective. We begin with a stylized framework that illustrates why targeting matters for pricing in the context of a targeted subsidy. In particular, we show that prices may decrease when the subsidy is directed toward low-WTP consumers. We then extend this framework to accommodate broader forms of financial aid beyond direct subsidies, such as student loans, and use it to formally define the *direct* and *composition* effects.

2.1 The price effects of a targeted subsidy in a linear demand system

Consider a market with a monopolist selling a single product at price p and zero marginal cost. A continuum of consumers, indexed by $i \in \mathcal{I}$, participate in the market. The government provides a specific subsidy in the form of a fixed transfer $\tau_i \in \mathbb{R}$, conditional on purchasing the good. The subsidy amount may vary across consumers, is known ex ante by both consumers and firms, and is independent of the price. The subsidy schedule is denoted by $\mathcal{T} \equiv \{\tau_i\}_{i \in \mathcal{I}}$. Consumers make a discrete choice between purchasing the good or not, with their WTP given by $\theta_i + \tau_i$, where θ_i is uniformly distributed over the interval $[0, 2]$.[7]

The monopolist maximizes profits by solving

$$p^* (\mathcal{T}) = \arg\max_p \, Q\big(p \big| \mathcal{T}\big) \cdot p, \tag{1}$$

[7] A microfoundation for this WTP arises from a consumer with quasi-linear utility $U_i[y, q] = v_i[q] + w_i[y - (p - \tau) \cdot q]$, where $v[\cdot]$ and $w[\cdot]$ are strictly increasing functions, y represents income, and q is an indicator for purchasing the good.

where Q denotes the quantity purchased. We assume that the subsidy schedule $\mathcal{T}$ is such that $Q(p|\mathcal{T})$ is twice-differentiable and decreasing in p, and that the profit function $Q(p|\mathcal{T}) \cdot p$ is strictly concave, which ensures a unique solution to Equation (1).

The first-order condition from Equation (1) implies that the firm sets p^* such that the absolute price elasticity of demand equals one:

$$\eta \left[p^* \left(\mathcal{T}\right) \middle| \mathcal{T}\right] = 1. \tag{2}$$

No subsidy: We begin with a benchmark schedule $\mathcal{T}_0$ in which no subsidy is provided:

$$\mathcal{T}_0 \equiv \left\{\tau_i = 0,\ \forall i \in \mathcal{I}\right\}.$$

This results in a linear demand curve given by $P = 2 - Q$. From Equation (2), the optimal price $\overline{p}$ is such that

$$\eta_0 \equiv \eta \left[\overline{p} \middle| \mathcal{T}_0\right] = 1,$$

where $\overline{p} \equiv p^* \left(\mathcal{T}_0\right)$.

Flat subsidy: Consider a schedule $\mathcal{T}_a$ under which all consumers receive a uniform subsidy τ:

$$\mathcal{T}_a \equiv \left\{\tau_i = \tau > 0,\ \forall i \in \mathcal{I}\right\}.$$

This raises consumer WTP to $\widetilde{\theta}_i = \theta_i + \tau$ and shifts the demand curve to $P = 2 - Q + \tau$, as illustrated in Panel A of Figure 1. The parallel shift reduces the price elasticity at the benchmark price $\overline{p}$ to

$$\eta_{a1} \equiv \eta \left[\overline{p} \middle| \mathcal{T}_a\right] = \frac{1}{1+\tau} < 1.$$

In response, the monopolist raises the price to $\widetilde{p}_a$ until the price elasticity is restored to one:

$$\eta_{a2} \equiv \eta \left[\widetilde{p}_a \middle| \mathcal{T}_a\right] = 1,$$

where $\widetilde{p}_a \equiv p^*(\mathcal{T}_a)$. As a result, the equilibrium price increases ($\widetilde{p}_a > \overline{p}$).

Targeted subsidy: Consider a schedule $\mathcal{T}_b$ with subsidies decreasing with consumers' baseline WTP:

$$\mathcal{T}_b \equiv \left\{\tau \cdot (\overline{\theta} - \theta_i),\ \forall i \in \mathcal{I}\right\},$$

where $\overline{\theta} = 1$ is the median WTP. Note that only consumers with WTP below median receive positive transfers.[8] This rotates the demand curve around the benchmark price $\overline{p}$, as shown in Panel B of Figure

[8]Consumers with $\theta_i > \overline{\theta}$ effectively face a tax. If, instead, these consumers received neither a tax nor a subsidy, the

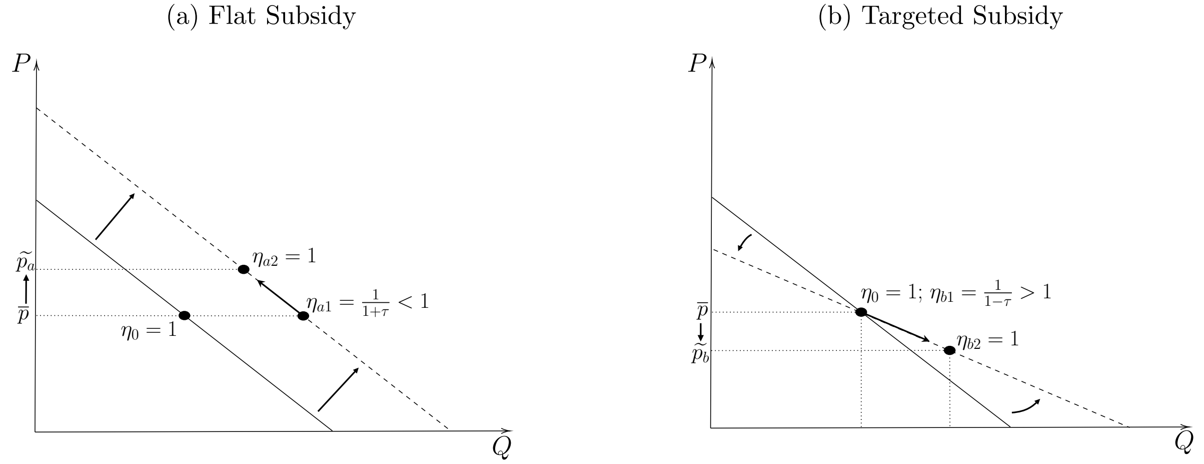

Figure 1: Flat and targeted subsidies

Notes: This figure illustrates how a subsidy τ_i affects the equilibrium price set by a single-product monopolist with zero marginal cost. Solid lines represent the baseline demand curve (no subsidy) and dashed lines represent demand with a subsidy. In both panels, $\overline{p}$ and $\widetilde{p}$ denote the equilibrium prices before and after the subsidy, respectively. Panel A depicts the case of a *flat subsidy*, which uniformly increases consumers' WTP, shifting the demand curve upward in parallel. This reduces the absolute price elasticity of demand, η, at the initial price $\overline{p}$ and induces the firm to raise its price to $\widetilde{p}_a$. Panel B considers a *targeted subsidy* that decreases with baseline WTP. This rotates the demand curve, making it flatter and increasing the absolute price elasticity of demand at $\overline{p}$. In response, the firm lowers its price to $\widetilde{p}_b$.

1. The new demand curve is flatter, increasing the price elasticity at $\overline{p}$ to

$$\eta_{b1} \equiv \eta\left[\overline{p}\middle|\mathcal{T}_b\right] = \frac{1}{1-\tau} > 1.$$

In response, the monopolist lowers the price to $\widetilde{p}_b$ until the price elasticity is restored to one:

$$\eta_{b2} \equiv \eta\left[\widetilde{p}_b|\mathcal{T}_b\right] = 1,$$

where $\widetilde{p}_b \equiv p^*(\mathcal{T}_b)$. As a result, the new equilibrium price decreases ($\widetilde{p}_b < \overline{p}$).

The intuition for why a price reduction is optimal follows from the impact of a targeted subsidy on the distribution of WTP across consumers. When the subsidy raises WTP among consumers who would not have purchased at baseline, it compresses the dispersion in WTP, and thus reduces the gap between these consumers and the original marginal consumer. That is, it shifts their WTP closer to $\overline{p}$ but still below it. Consequently, a small price reduction induces a substantial increase in demand, making a price decrease the firm's optimal response.

Other Subsidy Schemes: In both cases above, price changes are unambiguous: A flat subsidy shifts demand uniformly upward, prompting the monopolist to increase prices, whereas a targeted subsidy, as described above, rotates demand, leading the monopolist to reduce prices. More complex subsidy

resulting kink in the demand function would violate our differentiability assumptions. However, in this specific case, Equation (1) would still have a unique solution. Moreover, the equilibrium price remains unchanged regardless of whether above-median-WTP consumers face a tax, since inframarginal consumers do not influence the firm's optimal pricing decision.

designs may combine demand shifts and rotations, with the net price impact determined by the relative strength of the parallel shift (which raises prices) and the rotation (which lowers prices).

2.2 The price effects of targeted financial aid under a flexible demand system

We extend the linear-demand framework from Section 2.1 to accommodate more flexible demand systems and other forms of financial aid, such as student loans. This extended framework serves two key purposes. First, it decomposes the price effects of targeted aid into direct and composition effects. Second, it identifies the parameters that determine the relative magnitudes of these effects, and provides the motivation for the empirical analysis in Section 4.3.

Consider a market with only one college (a single-product monopolist) charging price p and with a marginal cost of c. There is a unit mass of students indexed by $i \in [0, 1]$ and characterized by χ_i. A fraction ρ of students receive financial aid, and students are ordered by their propensity to receive aid; that is, student i receives aid if $i \leq \rho$. For example, if aid is perfectly targeted by income, $i = 0$ is the poorest student. Students' utilities for attending college depend on their loan-holder status $k \in \{L, NL\}$, where L represents loan-holders and NL represents non-loan-holders.

We model financial aid as flexibly influencing enrollment decisions by allowing the utility of attending college to depend on loan status. Specifically, the utility of enrolling for a student i with loan-holder status k is denoted by u_{i1}^k, while the utility of not enrolling is u_{i0}. These utilities are drawn from a joint distribution $F_u(\cdot|p, \chi_i)$, which depends on student characteristics χ_i and tuition price p. The probability that a student enrolls in college, conditional on loan status k, characteristics χ_i, and price p, is given by

$$s_i^k(p) = \Pr[u_{i1}^k > u_{i0} \mid p, \chi_i, k].$$

From the firm's perspective, the optimal price, conditional on ρ, is

$$p^*(\rho) = \arg\max_p \; S(p, \rho) \cdot (p - c), \tag{3}$$

where total enrollment (i.e., market share), $S(p, \rho)$, is given by

$$\overbrace{S(p,\rho)}^{\text{Enrollment}} \equiv \overbrace{\int_0^\rho s_i^L(p)di}^{\text{with aid}} + \overbrace{\int_\rho^1 s_i^{NL}(p)di}^{\text{without aid}}, \tag{4}$$

and s_i^L and s_i^{NL} are the probabilities that student i enrolls in college with and without a loan, respectively.

To decompose the impact of student financial aid on prices into the direct and composition effects, we compute the effects of a marginal expansion in the financial aid program. In Supplemental Appendix

B, we show that differentiating Equation (3) with respect to ρ yields the following expression:

$$\frac{\mathrm{d}\log p^*}{\mathrm{d}\rho} = \Omega \cdot \left[\underbrace{\left(\eta_\rho^{NL} - \eta_\rho^{L}\right)}_{\text{direct effect}} - \overbrace{\left(\eta_\rho^{L} - \eta\right) \cdot \frac{s_\rho^{L} - s_\rho^{NL}}{s_\rho^{NL}}}^{\text{composition effect}} \right], \tag{5}$$

where η represents the price elasticity of overall demand S at price p^*. The enrollment probability of the marginal loan holder is denoted by s_ρ^L, while their enrollment probability in the absence of a loan would be s_ρ^{NL}. Similarly, η_ρ^L and η_ρ^{NL} denote the price elasticities of marginal loan holders with and without loans, respectively. The term Ω captures the role of price elasticity and demand curvature in passthrough. In particular, a more inelastic demand increases Ω, consistent with prior research on tax incidence under imperfect competition (Weyl and Fabinger, 2013).[9]

The elements in Equation (5) help us analyze the factors that determine the relative importance of the direct and composition effects. We discuss each in turn.

The direct effect captures how financial aid alters the price elasticity of demand for the marginal group of students receiving loans. Specifically, it depends on the difference in elasticity when these students hold a loan versus when they do not (i.e., $\eta_\rho^{NL} - \eta_\rho^L$). A larger difference strengthens the direct effect, pushing prices upward. The intuition behind this effect mirrors the forces discussed in Section 2.1 and illustrated in Panel A of Figure 1.

The composition effect is governed by two key components. First, the extent to which financial aid increases demand among the marginal group, relative to when they do not have a loan (i.e., $s_\rho^L - s_\rho^{NL}$); a larger increase in demand amplifies the composition effect. Second, the difference between the price elasticity of the marginal group when loans are provided and the overall elasticity, $\eta_\rho^L - \eta$. If the marginal group is more price-sensitive than the market average, the composition effect reduces prices; conversely, if they are less price-sensitive, it increases prices. The intuition behind the composition effect parallels the forces discussed in Section 2.1 and illustrated in Panel B of Figure 1. Expanding demand among a more price-sensitive group implies that small price reductions lead to substantial increases in enrollment, which flattens the demand curve and increases its elasticity. As a result, a more elastic demand induces the college to lower tuition.

The central insight of the analysis is that the price effects of targeted financial aid can be decomposed into direct and composition effects, with the latter depending on the targeting design. Both effects are governed by price elasticity and demand curvature, as captured by the term Ω, which aligns with

[9] Formally, the price elasticity of marginal loan holders, given $k \in \{L, NL\}$, is defined as $\eta_\rho^k \equiv -\frac{p}{s_\rho^k}\frac{\partial s_\rho^k}{\partial p}$. The term Ω is a function of market shares, price elasticity η, and the curvature of the demand function λ: $\Omega \equiv \frac{s_\rho^{NL}}{S}\frac{1}{\eta^2}\frac{1}{2-\lambda}$, where $\lambda \equiv S\frac{\partial^2 S}{\partial p^2}/(\frac{\partial S}{\partial p})^2$. We assume that $\lambda < 2$, which ensures that $\Omega > 0$. This restriction is a standard assumption in the literature (Weyl and Fabinger, 2013) and is satisfied by commonly used discrete choice models, including logit and nested logit, except in degenerate cases.

established results from the literature on incidence under imperfect competition (Delipalla and Keen, 1992; Anderson et al., 2001; Weyl and Fabinger, 2013; Miravete et al., 2018; Kroft et al., 2024a). However, prior studies have primarily focused on uniform taxes or subsidies that apply equally across individuals, overlooking the role of targeting in shaping incidence.

3 Setting, data, and sample

This section provides an overview of our empirical setting: the Brazilian higher education market and the federal student loan program. It also details the datasets used in our analysis.

3.1 Setting

Private higher education. The Brazilian higher education market comprises both public and private institutions, with private enrollment more than tripling over the past two decades. By 2014, 75% of the 8.8 million in-person students attended private colleges, and for-profit institutions accounted for 51% of private enrollment. The sector is highly concentrated, with multiple colleges often owned by the same firm: The top 10 firms operating in the private higher education sector account for 42% of total private enrollment, and the top 100 for 72%. The four largest firms (*Anima Holding*, *Kroton*, *SER Educacional*, and *Estácio*) are publicly traded and represent 29% of private enrollment. Unlike public institutions—which are tuition-free, more prestigious, and oversubscribed—private colleges charge tuition and tend to operate with substantial excess capacity. According to administrative data, private degree programs fill only 48% of their reported seats, on average, and nearly 90% enroll fewer than 80% of their capacity. As a result, most private colleges are effectively non-selective.

The federal student loan program (FIES). FIES provides subsidized loans to students attending private colleges. Established in 1999, the program was substantially expanded by a major reform in 2010. New loan contracts increased from fewer than 20,000 in 2009 to around 550,000 in 2014. For FIES-supported students, the government pays colleges directly, covering up to 100% of tuition, with repayment beginning 18 months after graduation or dropout. By 2014, FIES disbursed USD 5.8 billion annually—15% of the Ministry of Education's budget and roughly one-sixth of private colleges' revenues. In 2014, the median program enrolled about 25% of its students using FIES loans, with some programs relying on FIES for all of their enrollment.

To qualify for FIES, students were required to take the ENEM, a high-stakes exam also used for admission to federal universities, though no minimum score threshold was imposed. Eligibility and the share of tuition covered by aid were means-tested. Loans covered between 50% and 100% of tuition costs depending on the ratio of tuition costs to per capita monthly family income; students were ineligible if tuition represented less than 20% of their per capita monthly income.[10] In 2015, fiscal constraints led

[10] See Supplemental Appendix C for details.

to reforms that tightened FIES eligibility for both students and institutions. We describe these policy changes in Section 4.1.

Tuition discounts. In addition to government aid, private colleges frequently offer tuition discounts as a form of financial aid. Discounts are common: In 2014, 14% of incoming private college students received a discount, with an average discount rate of 33%. These discounts are not merit-based scholarships but rather part of a price discrimination strategy aimed at students who otherwise could not afford tuition.[11] Discounts are typically retained throughout enrollment, conditional on good academic standing.

Colleges do not explicitly target specific demographics for tuition discounts. Instead, they employ dynamic pricing strategies similar to those used by airlines and hotels. A prominent example is the use of online marketplaces, which function like travel fare aggregators (e.g., Expedia or Hotwire). These platforms post temporary tuition discount offers and require students to actively search for and secure them. Supplemental Appendix Figure A.1 illustrates the interface of such a platform. Offers on these marketplaces are typically short-lived. Using administrative records from *QueroBolsa*, the largest marketplace in the country, we find that 25% of offers are available for less than one week, and 50% are removed within three weeks.[12] Consistent with these institutional features, Supplemental Appendix E.2 presents descriptive evidence showing that discounts are not targeted on observable student characteristics. It is also important to note that FIES regulations explicitly prohibit colleges from denying tuition discounts to students receiving government loans.

3.2 Data

Student-level data. We construct a comprehensive dataset on students' college education, financial aid access, and demographics by merging data from three sources using individual-level identifiers. The first source is the Census of Higher Education, which covers the universe of students enrolled in higher education. It provides detailed information on students' colleges, degree programs, and whether they receive tuition discounts, though the exact discount amounts are not reported. The second source is administrative records from FIES, which tracks students who receive loans each year and the loan amounts. The third source is the ENEM dataset, which includes all students taking Brazil's national standardized exam annually. It provides student test scores and responses to a comprehensive socioeconomic survey. The survey collects demographics such as race, gender, family income, high school type, parental education, and students' plans to apply for a FIES loan.

Tuition fees. In Brazil, universities are not required to report tuition fees to regulatory authorities, so

[11] Students receiving discounts have, on average, 25% lower per capita family income and 0.45 standard deviations lower ENEM scores, and are 25% less likely to have college-educated parents.

[12] QueroBolsa accounts for 18% of all tuition discounts.

we rely on four data sources to construct tuition fee records. The first two sources use administrative data from government-funded grants (PROUNI) and loans (FIES) programs. Records from the National Education Fund (FNDE) track payments made for students in these programs, which allows us to calculate tuition fees for participating institutions. The third source is a nationally representative survey by Hoper, a consultancy specializing in higher education, which provides tuition data across a broad sample of institutions. The fourth is administrative data from QueroBolsa, Brazil's largest degree search platform, which offers detailed tuition prices for participating degree programs. Both the Hoper and QueroBolsa datasets are novel and typically unavailable to researchers, but we secured access through partnerships with these companies. Supplemental Appendix D describes how we combine these sources to construct full and discounted tuition prices for each program. We recover year-specific prices for 95% of degree–year observations, covering 98% of total enrollment. Supplemental Appendix E.1 shows that approximating each degree with only two prices—full and discounted—is a reasonable representation of the data.

3.3 Sample

We define a *degree* as a combination of major, institution, and shift.[13] For instance, an economics program at Pythagoras University offered during the night shift constitutes a distinct degree. The term *college* refers to a firm that may operate multiple degree programs across different campuses and regions. We define *regions* according to the classification of the Brazilian Institute of Geography and Statistics (IBGE), which groups Brazil's 5,568 municipalities into 137 meso-regions based on geographic proximity and shared socioeconomic characteristics.

Our analysis covers the period from 2012 to 2017. For computational tractability, we impose several sample restrictions. While our dataset includes all ENEM takers, we limit the analysis to students residing in regions with at least 5,000 ENEM takers and 1,000 incoming college students per year. This yields a sample of 69 regions. We further restrict the data to in-person private degree programs with at least 15 incoming students annually. Students enrolled in degrees that fall below this threshold remain in the dataset but are assigned to the outside option. The same applies to students who enroll in degrees outside the region where they took the ENEM.[14] Additionally, we exclude students with ENEM scores below 400 or above 700, as they comprise only about 3% of private college enrollment and are not central to colleges' pricing decisions.[15] Finally, we exclude students receiving government grants from the sample.[16] After these restrictions, our final sample consists of over 20 million ENEM takers and

[13]Shift refers to the time of day the degree is offered, either daytime or nighttime.

[14]Fewer than 5% of students enroll in a private degree program outside the region where they took the ENEM, indicating that cross-region enrollment in the private sector is limited. This supports our treatment of region-year combinations as separate markets.

[15]Students scoring above 700 represent just 1.5% of ENEM takers, with 75% of them enrolling in public universities. Those scoring below 400 make up 6% of test-takers, but fewer than 5% pursue higher education.

[16]The federal government's grant program, PROUNI, provides scholarships to low-income students enrolled in private

13,567 unique degree programs offered by 695 colleges, covering 88% of ENEM takers nationwide.

Throughout the paper, we index degrees by j, regions by r, and years by t. Each degree j belongs to a unique region r but is observed across multiple years t.

4 Descriptive evidence

4.1 Background: The 2015 FIES Reform

Our empirical analysis centers on the 2015 FIES reform—a major policy shift driven by federal budget constraints that substantially tightened loan eligibility criteria for both students and institutions. For students, the reform introduced a maximum per capita family income threshold of 2.5 times the federal minimum wage and a minimum ENEM score of 450.[17] For institutions, the reform introduced a cap on the number of FIES-funded students per degree, with priority given to high-quality programs, health-related fields, and institutions located in low-income regions.[18]

In most degree programs, demand for loans exceeded the cap and resulted in the allocation of loans through a deferred acceptance mechanism that relied on ENEM scores to create degree-specific cutoffs for loan eligibility. Students scoring above the cutoff qualify for loans and retain them as long as they stay in the same program. It is important to note that admission to private degree programs in Brazil is not centralized, and the described mechanism only applies to federal financial aid allocation. Also, most private sector degrees are non-selective, which means that students who do not qualify for a loan due to a low ENEM score can still enroll by paying tuition out of pocket.

The FIES reform led to a sharp contraction in the number of new loans, dropping from around 550,000 in 2014 to approximately 150,000 in 2017. Figure 2 presents the key aggregate trends surrounding the reform. Panel A shows this steep drop in loan volume alongside a decline in posted tuition prices, which fell by 5.2% over the same period. This decline in prices suggests weakened demand for higher education following the reduction in loan availability. Note that posted tuition only declines in 2016, as tuition fees for 2015 were set prior to the announcement of the policy change.

Panel B shows the impact of the reform on the stock prices of the four largest firms in the higher education sector (Anima Holding, Kroton, SER Educacional, and Estácio). On the day the reform was announced, stock prices fell sharply, which indicates that the policy change was largely unexpected and had significant implications for market expectations regarding future profits. The persistent lower stock

colleges. Between 2012 and 2017, 4.1% of incoming students receive a full PROUNI grant and 1.3% receive a partial grant. Participating colleges are exempt from federal taxes and must reserve a fixed share of seats for grant recipients. Full-grant students pay zero tuition and therefore do not affect colleges' pricing incentives. Partial-grant students pay the remaining tuition and could affect pricing incentives, but they represent a small share of the market. For simplicity, we exclude all PROUNI recipients from the estimation sample.

[17] In our sample, 68% of ENEM takers report a family income below 2.5 times the minimum wage. An ENEM score of 450 corresponds approximately to the 25th percentile. See Supplemental Appendix C.1.2 for details on the formula determining the share of tuition covered.

[18] For a detailed description of the loan allocation process, see Supplemental Appendix C.2.

prices in subsequent months suggest that the reform triggered a structural reassessment of the sector's financial outlook.

(a) Number of loans and tuition costs

(b) Stock prices

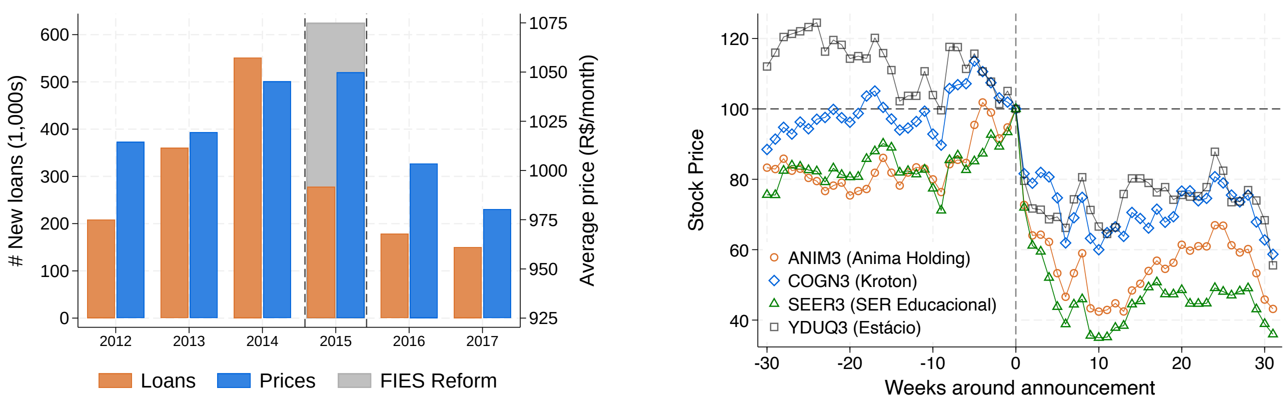

Figure 2: The 2015 FIES Reform: Aggregate Trends

Notes: This figure illustrates aggregate trends surrounding the 2015 FIES reform. Panel A presents the number of incoming students receiving FIES loans (orange bars) and the average monthly full tuition price across all private degrees (blue bars) for each year. The shaded gray area indicates the year the reform was being implemented. Tuition prices are computed as the enrollment-weighted average of full prices for each degree and are deflated to 2014 price levels. Panel B illustrates the stock prices of the four largest higher-education conglomerates, which collectively serve 30% of students with federal loans. Stock prices are normalized to 100 on the day preceding the reform announcement, and the x-axis is set so that week 0 corresponds to the announcement week. Stock price data were sourced from GoogleFinance.

4.2 Effects of financial aid on tuition and enrollment

We use the FIES reform to examine how tuition fees and enrollment respond to changes in financial aid availability. To measure exposure to the reform, we exploit the introduction of caps on the number of students with loans who are allowed to enroll in each degree program. Importantly, these caps were set independently of the number of students with loans in each degree before the reform.[19] As a result, degree programs with higher pre-reform loan enrollment experienced larger reductions in loan availability.

To capture this variation, we follow an approach similar to Black et al. (2023) and construct a degree-level exposure measure that reflects the extent to which the cap constrained loan availability in each program. Specifically, we define exposure as

$$\text{Exp}_j = \frac{\max\{N^L_{j,2012} - \bar{N}_j, 0\}}{N_{j,2012}}, \tag{6}$$

where $N^L_{j,2012}$ represents the number of FIES-funded students in degree j in 2012, $N_{j,2012}$ the total

[19]The number of students who applied for FIES loans in a given region in 2015 enters the cap calculation, but not the specific number allocated to each degree within that region. See Supplemental Appendix C.2 for details, including references to the laws that define the allocation formulas. All analyses using this measure of exposure include region-year fixed effects to address any potential endogeneity arising from region-level variation in the total number of loans.

enrollment in the same year, and $\bar{N}_j$ the degree-specific loan cap imposed by the policy.[20] Supplemental Appendix Figure A.2 shows that our exposure measure, Exp_j, is strongly associated with the change in the share of students with loans between 2014 and 2017: A one standard deviation increase in exposure corresponds to a 16 percentage point decline in loan usage, confirming that more exposed degrees experienced larger reductions in FIES-funded enrollment.

A potential concern with this exposure measure is mean reversion: Degree programs with high pre-reform loan uptake may trend toward the mean over time. To mitigate this, we define exposure using 2012 data, 3 years before the reform. Also, we restrict our estimation to 2013 onward to avoid confounding shocks that may have influenced both high loan uptake and other variables in 2012.

Building on this variation, we examine how the 2015 reform affected a range of outcomes across degrees with different levels of exposure. Specifically, we estimate the following regression:

$$\text{Outcome}_{jt} = \gamma_j + \gamma_{rt} + \sum_l \beta_l \cdot \mathbb{1}\{t = l\} \cdot \text{Exp}_j + \epsilon_{jt}, \tag{7}$$

where β_l are our coefficients of interest and capture the effect of exposure over time. The term γ_j represents degree fixed effects, γ_{rt} accounts for region-year fixed effects, and ϵ_{jt} is an idiosyncratic error term.

Figure 3 presents the estimated coefficients for several outcomes. The results show parallel pre-trends, and thus support the validity of attributing post-2015 changes to the reform and mitigate concerns about mean reversion. Panel A confirms that more exposed degree programs experienced a decline in the share of students using loans. Panel B shows that these programs also reduced tuition fees, suggesting that greater loan availability is associated with higher prices. As discussed in Section 4.1, tuition prices for 2015 were already set when the reform was announced, so tuition effects are only observed starting in 2016. Panel C reveals that more exposed degrees also faced a decline in enrollment. The effects are economically meaningful: in 2017, a one standard deviation increase in Exp_j is associated with a 2.1% decline in tuition and a 9.2% drop in enrollment. Finally, Panel D shows that more exposed degrees experienced an increase in the number of students paying out of pocket—likely reflecting both the tuition reduction and the enrollment of students who take out loans when available but still choose to enroll without them.

In parallel with the FIES reform, the federal government implemented a broader fiscal consolidation in 2015, generating economy-wide adjustments. Although FIES was the only major change in higher-education policy during this period,[21] macroeconomic conditions could nonetheless affect the estimates

[20] We do not observe caps for programs that did not participate in FIES. Instead, we predict their caps based on field of study and region, which are the primary determinants of cap allocation in the government's formula.

[21] Beyond FIES, the two main federal higher-education expenditures are PROUNI grants and public universities. PROUNI provides tuition scholarships for students enrolled in private colleges, while public universities are tuition-free and publicly funded. There were no major policy changes affecting PROUNI or public universities around the 2015 FIES reform. Consistent with this, enrollment in PROUNI and public institutions remains stable over this period, whereas FIES participation drops sharply in 2015 (see Supplemental Appendix Figure A.3).

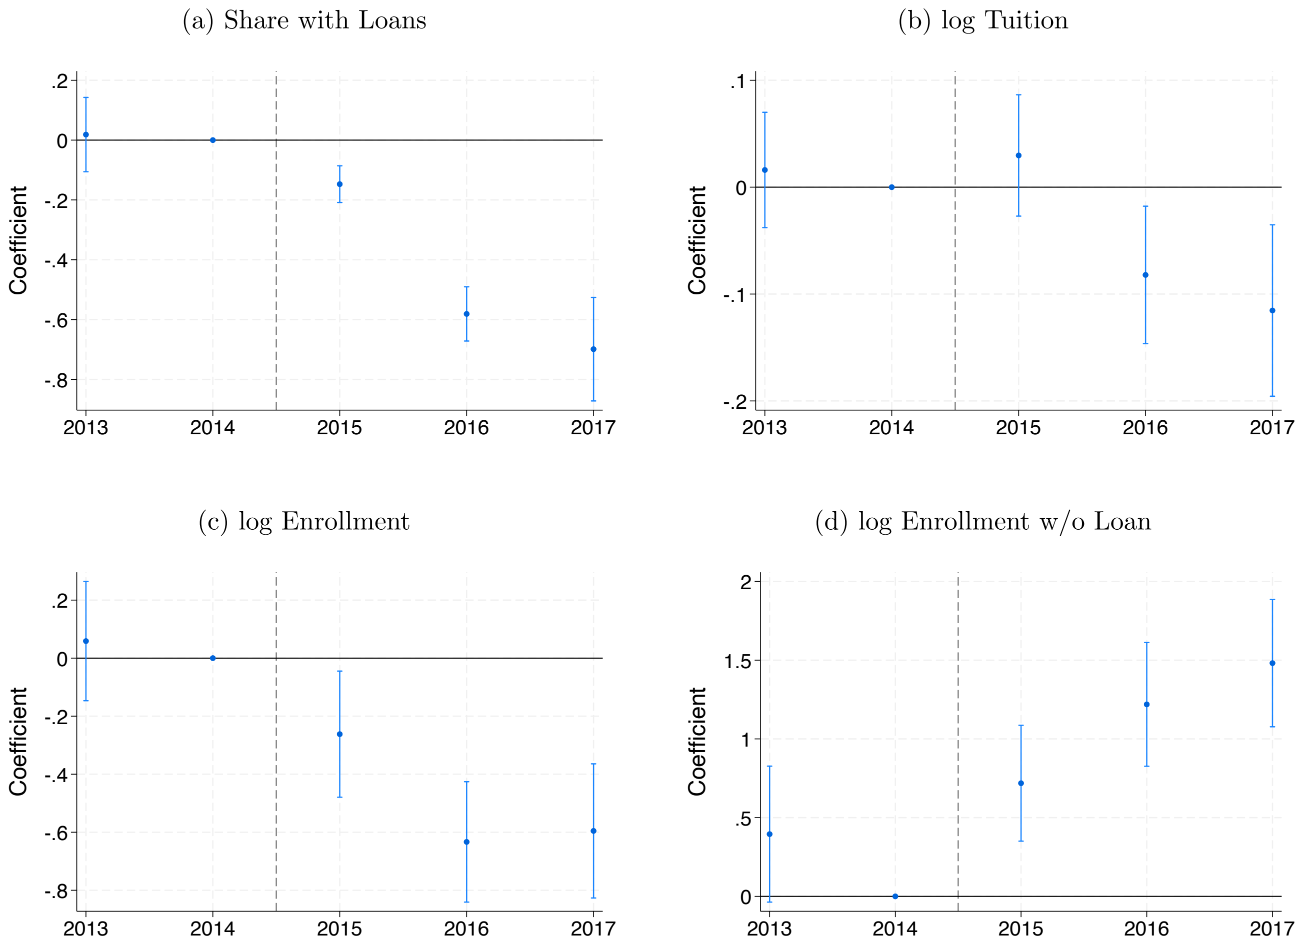

Figure 3: Differential Exposure to the 2015 FIES Reform

Notes: This figure reports OLS estimates of β_l from Equation (7). The error bars represent 95% confidence intervals and standard errors are clustered at the college-year level. The outcome in Panel A is the share of incoming students with loans; in Panel B, the log monthly tuition fee; in Panel C, the log number of incoming students; and, in Panel D, the log number of incoming students not using a government loan. Tuition prices are calculated as enrollment-weighted averages of full and discounted prices for each degree-year and are deflated to 2014 price levels. The vertical line marks the reform's announcement.

in Figure 3. For instance, programs more exposed to the reform may be concentrated in fields linked to occupations that were differentially affected by other contemporaneous policies or economic shocks. To assess whether our results reflect such concurrent factors, we augment Equation (7) with region–major–year fixed effects, which flexibly control for region-specific trends across fields of study. The resulting estimates, reported in Supplemental Appendix Figure A.4, are consistent with our baseline results. This robustness suggests the effects are not driven by other simultaneous policy changes or by differential exposure of more FIES-intensive degrees to broader macroeconomic shocks.

Another potential concern is the expansion of online education during our sample period. If degrees more exposed to the FIES reform were also more exposed to online competition, our estimates could be biased. Supplemental Appendix F addresses this concern by controlling for exposure to online competitors. We find that greater exposure to online programs is associated with lower tuition for in-person degrees, consistent with increased competitive pressure, as documented by Barahona et al. (2025). However, the estimated effect of the FIES reform is virtually unchanged once this control is included, suggesting that the expansion of online education is not driving our results.

4.3 Direct and composition effects

We now examine whether the tuition adjustments observed after the FIES reform align with the predictions of our conceptual framework. In particular, we test whether tuition fees respond differently in markets in which the composition effect is expected to be more pronounced. According to our framework, the impact of financial aid on tuition depends on the difference between the overall price elasticity of demand and the elasticity among aid recipients—captured by the term $\eta_\rho^L - \eta$ in Equation (5). The theoretical prediction is that a larger value of $\eta_\rho^L - \eta$ strengthens the composition effect, thereby offsetting the downward pressure on prices from the direct effect of reduced loan availability.

A key challenge, however, is that these elasticities are not directly observable. While we estimate them formally later in the paper using an equilibrium model of supply and demand for higher education, a descriptive analysis requires a tractable approximation based on observable data. Specifically, we assume that the price elasticity of demand is inversely related to income and use the average income of enrolled students to proxy for the overall elasticity. Similarly, we approximate the elasticity among aid recipients using the average income of FIES loan recipients. Based on this logic, we construct $\Delta\eta_j$ as the log difference between the average income of all students at degree j's college and the average income of FIES loan recipients at the same college, using 2012 data—the first year of our sample. We interpret this measure as a proxy for the elasticity gap, $\eta_\rho^L - \eta$.[22]

This measure provides meaningful insight into differences in price elasticity under two assumptions. First, income must serve as a valid proxy for price sensitivity.[23] Second, the income distribution of enrolled students should mirror the characteristics of the broader pool of prospective students that a college might attract. When these conditions are met, a larger gap between the average student's income and that of loan recipients should correspond to a greater difference in the respective price elasticities.

Using this variable, our model predicts that at institutions where loan recipients have incomes similar to those of the average student, increased financial aid availability should lead to higher tuition. In contrast, at colleges where loan recipients come from substantially lower-income backgrounds, an expansion in loan usage should result in a more muted tuition response—or even a decline—due to the stronger composition effect.

To empirically assess the role of this mechanism, we leverage the FIES reform and extend Equation

[22] Supplemental Appendix Figure A.5 shows substantial variation in $\Delta\eta_j$. Among degrees in the top 5% of this distribution, loan recipients are, on average, 75% poorer than the overall student body. In contrast, in the bottom 5% they are only 8% poorer.

[23] This assumption is supported by the model estimates reported in Supplemental Appendix Table A.2, which indicate that lower-income students are more price sensitive than higher-income ones.

(7) to allow for heterogeneity in tuition responses driven by differences in $\Delta\eta_j$. Specifically, we estimate:

$$\log p_{jt} = \gamma_j + \gamma_{rt} + \beta_0 \cdot \text{post}_t \cdot \text{Exp}_j + \beta_C \cdot \text{post}_t \cdot \text{Exp}_j \cdot \Delta\eta_j + \epsilon_{jt}, \tag{8}$$

where $\log p_{jt}$ denotes the log of the average tuition fee for degree j in year t, post_t is a dummy variable equal to one for years following the 2015 reform, and the remaining terms follow the notation from Equation (7). Our key parameters of interest are β_0 and β_C. The coefficient β_0 captures the average association between loan usage and tuition prices. Meanwhile, β_C measures the interaction between loan usage and income differences between students with and without loans, which captures the composition effect of financial aid targeting.

Our estimation sample excludes 2012 for the reasons discussed in Section 4.2. In addition, in 2015, colleges had already set their prices when the reform was announced but students had not yet made their enrollment decisions, which makes that year difficult to interpret. As a result, our preferred specification also excludes 2015 and focuses on two pre-reform years (2013 and 2014) and two post-reform years (2016 and 2017). As a robustness check, Supplemental Appendix Table A.1 reports estimates using the full sample (2012-2017). The results closely align with our baseline findings, which indicates that these sample restrictions are not quantitatively consequential.

Table 1 reports OLS estimates of Equation (8). Column (1) examines the relationship between financial aid availability and tuition without accounting for heterogeneity in $\Delta\eta$. The result indicates that more exposed degrees reduced tuition following the reform. This is consistent with Figure 3, Panel B, and reinforces the finding that greater loan availability is associated with higher tuition. The effect is substantial: A one standard deviation increase in exposure corresponds to a 2.0% decline in tuition post-reform.[24]

Column (3) introduces heterogeneity by $\Delta\eta$ and reveals a substantial composition effect. As $\Delta\eta$ increases, the relationship between exposure and tuition becomes less negative. The magnitude of this effect is considerable: Our estimates indicate that for degrees in the bottom 5% of $\Delta\eta$, a one standard deviation increase in exposure is associated with a 3.9% *decrease* in tuition post-reform. In contrast, for degrees in the top 5%, the same increase in exposure corresponds to a 0.1% *increase* in tuition. This suggests that in high $\Delta\eta$ programs, the composition effect outweighs the direct effect.[25]

To address potential confounding factors, Columns (2) and (4) introduce region-major-year fixed effects to control for broader policy changes in 2015 that may have impacted the education sector. Also, Column (5) accounts for potential pre-trends by including year fixed effects interacted with $\Delta\eta$. Across all specifications, the results remain consistent and statistically robust, reinforcing the validity of our

[24] As discussed in Section 4.2 and shown in Supplemental Appendix Figure A.2, a one standard deviation increase in exposure corresponds to a 16 percentage point decline in loan usage between 2014 and 2017.

[25] The differences remain substantial at less extreme points of the $\Delta\eta$ distribution. A one standard deviation increase in exposure is associated with a 2.8% decrease in tuition for degrees in the bottom 25% of $\Delta\eta$, compared with a 0.9% decrease for those in the top 25%.

findings.

Overall, our results align well with the predictions of our conceptual framework. The findings highlight the critical role of financial aid targeting in shaping tuition responses, which provides suggestive evidence of a composition effect. These results suggest that financial aid programs can have markedly different price effects depending on the characteristics of aid recipients, underscoring the importance of considering student composition when designing and evaluating financial aid policies.

Table 1: Direct and Composition Price Effects of the 2015 FIES Reform

Dep var: $\log p_{jt}$	(1)	(2)	(3)	(4)	(5)
post × Exp	-0.106***	-0.102***	-0.099***	-0.100***	-0.098***
	(0.028)	(0.020)	(0.022)	(0.019)	(0.020)
post × Exp × $\Delta\eta$			0.503***	0.395***	0.359**
			(0.151)	(0.142)	(0.155)
Observations	13447	11695	13447	11695	11695
Degree FE	X	X	X	X	X
Region-Year FE	X	X	X	X	X
Major-Region-Year FE		X		X	X
Δ Inc-Year FE					X

Notes: This table reports OLS estimates of Equation (8). The dependent variable is the log of tuition for each degree-year. Tuition is calculated as the enrollment-weighted average of full and discounted prices, deflated to 2014 levels. "Exposure" measures a degree's exposure to the 2015 FIES reform; "post" indicates post-reform years; and $\Delta\eta$ is the log difference between the average income of all students and the average income of FIES loan recipients at each college in 2012. The sample includes two pre-reform years (2013 and 2014) and two post-reform years (2016 and 2017). Standard errors, clustered at the college-year level, are in parentheses. Asterisks indicate statistical significance: * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$.

5 An equilibrium model of higher education with targeted aid

We now introduce an expanded equilibrium model of higher education, building on the framework outlined in Section 2.2. While the core economic forces remain the same, the extended model incorporates features that more accurately reflect the complexities of higher education markets. In particular, we allow for multiple colleges competing within each market, each offering a portfolio of degree programs and engaging in price discrimination. These extensions enable us to conduct counterfactual simulations and examine the equilibrium effects of alternative financial aid designs, such as need-based and merit-based targeting.

5.1 Demand

Our demand model consists of a continuum of risk-neutral students, indexed by $i \in \mathcal{I}$, each characterized by their household income w_i and exam score h_i. We denote geographic regions by r and years by t, with each region-year combination treated as a distinct market indexed by rt. Each degree j is specific to one region r and may span multiple years t. Each student belongs to only one market and makes a one-time enrollment decision. The sets of students and degrees in each market are represented by $\mathcal{I}_{rt}$ and $\mathcal{J}_{rt}$, respectively. Degrees are non-selective, meaning that students can enroll in any degree available in their market or choose the outside option.[26]

We assume that the utility derived by student i from enrolling in degree j consists of three main components. The first two capture the experience utility associated with the degree: a mean utility term and an idiosyncratic component that reflects individual heterogeneity. Together, these two terms account for factors such as expected labor market returns and the consumption value of education. The third component captures the disutility associated with the financial cost of attending the program.

Formally, the utility function is given by

$$U_{ijrt} = \underbrace{\delta_{jrt}}_{\text{mean utility}} + \underbrace{\mu_{ijrt}}_{\text{individual utility}} - \underbrace{(\alpha^0_{ijrt} + \alpha^1_{ijrt} \cdot p_{ijrt})}_{\text{financial costs}}. \tag{9}$$

The first component, δ_{jrt}, represents the mean utility value of the degree and depends on program-specific characteristics and broader market conditions. We parameterize it as

$$\delta_{jrt} = \delta_{jr} + \delta_{rt} + \xi_{jrt}, \tag{10}$$

where δ_{jr} is a degree-region fixed effect that absorbs time-invariant characteristics of degree j in region r. The term δ_{rt} accounts for region- and time-specific market conditions that influence demand. Lastly, ξ_{jrt} represents unobserved, time-varying demand shocks that affect the degree's overall attractiveness.

The second component, μ_{ijrt}, captures individual-specific utility derived from the degree, beyond the mean value. We define this utility component, which depends on individual characteristics, as

$$\mu_{ijrt} = \beta^h \cdot h_i + \beta^w \cdot w_i + \varepsilon_{ijrt}, \tag{11}$$

where β^h and β^w measure the impact of a student's exam scores h_i and income w_i on their baseline utility relative to the outside option—defined as the utility of not enrolling in a private degree program, with $U_{i0rt} = \varepsilon_{i0rt}$.[27] Finally, ε_{ijrt} is a student-specific demand shock, which we assume follows an extreme value distribution and allows for unobserved heterogeneity in individual preferences.

The third component, $\alpha^0_{ijrt} + \alpha^1_{ijrt} p_{ijrt}$, captures the financial cost of enrolling in a given program. We allow this component to depend on tuition through α^1_{ijrt}, which reflects the student's price sensitivity.

[26] The outside option includes both students who do not enroll in college and those who enroll in a public institution. Supplemental Appendix G shows that this formulation is equivalent to a model that explicitly incorporates public colleges as separate alternatives. It also documents that our main results are robust to excluding students who enroll in public institutions from the estimation sample.

[27] In our context, it is essential to link utility to scores and income, as we do in Equation (11), because public universities—included in the outside option—admit students based on exam scores, h. Also, income, w, might influence preferences for public universities. Thus, a student's utility from the outside option depends directly on their exam performance and income.

The term α^0_{ijrt} captures fixed financial burdens of enrollment—such as living expenses or application fees—that affect enrollment decisions regardless of tuition.

The price, p_{ijrt}, faced by student i depends on whether they qualify for discounts, as specified below:

$$p_{ijrt} = p_{jrt}\big(D_{ijrt}\big) = \begin{cases} p^F_{jrt}, & \text{if } D_{ijrt} = 0, \\ p^D_{jrt}, & \text{if } D_{ijrt} = 1, \end{cases}$$

where p^F_{jrt} and p^D_{jrt} are the full and discounted tuition rates for degree j in region r and year t, and D_{ijrt} indicates whether student i has a discount in degree j in region r and year t.[28]

We allow student loans to influence the financial-cost component of utility through two channels: by altering price sensitivity, α^1_{ijrt}, and by affecting the baseline utility cost of education financing, α^0_{ijrt}. Loans reduce price sensitivity by lowering the present value of expected tuition payments through two mechanisms: subsidized interest rates and the possibility of loan default. Also, loans may influence the baseline financing cost by easing liquidity constraints and serving as a behavioral nudge that encourages enrollment.

We capture these various channels by allowing the financial parameters, α^0_{ijrt} and α^1_{ijrt}, to flexibly depend on the student's loan status and student characteristics. Specifically, we define:

$$\alpha^0_{ijrt} = \alpha^0_L \cdot L_{ijrt} + \alpha^0_{wL} \cdot w_i \cdot L_{ijrt} \tag{12}$$

$$\log \alpha^1_{ijrt} = \alpha^1 + \alpha^1_w \cdot w_i + \alpha^1_L \cdot L_{ijrt} + \alpha^1_{wL} \cdot w_i \cdot L_{ijrt}, \tag{13}$$

where L_{ijrt} indicates whether student i would use a loan if they enroll in degree j in region r and year t. The parameters α^0_L and α^0_{wL} capture how baseline utility depends on loan status and its interaction with income. The terms α^1 and α^1_w reflect average price sensitivity and its variation with income, while α^1_L and α^1_{wL} allow loans to affect price sensitivity differently across income levels.

Students choose among the degrees available in their market or the outside option to maximize their utility. Therefore, enrollment decisions are given by

$$q_{ijrt} = \mathbb{1}\big\{U_{ijrt} > U_{ikrt}, \forall\, k \in \mathcal{J}_{rt}\big\},$$

where q_{ijrt} represents whether student i enrolls in degree j in market rt. We can then calculate the share of students enrolled in degree j in market rt with financial aid status (l, d) as

$$s_{jrt}(l,d) = \mathbb{E}\left[\mathbb{1}\big\{L_{ijrt} = l\big\} \cdot \mathbb{1}\big\{D_{ijrt} = d\big\} \cdot q_{ijrt} \big| i \in \mathcal{I}_{rt}\right], \tag{14}$$

where $\mathbb{E}$ denotes the expectation operator over preference shocks, ε_{ijrt}, and financial aid shocks,

[28] Supplemental Appendix E.1 shows that approximating within-degree price variation with a two-price structure provides a good fit to the data.

$\vartheta^L_{ijrt}, \vartheta^D_{ijrt}$. The total market share for degree j in market rt is given by

$$s_{jrt} = \sum_{l \in \{0,1\}} \sum_{d \in \{0,1\}} s_{jrt}(l, d).$$

5.2 Financial aid allocation

Financial aid is allocated at the student-degree pair level, and students choose their enrollment knowing both their loan status, L_{ijrt}, and discount status, D_{ijrt}, for every degree. We model financial aid allocation using the following latent variable framework:

$$L_{ijrt} = \mathbb{1}\{\rho^L_w w_i + \rho^L_{ft} \geq \vartheta^L_{ijrt}\} \cdot \mathbb{1}\{h_i \geq \overline{h}_{jrt}\} \quad (15)$$

$$D_{ijrt} = \mathbb{1}\{\rho^D_w w_i + \rho^D_{ft} \geq \vartheta^D_{ijrt}\}, \quad (16)$$

where ϑ^L_{ijrt} and ϑ^D_{ijrt} are idiosyncratic error terms that follow a logistic distribution and capture unobserved heterogeneity in students' propensity to search for or apply for a loan or discount for a given degree program. Students' likelihood of receiving loans and discounts also depends on their income through the parameters ρ^L_w and ρ^D_w, reflecting the fact that both eligibility and the propensity to seek financial aid may vary with household income. In addition, the government sets degree-region-specific loan score thresholds, denoted by $\overline{h}_{jrt}$.[29]

The parameters ρ^L_{ft} and ρ^D_{ft} capture systematic unexplained variation in the share of students receiving loans and discounts within a college f. This specification is motivated by the substantial variation in financial aid usage across degree programs even after controlling for student characteristics. Such variation may reflect factors not explicitly modeled, including social stigma associated with receiving aid (Moffitt, 1983; Finkelstein and Notowidigdo, 2019) and differences in students' access to information about financial aid (Castleman and Page, 2015).[30]

Importantly, in this formulation L_{ijrt} and D_{ijrt} should be interpreted as the student *receiving* financial aid, which in practice requires both *eligibility* and *take-up*. Because our primary objective is to understand supply-side responses to financial aid, this reduced-form approach provides a parsimonious way to capture how financial aid shifts the demand curve and, in turn, how colleges adjust prices. One implication of this formulation is that aid receipt is modeled independently of tuition levels. In reality, take-up decisions may respond to prices—for example, a student who declines a government loan when

[29] In practice, the government sets caps on the number of loans for each degree (see Section 4.1). These caps generate degree–year-specific score cutoffs that clear the market. To keep the model computationally tractable, we treat the resulting degree–year cutoffs $\overline{h}_{jrt}$ as given rather than modeling them as equilibrium outcomes.

[30] Stigma may be lower at colleges where aid usage is more common, making students more willing to accept aid when surrounded by peers who do the same. Access to information can also vary across institutions. As discussed in Section 3, many students rely on online platforms such as QueroBolsa to learn about financial aid opportunities. Colleges with lower name recognition tend to attract more students through these platforms, which can lead to higher aid usage. Consistent with this pattern, discount shares are negatively correlated with program age—a proxy for reputation. On average, an additional 10 years of program age is associated with a 2.1 percentage point decline in discount shares. Among the oldest 5% of programs, the average discount share is 16%, compared to 26% among the youngest 5%.

tuition is low due to stigma may choose to take it up if prices increase. Our framework abstracts from this margin in order to keep the analysis tractable.

5.3 Supply

Each college f offers a set of degrees $\mathcal{J}_{ft}$ in each year t and each degree is associated with a time-varying marginal cost c_{jrt}. In each year t, colleges choose both full and discounted prices for each of their degrees to maximize expected profits:

$$\max_{\{p^F_{jrt}, p^D_{jrt}\}_{j \in \mathcal{J}_{ft}}} \sum_{j \in \mathcal{J}_{ft}} s^F_{jrt} \cdot (p^F_{jrt} - c_{jrt}) + s^D_{jrt} \cdot (p^D_{jrt} - c_{jrt} + \kappa_{jrt}), \tag{17}$$

where $s^F_{jrt} \equiv \sum_{l \in \{0,1\}} s_{jrt}(l, 0)$ and $s^D_{jrt} \equiv \sum_{l \in \{0,1\}} s_{jrt}(l, 1)$ represent the enrollment of students paying full and discounted prices, respectively. The term κ_{jrt} captures degree-region-year-specific factors that shift the relative pricing of discounted versus full-price enrollment (e.g., branding or reputation differences across degrees), effectively acting as a reduced-form shifter of the profitability of discounted enrollment.

We decompose marginal cost as follows:

$$c_{jrt} = c_{jr} + c_{rt} + \omega_{jrt}, \tag{18}$$

where c_{jr} are degree-region fixed effects, c_{rt} are market fixed effects, and ω_{jrt} is a degree-region-year cost shock.

A total of 42% of degrees in our sample are offered by institutions legally classified as not-for-profit. A natural concern is whether these institutions follow a different objective function. Empirically, however, we find no evidence that they adopt systematically different pricing strategies from for-profit institutions. Although degrees offered by not-for-profit colleges are, on average, 11.8% more expensive, this difference is fully explained by degree age, major, and shift. We likewise find no meaningful differences in price discrimination: the average share of students receiving discounts is 19% in not-for-profit colleges, compared with 20.7% in for-profit institutions. Indeed, existing research notes that many not-for-profit colleges in Brazil operate in practice much like for-profit institutions (Schwartzman, 2026). These patterns support modeling all colleges under a common profit-maximization framework, regardless of legal ownership status.

As discussed in Section 3.1, tuition discounts are prevalent in the Brazilian higher education market. As a result, students with and without loans may face different net prices, which weakens the composition effect. To capture this interaction between loans and discounts, our framework departs from standard Bertrand–Nash oligopoly models (e.g., Berry et al., 1995; Nevo, 2001) by allowing colleges to set two prices for each program: one with discounts and one without. Nonetheless, the mechanics of the model remain similar to the Bertrand–Nash framework, since enrollment with and without discounts can be treated as distinct products offered by the same college.[31]

To map the standard framework (one price per product) to our model (two prices per product), we introduce some notation. Let $\vec{s}_{rt}$ be a vector that collects the market share of each degree in market rt, with and without discounts:

$$\vec{s}_{rt} \equiv \left(s^D_{jrt},\ s^F_{jrt}\right)'_{j\in\mathcal{J}_{rt}}$$

Note that $\vec{s}_{rt}$ contains twice as many elements as there are degrees in market rt. Analogously, $\vec{p}_{rt}$, $\vec{c}_{rt}$, and $\vec{\kappa}_{rt}$ collect prices, marginal costs, and κ_{jrt}'s of degrees in market rt, and are ordered as $\vec{s}_{rt}$. The elements of these vectors are indexed by ι, so $s_{\iota rt}$ denotes the ιth element of $\vec{s}_{rt}$, and similarly for $p_{\iota rt}$, $c_{\iota rt}$, and $\kappa_{\iota rt}$. Finally, each element $d_{\iota rt}$ of $\vec{d}_{rt}$ is a dummy equal to one if element ι refers to the discounted version of a degree.

Taking the first-order conditions from Equation (17), the optimal pricing equation is

$$p_{\iota rt} = c_{\iota rt} + \left(\Delta_{rt}^{-1}\vec{s}_{rt}\right)_\iota - \kappa_{\iota rt}d_{\iota rt} \tag{19}$$

where Δ_{rt} is the matrix of demand derivatives, with (ι,υ) element given by

$$(\Delta_{rt})_{\iota\upsilon} = \begin{cases} -\frac{\partial s_{\upsilon rt}}{\partial p_{\iota rt}} & \text{, if } \iota \text{ and } \upsilon \text{ belong to the same college,} \\ 0 & \text{otherwise.} \end{cases} \tag{20}$$

Equation (19) features the standard forces that drive prices in an imperfect competition environment. First, degrees with higher marginal costs charge higher prices. Second, the markup term, $\left(\Delta_{rt}^{-1}\vec{s}_{rt}\right)_\iota$, represents the inverse of the demand sensitivity matrix and captures how changes in prices influence enrollments across all programs offered by the college. If students who receive discounts are more price elastic—perhaps because they tend to have lower income—then discounted markups will generally be smaller than full-price markups. The final term, $\kappa_{\iota rt}d_{\iota rt}$, deviates from the standard framework and captures factors not explicitly modeled that shift the relative pricing of discounted versus full-price enrollment.

The pricing equation also clarifies how price discrimination alters the composition effect. Suppose there are two types of students: high and low price sensitivity, and that only high-sensitivity students receive loans. Under uniform pricing, loan holders are therefore more price sensitive than the average student, generating a composition effect that pushes prices down. Now suppose colleges price discriminate and all high-sensitivity students receive discounts while low-sensitivity students pay full price. In

[31] Treating enrollment with and without discounts as distinct products offered by the same college ensures that our framework remains formally equivalent to a standard Bertrand–Nash model with differentiated products. As a result, our pricing game raises no additional concerns regarding equilibrium existence or uniqueness beyond those already present in conventional BLP-style models. The computational approach likewise follows the standard methods used in that literature.

this case, the composition effect disappears: no full-price students receive loans, and among discounted students loan holders are no more price sensitive than other students.

Even though our supply model is flexible enough to capture several key features of this market, it abstracts from some potentially relevant margins. First, colleges take demand as given. This assumption simplifies the analysis but rules out the possibility that colleges may influence demand through investments in educational quality, infrastructure, or marketing.

Second, the framework incorporates price discrimination in a tractable way by relying on three simplifying assumptions: (1) within each degree, prices can be summarized by a two-price structure; (2) colleges choose the size of the available discount but not which students receive it; and (3) the components governing discount magnitude and availability—κ and ρ^D—are policy-invariant. Supplemental Appendix E provides empirical support for these three assumptions.

6 Model estimation

We estimate our model using the generalized method of moments (GMM) approach developed by Berry et al. (1995) (BLP hereafter) and extended by Petrin (2002), which combines instrumental variables and micro-moments to identify the model's parameters. We first outline the moments and instruments used in the estimation, followed by a detailed explanation of the construction of the estimator. Finally, we present the parameter estimates and evaluate the model's fit.

6.1 Moments

Price instrument. A critical challenge in estimating demand is the potential correlation between prices and the unobserved demand shock, ξ_{jrt}, which requires the use of an instrument for prices. To address this issue, we leverage the 2015 federal loans reform, which introduced a maximum number of students with loans per degree program. Because this cap was independent of degrees' pre-reform loan usage, degrees with more students using loans prior to the reform were disproportionately affected. In Section 4.2, we use Equation (6) to define an exposure measure, Exp_j, based on each degree's pre-reform loan usage, and show that degrees more affected by the reform experienced larger tuition reductions. Motivated by these findings, we define the following instrument:

$$Z_{jrt} = \text{Exp}_j \cdot \text{post}_t,$$

where Exp_j is the exposure of degree j to the reform, as defined in Equation (6), and $\text{post}_t \equiv \mathbb{1}\{t > 2015\}$ is an indicator for post-reform years.

We use this instrument to estimate price elasticities, assuming that Z_{jrt} is uncorrelated with unobserved demand and supply shocks, conditional on region-year and degree fixed effects. Formally, we impose the moment conditions $\mathbb{E}[Z_{jrt} \cdot \xi_{jrt}] = 0$ and $\mathbb{E}[Z_{jrt} \cdot \omega_{jrt}] = 0$. These conditions imply that the correlation between price changes and policy exposure observed in the data reflects equilibrium responses, rather than a spurious correlation between unobserved shocks and the exposure instrument Z_{jrt}.[32] Section 4 provides two pieces of supporting evidence. First, degrees with higher and lower exposure followed parallel trends in a range of outcomes prior to the reform. Second, the reduced-form relationship between exposure and post-reform outcomes remains robust when controlling for major-region-year fixed effects.

Exogenous variation in loan status. To estimate the effects of loans on demand, we exploit the discontinuity in loan access at each eligibility score threshold. As described in Section 3.1, students apply for loans through a centralized admission system that generates degree-specific eligibility cutoffs (see Supplemental Appendix C for a detailed description of the allocation system).[33]

As discussed in Section 5.1, there are two channels through which access to loans may influence enrollment decisions: an increase in baseline utility, α_L^0, and a reduction in price sensitivity, α_L^1. While both effects can raise the likelihood of enrollment, a reduction in price sensitivity may also induce students to select more expensive programs. To separately identify these mechanisms, we construct two moment conditions: the first is the log difference in enrollment rates for students just above and just below the eligibility threshold; the second is the log difference in total tuition expenditures. We construct these moments by pooling all degree-specific discontinuities.

To account for heterogeneity in loan effects by income, captured by α_{wL}^0 and α_{wL}^1, we estimate both moments separately for students in four household income bins. This yields eight empirical moments, which we include as micro-moments in our GMM estimation. These moments discipline the model to reproduce the observed patterns in enrollment and spending across income groups. Supplemental Appendix H.1.2 formally defines these moments, and Supplemental Appendix Figure H.1 provides a visualization of the underlying variation in the data.

These moments identify the effects of loans on demand because, in the model, the only factor affecting enrollment that changes discontinuously at the eligibility cutoffs is access to loans. This identification strategy relies on two key assumptions. The first is that there are no systematic differences between students just above and just below the eligibility cutoffs. In our setting, scores are determined by a centralized national exam, and program-specific cutoffs vary widely and are not known ex ante. Therefore, there is no scope for manipulation that could generate discontinuous differences at the thresholds. Supplemental Appendix H.1.3 provides empirical support for this assumption: The score distribution is smooth, and a series of predetermined student characteristics are balanced across the thresholds. The second assumption is the exclusion restriction, which states that crossing the loan eligibility cutoff

[32] This is consistent with Kargar and Mann (2023), who show that the price effects of federal student loan reforms in the U.S. are driven by changes in markups rather than costs.

[33] PROUNI, a federal grant program for low-income students at private colleges, also allocates scholarships using ENEM scores, and some degrees feature PROUNI-specific cutoffs. This could in principle confound the loan discontinuity. In practice, however, PROUNI cutoffs are typically higher than the corresponding FIES cutoffs, and the two coincide in fewer than 3% of degrees in our sample. In addition, FIES and PROUNI operate through separate application systems.

affects enrollment solely through access to the loan itself. No other government policy is linked to these thresholds, and regulations prohibit colleges from conditioning admission or tuition discounts on loan status. These institutional features support the validity of the exclusion restriction.

Other micro-moments. Our model incorporates heterogeneity in students' baseline utility based on exam scores, β^h, and income, β^w, as well as in price sensitivity by income, α_w^1. To estimate these parameters, we include the following three micro-moments: the average score of enrolled students, the average income of enrolled students, and the average income of students enrolled in degrees with a price above the median. To estimate the effect of income on the likelihood of receiving a loan, ρ_w^L, or a discount, ρ_w^D, we incorporate two additional micro-moments: the average income of students who receive loans and the average income of those who receive discounts. Equations that define these micro-moments are provided in Supplemental Appendix H.2.

6.2 Estimation

We estimate the model using the sample described in Section 3.2, which encompasses the in-person private higher education market across the 69 largest regions in Brazil. The sample comprises 13,567 unique degree programs offered by 695 colleges. We exclude 2012 and 2015 for the reasons discussed in Section 4. The year 2012 is omitted to avoid a mechanical correlation between structural shocks and exposure to the FIES reform, as exposure is measured in 2012. The year 2015 is excluded because the FIES reform was announced after tuition prices for that year had already been set. Consequently, 2015 loan allocations followed post-reform rules, while tuition prices were determined under the expectation of pre-reform rules. This inconsistency complicates the interpretation of college and student behavior during that year.

Let θ denote the vector of all model parameters, encompassing the preference parameters defined in the utility function in Equations (10), (11), and (13); financial aid usage parameters from Equations (15) and (16); and supply parameters in Equation (17). To estimate these parameters, we collect the moments described in Section 6.1 in a vector $m(\theta)$, which we then use to define the GMM objective function:

$$Q(\theta) = \frac{1}{2} m'(\theta) W m(\theta), \tag{21}$$

where W is the optimal GMM weighting matrix.

Our estimation approach builds on the methodology of BLP, while extending it to incorporate additional high-dimensional parameters specific to our model. In standard BLP estimation, θ includes two sets of high-dimensional parameters: mean utilities, δ_{jrt}, and marginal costs, c_{jrt}. These parameters are recovered as follows. First, δ_{jrt} is estimated by a fixed-point contraction that ensures that the market shares predicted by the model match those observed in the data. Second, c_{jrt} is recovered by inverting the firms' first-order conditions (FOCs). From δ_{jrt} and c_{jrt} one can recover the model residuals, ξ_{jrt}

and ω_{jrt}, and use them to construct moment conditions $m(\theta)$. Thus, the algorithm minimizes $Q(\theta)$ subject to the constraint whereby equilibrium prices and market shares in the model must align exactly with their counterparts in the data.

In contrast, our model requires estimating five sets of high-dimensional parameters: mean utilities, δ_{jrt}; loan usage propensities, ρ^L_{ft}; discount usage propensities, ρ^D_{ft}; marginal costs, c_{jrt}; and the structural errors that govern discount magnitudes, κ_{jrt}. To estimate these parameters, we extend the BLP approach. The first three sets of parameters are recovered using fixed-point contractions that ensure that specific quantities in the model match their observed counterparts in the data. Specifically, matching degree-year market shares yields δ_{jrt}, while matching the share of students with loans and discounts in each college provides ρ^L_{ft} and ρ^D_{ft}, respectively. We recover the remaining two sets of parameters, c_{jrt} and κ_{jrt}, by inverting the firms' FOCs, leveraging the fact that firms optimize two choice variables (p^F_{jrt} and p^D_{jrt}). Similar to the standard BLP, our estimator minimizes $Q(\theta)$ subject to the constraint whereby these five model-predicted quantities must match their empirical counterparts. In Supplemental Appendix H.3, we formally define the estimator.

6.3 Model estimates

We report the estimated preference and targeting parameters in Supplemental Appendix Table A.2 and the distributions of price elasticities, marginal costs, and markups in Supplemental Appendix Figure A.6. The median price elasticity is -3.46, consistent with prior research on the private higher education sector (Armona and Cao, 2024; Barahona et al., 2025). Our estimates confirm that price sensitivity varies with both student income and loan status. The median price elasticity for students with below-average income is -5.07, whereas for those with above-average income it is -2.96. Also, taking a loan reduces individual price sensitivity, α^1_{ijrt}, by 11.1% on average.

The median yearly tuition price and marginal cost are \$4,152 and \$2,945, respectively. The implied median markup is 0.29, which indicates that colleges exert substantial market power. The median structural error that governs unexplained differences between full and discounted prices, κ_{jrt}, is \$852, which implies that colleges behave as if educating discounted students entailed a lower marginal cost. Such behavior could reflect true differences in the cost of serving different types of students or, alternatively, colleges may derive additional utility from enrolling these students—perhaps due to reputational gains or a desire to signal a commitment to social inclusion.[34]

Our estimates of the financial-aid targeting parameters, reported in Panel B of Supplemental Appendix Table A.2, indicate that loans are effectively targeted toward low-income students, whereas tuition discounts are not. A one-standard-deviation increase in family income reduces the probability of receiving a loan by 12.9 percentage points. By contrast, the relationship between income and the probability of receiving a tuition discount is small and not statistically significant. These results are

[34]Markup is defined as $\frac{p_{jrt}-c_{jrt}}{p_{jrt}}$. Tuition prices, marginal costs, and κ are reported in USD per year.

consistent with descriptive patterns in the data: loan receipt is substantially more correlated with family income than tuition discounts, as shown in Supplemental Appendix Figure E.1.

We assess the model's fit by comparing its predictions to the corresponding empirical moments. Supplemental Appendix Table A.3 reports the targeted micro-moments alongside their empirical counterparts. The close alignment between model-implied and observed moments indicates that the model captures key features of the Brazilian higher-education market.

Interpreting the estimates: Direct and composition effects. We now interpret the model estimates in the context of the relative strengths of the direct and composition effects. As shown in Equation (5), the direct effect depends on how much loans reduce recipients' price elasticity, while the composition effect is stronger when students likely to receive loans are substantially more price-elastic than the market average. While these two objects are related, they capture distinct forces: The direct effect reflects the impact of loans on individual price elasticities, whereas the composition effect depends on how loans are targeted across students with heterogeneous baseline elasticities.

To quantify these forces, we proceed as follows. Using Equation (15), we first compute each student's probability of receiving a loan and organize students into groups based on these probabilities.[35] We then estimate the price elasticity of demand for each group under two scenarios: one in which all students receive loans and another in which none do, holding tuition fixed at observed levels. Comparing these scenarios isolates the direct effect of loan receipt on demand, holding fixed the composition of students within each group. In contrast, comparing elasticities across groups within the "no loan" scenario reveals heterogeneity in price sensitivity, which is the source of the composition effect.

The results, presented in Figure 4, reveal two key patterns. First, loans reduce price elasticity. Second, students with a higher probability of receiving a loan exhibit greater price elasticity. Notably, the difference in elasticity between high- and low-loan-probability groups is substantially larger than the within-group effect of loan receipt, suggesting that the composition effect plays a quantitatively important role.

We next assess the role of price discrimination as an additional channel that could attenuate the composition effect. As discussed in Section 5.3, if tuition discounts were targeted similarly to loans, students with and without loans would face different net prices, potentially weakening the composition effect. In practice, however, loan availability is strongly correlated with income, while tuition discounts are less so. As a result, loan and discount usage are only weakly correlated: -0.09 in the raw data, and close to zero (-0.0295) using structural estimates that account for endogenous selection into institutions that offer more discounts. This weak correlation suggests that price discrimination does little to attenuate the composition effect in our setting.

[35] Because the probability of receiving a loan depends on a student's chosen degree, we compute, for each student, the average loan probability across all degree options available in their market.

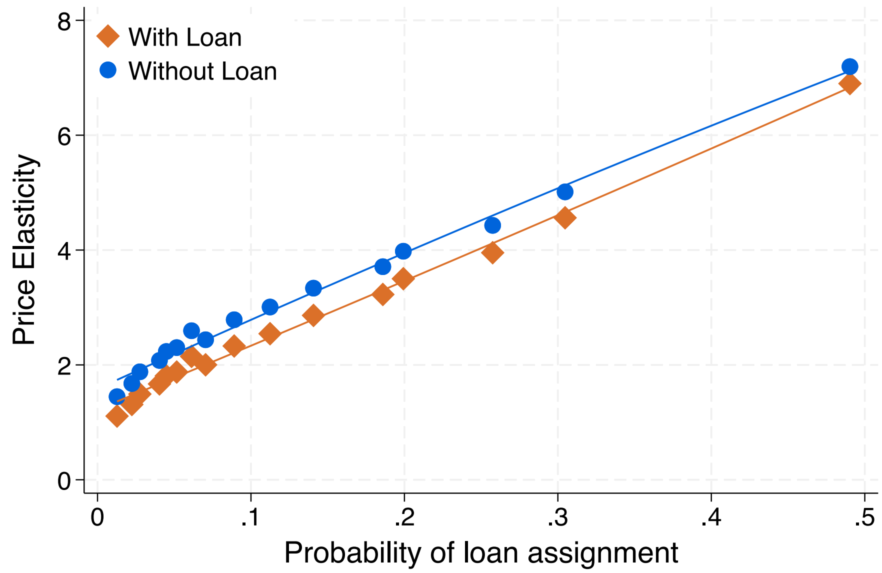

Figure 4: Price elasticity and loan targeting

Notes: This figure illustrates the relationship between price sensitivity and loan targeting, as well as the effect of loans on price sensitivity. The x-axis reports the probability of receiving a loan, computed using Equation (15) and the estimated parameters from Section 6.3. Students are grouped based on this probability, as represented by the dots and squares. The y-axis shows the estimated price elasticity of demand for each group under two scenarios: one in which all students receive loans (orange squares) and another in which no student does (blue dots). Prices are held fixed at their observed levels in the data. The lines represent quadratic fits.

7 Equilibrium effects of targeted student loans

In this section, we use our estimated model to quantify how loan availability affects tuition fees and college enrollment in equilibrium, and examine how different aid-targeting strategies shape these outcomes. The analysis proceeds in four steps. First, we define the sample and outline the simulation procedure. Second, we introduce the aid-targeting schemes. Third, we evaluate how targeting influences tuition prices. Fourth, we assess college enrollment effects, distinguishing between the direct demand response to aid and the equilibrium adjustments driven by tuition changes.

7.1 Sample and simulation procedure

Throughout this exercise, we simulate counterfactual scenarios using 2014 data, which capture the pre-reform market environment. Accordingly, our counterfactual sample includes all 2013 ENEM takers who participated in the 2014 enrollment cycle, subject to the sample restrictions described in Section 3.3. Since we focus on a single year, we omit the t indices throughout this section. Additionally, since each degree j belongs to a unique region r, region indices are also suppressed throughout.

To simulate market shares and tuition fees, we use estimated parameters from Equations (9) and (17), while holding demand and supply shocks (ξ_j, ω_j, and κ_j) constant in order to replicate the 2014 market conditions. Discount allocations are simulated based on the estimated parameters from Equation

(16), with firm-specific discount allocation parameters, ρ_f^D, also fixed at their 2014 levels.[36]

While our model allows discount sizes, $p_j^F - p_j^D$, to adjust endogenously in response to different loan policies, the mechanism used to allocate discounts remains fixed. Specifically, the discount allocation parameters, ρ_w^D and ρ_f^D, do not vary across counterfactual scenarios. As a result, while students' enrollment decisions—and thus the overall number of discounts used—may change, the eligibility criteria for discounts, D_{ij}, remain constant. In other words, institutions cannot modify how discounts are distributed across students, but shifts in enrollment choices may still affect the total number of discounts used. Supplemental Appendix E.3 provides evidence that institutions do not adjust their discount targeting in response to changes in loan policy.

7.2 Loan targeting schemes

We consider five scenarios for allocating loans. In the first, no loans are distributed. In the remaining four, we vary the targeting scheme used to assign loans:

(0) *No-loans*: This scenario represents a world without student loans.

(1) *Baseline:* The scheme is based on the loan allocation observed in the data. Specifically, we use Equation (15) and assign loans based on the estimated parameters $\hat{\rho}_w^L$ and $\hat{\rho}_f^L$.[37]

(2) *Random:* Loans are randomly assigned to students, with each student having an equal probability $\rho \in [0, 1]$ of qualifying for a loan across all degree programs. A higher value of ρ increases the share of students with loans.

(3) *Need-based:* Loans are allocated based on students' income levels w_i. Coverage is determined by an income percentile threshold $\rho \in [0, 1]$: Students with $pctile(w_i) \leq \rho$ qualify for a loan across all degree programs, where $pctile(\cdot)$ is the percentile function. Lower values of ρ restrict aid to the most financially disadvantaged, while higher values expand coverage.

(4) *Merit-based:* Loans are allocated based on academic performance h_i. Coverage depends on a percentile threshold $\rho \in [0, 1]$: Students with $1 - pctile(h_i) \leq \rho$ qualify for a loan across all degree programs. Smaller values of ρ limit loans to top-performing students, while higher values expand coverage.

In loan allocation schemes (2)–(4), the parameter ρ captures the generosity of the loan program. In these schemes, ρ corresponds to the *loan coverage rate*, defined as the share of potential students who

[36] We simulate discount allocations by drawing the student-degree-specific discount propensity, ϑ_{ij}^D, from a logistic distribution. Draws are independent across student-degree pairs and held fixed across all counterfactuals.

[37] Since this allocation corresponds to the pre-reform period, there is no degree-specific loan cutoff, i.e., $\overline{h}_j = 0$. We simulate loan allocations by drawing the student-degree-specific loan propensity, ϑ_{ij}^L, from a logistic distribution. The draws are independent for each student-degree pair. By construction, the baseline simulation matches the observed firm-level share of students receiving loans.

would receive a loan if they enroll. This follows by construction because the model does not distinguish eligibility from take-up: all students assigned a loan take it up conditional on enrollment.

For each scheme k and loan coverage rate ρ, we define the total cost of the loan program as a function of market shares and tuition fees, normalized relative to the cost of the baseline program.[38] Specifically, the budget function $B^k(\rho)$ is given by

$$B^k(\rho) = \frac{\sum_{j \in J} \left[N_j^k(1,0)(\rho) \cdot p_j^{Fk}(\rho) + N_j^k(1,1)(\rho) \cdot p_j^{Dk}(\rho) \right]}{\overline{B}}, \tag{22}$$

where $N_j^k(l,d)(\rho)$ denotes the number of students in degree j with loan status l and discount status d under scheme $k \in \{2,3,4\}$, given a loan coverage rate ρ. The functions $p_j^{Fk}(\rho)$ and $p_j^{Dk}(\rho)$ represent the full and discounted tuition prices, respectively. The denominator $\overline{B}$ denotes the total program cost in the baseline scenario.[39] To facilitate comparison across schemes, we define a scenario as *budget-neutral* when $B^k(\rho) = 1$, which indicates that total spending is held fixed at the baseline level. Note that our definition of budget neutrality accounts for endogenous tuition changes.

7.3 Impact of loans on tuition fees

To evaluate the overall impact of the loan program on tuition prices, we simulate equilibrium outcomes under the scenarios described above. Tuition fees—both full prices and discounts—are determined by Equation (19), while market shares are governed by Equation (14). Our analysis focuses on average tuition, calculated as the enrollment-weighted mean of full and discounted prices across degree programs.

Figure 5 shows simulated tuition changes, expressed as percentage differences relative to the no-loan benchmark. The x-axis plots the budget of the loan program under each counterfactual, normalized by the baseline budget, as defined in Equation (22). The vertical dashed line indicates budget-neutral loan programs—i.e., with the same total budget as the baseline.

We begin by examining the price effects under the *baseline* loan allocation, indicated by the black triangle in Figure 5. In this counterfactual, loans are assigned according to their empirical distribution in the data, in which loans are imperfectly targeted toward low-income students (see Supplemental Appendix Figure E.1). The resulting net price effect reflects the interaction of two opposing forces: an upward direct effect and a downward composition effect, since loan recipients tend to be more price elastic (see Section 6.3 and Figure 4). Overall, relative to the no-loan scenario, this allocation increases equilibrium tuition by 0.1%. This indicates that, under the baseline loan allocation, the direct and composition effects approximately offset each other.[40]

[38] This cost reflects the government's expenditure on tuition for loan recipients while they are enrolled. In all schemes, we assume the government covers 100% of tuition costs and do not account for future student repayments.

[39] Formally, the baseline cost is $\overline{B} = \sum_{j \in J} \left[N_j^1(1,0) \cdot p_j^{F1} + N_j^1(1,1) \cdot p_j^{D1} \right]$, where $N_j^1(l,d)$ denotes the number of students in degree j with loan status l and discount status d in the baseline scenario, and p_j^{F1} and p_j^{D1} are the corresponding full and discounted tuition prices.

[40] The effects reported here are not directly comparable to those in Table 1. In this section, we compare the 2014 loan

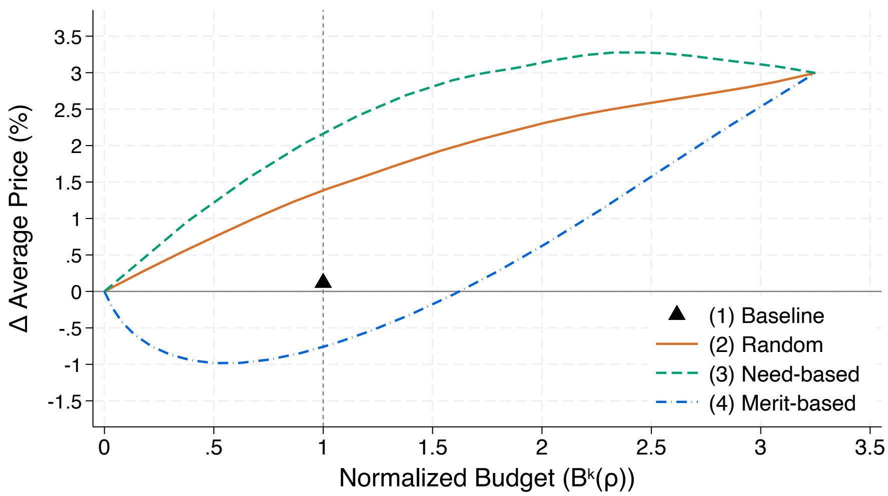

Figure 5: Impacts of loans on prices

Notes: This figure shows the effect of student loans on tuition under alternative targeting schemes. The y-axis reports the percentage change in average tuition relative to a scenario with no loans. Average tuition is computed as the enrollment-weighted average of full and discounted prices across degrees. Results are based on the baseline estimated parameters (black triangle) and counterfactual simulations under three targeting schemes: random allocation (orange solid line), need-based targeting (blue dashed line), and merit-based targeting (green dash-dot line). The vertical dashed line indicates budget-neutral loan programs (i.e., with the same total budget as the baseline). The x-axis reports the budget of the loan program under each counterfactual, normalized by the baseline budget, as defined in Equation (22).

Next, we examine the effects of *random* targeting, in which loans are allocated independently of student characteristics. In this case, tuition responds solely to the direct effect of increased financial aid, since the composition effect—present when aid is targeted—is absent. Under a budget-neutral constraint, tuition rises by 1.4% relative to the no-loan scenario. Since only the direct effect operates, tuition increases monotonically with the size of the loan program, as illustrated by the solid orange line.

We then consider two widely used financial aid allocation strategies: *need-based* and *merit-based* loan targeting. Unlike random targeting, these approaches generate both direct and composition effects, which affect tuition through distinct channels. Under *need-based* targeting, loans are directed toward low-income students, generating a negative composition effect that puts downward pressure on tuition. As shown by the blue line in Figure 5, this effect dominates when the program budget is up to 1.6 times the baseline level, resulting in equilibrium tuition below the no-loan benchmark. As the program expands and the income of the marginal loan recipient rises, the composition effect weakens. Beyond this threshold, the direct effect dominates, leading tuition to increase.

By contrast, merit-based targeting allocates loans to students with the highest test scores—a group that, as shown in Supplemental Appendix Figure A.7, is disproportionately drawn from higher-income backgrounds. These students tend to be less sensitive to tuition prices, which means the composition

allocation to a counterfactual without loans. By contrast, Table 1 exploits the 2015 reform and compares equilibrium prices before and after the policy change, which both reduced the overall volume of loans and modified their targeting across students.

effect works in the same direction as the direct effect, reinforcing upward pressure on tuition. This pattern is reflected in the green line in Figure 5, which shows that tuition increases at a higher rate under merit-based targeting than in the random scenario as the loan program expands.

Overall, we find substantial heterogeneity in tuition responses across targeting schemes, reflecting the interaction of direct and composition effects. Under *budget-neutral* allocations—indicated by the vertical dashed line in Figure 5—*merit-based* targeting increases tuition by 2.2%, whereas *need-based* targeting leads to a 0.8% decrease. The large gap in price effects across these scenarios highlights that tuition responses can either amplify or dampen the overall impact of loan programs on college enrollment, depending on how loans are targeted. We examine these enrollment effects in the next section.

7.4 Equilibrium impacts of loans on enrollment

We now examine how loan targeting affects college enrollment, focusing on two outcomes: total enrollment and enrollment in high-quality degree programs. We measure program quality using the average earnings of a degree's graduates, computed from RAIS, a matched employer–employee administrative dataset covering the universe of formal-sector employment in Brazil. We classify "top" programs as those in the top 10% of the quality distribution.[41] All effects are expressed as percentage changes relative to enrollment under the no-loan scenario.

To disentangle demand- and supply-side effects, we construct two counterfactual enrollment measures for each targeting scheme. The first, *demand-only*, captures the effect of loan availability on students' enrollment decisions while holding tuition fixed in the no-loan scenario. The second, *equilibrium*, incorporates supply-side responses by allowing tuition to adjust endogenously to changes in demand, thereby capturing the full impacts of the loan program. For clarity, we focus on the case of a budget-neutral loan program, with results summarized in Figure 6.

In the demand-only counterfactual, need-based targeting generates the largest increase in total enrollment (35.8%), followed by the baseline scenario (31.9%), random allocation (28.6%), and merit-based targeting (24.2%). These differences are driven by two factors. First, the number of students who receive a loan but would have enrolled even without one. Under a budget-neutral constraint, allocating more loans to such inframarginal students reduces the number of loans available for students who enroll only if they receive financial aid. Second, targeting schemes that concentrate loans on more expensive degrees exhaust the budget more quickly, which results in fewer loans being offered overall. Indeed, students with above-median test scores are 1.8 times more likely to enroll without a loan than those with below-median scores, and enroll in degrees that are 16% more expensive, on average. As a consequence, the merit-based scheme—which targets high-score students—generates a weaker impact

[41] Supplemental Appendix Figure A.8 reports robustness results using the average entrance exam score of incoming students as an alternative measure of quality. All quality measures are constructed using the 2010 incoming cohort and are held fixed throughout the analysis.

on enrollment.

While overall enrollment responses are larger under need-based targeting, enrollment in top degrees increases more under merit-based targeting when considering demand-only counterfactuals. This reversal is driven by sorting: even when liquidity constraints are relaxed, low-income students are more likely to enroll in lower-tier programs, which attenuates the effect of *need-based* targeting on enrollment in top programs.[42]

In the *equilibrium* counterfactual, endogenous price adjustments either amplify or dampen enrollment effects, depending on the targeting scheme. Under need-based targeting, the associated price decline increases the enrollment gain from 35.8% to 39.6%. By contrast, under merit-based targeting, price increases offset much of the enrollment gains observed in the demand-only counterfactual, reducing the effect from 24.2% to just 15.8%. These patterns underscore how endogenous price responses can critically shape the aggregate impact of loan programs.

The effects of supply-side responses on top-degree enrollment mirror those for overall enrollment. Under need-based targeting, endogenous tuition adjustments amplify the impact of loans on top-degree enrollment, increasing it from 26.8% in partial equilibrium to 31.9% in general equilibrium. By contrast, under merit-based targeting, the effect declines substantially, from 47.9% to 34.1%. As a result, tuition responses largely eliminate the sizable differences in top-degree enrollment effects observed in partial equilibrium.[43]

(a) Total enrollment (b) Enrollment in top degree programs

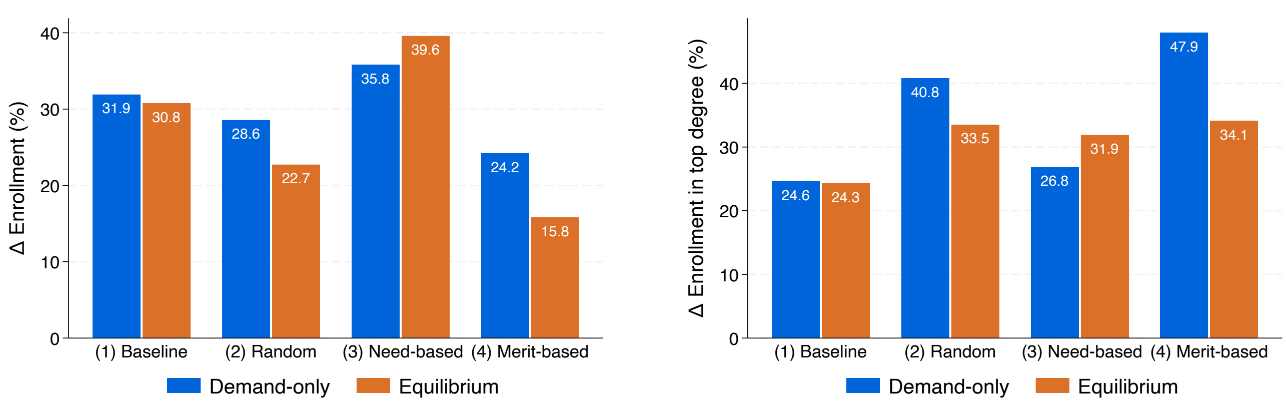

Figure 6: Equilibrium effects of loans on enrollment

Notes: This figure shows the effects of student loans on enrollment under alternative targeting schemes. Outcomes are reported as percentage changes relative to a scenario with no loans. For each targeting scheme, simulations are conducted at a coverage rate ρ that ensures a constant total loan program budget equal to the baseline level (budget-neutral). Panel A presents effects on total enrollment, while Panel B focuses on enrollment in programs in the top 10% of the quality distribution. Degree quality is measured by the average earnings of its graduates. Orange bars allow for supply-side responses by letting tuition adjust endogenously; blue bars hold prices fixed at their observed levels.

[42]This sorting reflects factors such as prices, geographic proximity, or personal preferences, rather than admissions constraints. In the data, top programs operate at only 57% of reported capacity, versus 47% for non-top programs—both far below full utilization.

[43]Supplemental Appendix Figure A.9 reports the effects of each loan scheme on tuition in top-degree programs.

Taken together, our findings highlight the central role of targeting in determining the effectiveness of student loan programs and have important policy implications. First, need-based programs that prioritize low-income students not only expand access to higher education for recipients but can also exert downward pressure on tuition, thus potentially improving affordability for all students—regardless of loan status. This dual effect underscores the strength of low-income targeting as a policy tool. Second, the results suggest that merit-based programs, which allocate aid based on academic achievement, may contribute to rising tuition and potentially undermine affordability and restrict broader access.

It is important to note that our analysis isolates the price effects of targeting and abstracts from other dimensions of financial aid design. For instance, students with higher academic achievement may have lower dropout rates, generate positive peer effects, and exhibit higher loan repayment rates—factors that could enhance the financial sustainability and broader benefits of aid programs. While our findings highlight the sensitivity of tuition responses to the structure of aid targeting, these price effects should be considered alongside other factors when evaluating the overall desirability of alternative financial aid designs.

8 Conclusion

In this paper, we make three key contributions. First, we show that the impact of targeted financial aid can be decomposed into a *direct* effect, which raises tuition, and a *composition* effect, which may increase or decrease tuition depending on the targeting scheme. Second, we provide empirical validation by leveraging a major reform in Brazil's federal student loan program. Third, we develop and estimate an equilibrium model of higher education. Using the estimated model, we show that a need-based loan program would *reduce* tuition by 0.8%, on average, relative to a no-loan scenario. In contrast, a merit-based scheme—i.e., allocating loans to students with high test scores, which are strongly correlated with high income—would *raise* tuition by 2.2%. Moreover, enrollment gains relative to the no-loan scenario range from 39.6% under need-based targeting to just 15.8% under merit-based targeting.

The insight that the incidence of subsidies includes a composition effect has broader relevance beyond higher education and offers a valuable lens for understanding price responses in other markets with imperfect competition. For instance, Polyakova and Ryan (2022) show that the effects of the Affordable Care Act on health insurance prices hinge on how subsidies are targeted. Similarly, in housing markets, targeted government subsidies often interact with privately set rents. Numerous studies have examined the impact of housing subsidies on rent prices, yielding heterogeneous results across settings (Gibbons and Manning, 2006; Fack, 2006; Eriksen and Ross, 2015; Brewer et al., 2019; Hyslop and Rea, 2019; Ramírez Sierra et al., 2024; Ochoa, 2025). Investigating the role of the composition effect in these contexts could help explain this variation and represents an important direction for future research.

References

Aguirre, Josefa, "Long-Term Effects of Grants and Loans for Vocational Education," *Journal of Public Economics*, December 2021, *204*, 104539.

Allende, Claudia, "Competition Under Social Interactions and the Design of Education Policies," *Working Paper*, 2021.

Anderson, Simon P, Andre de Palma, and Brent Kreider, "Tax Incidence in Differentiated Product Oligopoly," *Journal of Public Economics*, 2001, *81* (2), 173–192.

Armona, Luis and Shengmao Cao, "Redesigning Federal Student Aid in Sub-baccalaureate Higher Education," *Working Paper*, 2024.

Avery, Christopher and Caroline Hoxby, *College Choices: The Economics of Where to Go, When to Go, and How to Pay for It* National Bureau of Economic Research Conference Report, Chicago: University of Chicago Press, 2004.

Baird, Matthew, Michael S. Kofoed, Trey Miller, and Jennie Wenger, "Veteran Educators or For-Profiteers? Tuition Responses to Changes in the Post-9/11 GI Bill," *Journal of Policy Analysis and Management*, September 2022, *41* (4), 1012–1039.

Barahona, Nano, Caue Dobbin, Joaquin Fuenzalida, and Sebastian Otero, "The Effects of Widespread Online Education on Market Structure and Enrollment," 2025.

Bennett, W. J., "Our Greedy Colleges," *The New York Times*, February 1987, pp. 31–32.

Berry, Steven, James Levinsohn, and Ariel Pakes, "Automobile Prices in Market Equilibrium," *Econometrica*, July 1995, *63*, 44–65.

Black, Sandra, Lesley J. Turner, and Jeffrey Denning, "PLUS or Minus? The Effect of Graduate School Loans on Access, Attainment, and Prices," 2023.

Bo, Inácio and Rustamdjan Hakimov, "Iterative Versus Standard Deferred Acceptance: Experimental Evidence," *The Economic Journal*, July 2019, *130* (626), 356–392.

Bodere, Pierre, "Dynamic Spatial Competition in Early Education: An Equilibrium Analysis of the Preschool Market in Pennsylvania," *Working Paper*, 2023.

Brewer, Mike, James Browne, Carl Emmerson, Andrew Hood, and Robert Joyce, "The Curious Incidence of Rent Subsidies: Evidence of Heterogeneity from Administrative Data," *Journal of Urban Economics*, November 2019, *114*, 103198.

Bucarey, Alonso, "Who Pays for Free College? Crowding Out on Campus," *Working Paper*, 2018.

_ , Dante Contreras, and Pablo Muñoz, "Labor Market Returns to Student Loans for University: Evidence from Chile," *Journal of Labor Economics*, October 2020, *38* (4), 959–1007.

Calonico, Sebastian, Matias D. Cattaneo, and Rocio Titiunik, "Robust Nonparametric Confidence Intervals for Regression-Discontinuity Designs: Robust Nonparametric Confidence Intervals," *Econometrica*, November 2014, *82* (6), 2295–2326.

Castleman, Benjamin L. and Lindsay C. Page, "Summer Nudging: Can Personalized Text Messages and Peer Mentor Outreach Increase College Going among Low-Income High School Graduates?," *Journal of Economic Behavior & Organization*, July 2015, *115*, 144–160.

Cattaneo, Matias D., Michael Jansson, and Xinwei Ma, "Manipulation Testing Based on Density Discontinuity," *The Stata Journal: Promoting communications on statistics and Stata*, March 2018, *18* (1), 234–261.

Cellini, Stephanie Riegg and Claudia Goldin, "Does Federal Student Aid Raise Tuition? New Evidence on For-Profit Colleges," *American Economic Journal: Economic Policy*, November 2014, *6* (4), 174–206.

Chaisemartin, Clément and Luc Behaghel, "Estimating the Effect of Treatments Allocated by Randomized Waiting Lists," *Econometrica*, 2020, *88* (4), 1453–1477.

Cook, Emily, "Market Structure and College Access in the US," 2025.

Delipalla, Sofia and Michael Keen, "The Comparison between Ad Valorem and Specific Taxation under Imperfect Competition," *Journal of Public Economics*, December 1992, *49* (3), 351–367.

Diamond, P.A., "A Many-Person Ramsey Tax Rule," *Journal of Public Economics*, November 1975, *4* (4), 335–342.

Epple, Dennis and Richard E Romano, "Competition between Private and Public Schools, Vouchers, and Peer-Group Effects," *American Economic Review*, 1998, *88* (1), 33–62.

_ , Richard Romano, and Holger Sieg, "Admission, Tuition, and Financial Aid Policies in the Market for Higher Education," *Econometrica*, 2006, *74* (4), 885–928.

Eriksen, Michael D. and Amanda Ross, "Housing Vouchers and the Price of Rental Housing," *American Economic Journal: Economic Policy*, August 2015, *7* (3), 154–176.

Fack, Gabrielle, "Are Housing Benefit an Effective Way to Redistribute Income? Evidence from a Natural Experiment in France," *Labour Economics*, December 2006, *13* (6), 747–771.

Ferreyra, Maria Marta, Ciro Avitabile, Javier Botero Alvarez, Francisco Haimovich Paz, and Sergio Urzua, *At a Crossroads: Higher Education in Latin America and the Caribbean*, The World Bank, May 2017.

Fillmore, Ian, "Price Discrimination and Public Policy in the US College Market," *The Review of Economic Studies*, May 2023, *90* (3), 1228–1264.

Finkelstein, Amy and Matthew J Notowidigdo, "Take-Up and Targeting: Experimental Evidence from SNAP," *The Quarterly Journal of Economics*, August 2019, *134* (3), 1505–1556.

Fu, Chao, "Equilibrium Tuition, Applications, Admissions, and Enrollment in the College Market," *Journal of Political Economy*, 2014, *122* (2), 225–281.

Gibbons, Stephen and Alan Manning, "The Incidence of UK Housing Benefit: Evidence from the 1990s Reforms," *Journal of Public Economics*, May 2006, *90* (4-5), 799–822.

Hampole, Menaka V, "Financial Frictions and Human Capital Investments," 2024.

Hyslop, Dean R. and David Rea, "Do Housing Allowances Increase Rents? Evidence from a Discrete Policy Change," *Journal of Housing Economics*, December 2019, *46*, 101657.

Joensen, Juanna Schrøter and Elena Mattana, "Student Aid Design, Academic Achievement, and Labor Market Behavior: Grants or Loans?," 2021.

Kargar, Mahyar and William Mann, "The Incidence of Student Loan Subsidies: Evidence from the PLUS Program," *The Review of Financial Studies*, March 2023, *36* (4), 1621–1666.

Kroft, Kory, Jean-William Laliberté, René Leal-Vizcaíno, and Matthew J. Notowidigdo, "Efficiency and Incidence of Taxation with Free Entry and Love-of-Variety Preferences," *American Economic Journal: Economic Policy*, May 2024, *16* (2), 300–334.

_ , _ , _ , and Matthew J Notowidigdo, "Salience and Taxation with Imperfect Competition," *Review of Economic Studies*, January 2024, *91* (1), 403–437.

Londono-Velez, Juliana and Catherine Rodriguez, "Upstream and Downstream Impacts of College Merit-Based Financial Aid for Low-Income Students: Ser Pilo Paga in Colombia," *American Economic Journal: Economic Policy*, 2020, *12* (2), 193–227.

Long, Bridget Terry, "How Do Financial Aid Policies Affect Colleges? The Institutional Impact of the Georgia HOPE Scholarship," *The Journal of Human Resources*, 2004, *39* (4), 1045–1066.

Lucca, David O, Taylor Nadauld, and Karen Shen, "Credit Supply and the Rise in College Tuition: Evidence from the Expansion in Federal Student Aid Programs," *The Review of Financial Studies*, 2019, *32* (2), 423–466.

Mahoney, Neale and E Glen Weyl, "Imperfect Competition in Selection Markets," *The Review of Economics and Statistics*, 2017, *99* (4), 637–651.

McCrary, Justin, "Manipulation of the Running Variable in the Regression Discontinuity Design: A

Density Test," *Journal of Econometrics*, February 2008, *142* (2), 698–714.

Mello, João M. P. De and Isabela F. Duarte, "The Effect of the Availability of Student Credit on Tuition: Testing the Bennett Hypothesis Using Evidence from a Large-Scale Student Loan Program in Brazil," *Economía*, 2020, *20* (2), 179–222.

Miravete, Eugenio J., Katja Seim, and Jeff Thurk, "Market Power and the Laffer Curve," *Econometrica*, 2018, *86* (5), 1651–1687.

Mirrlees, J. A., "An Exploration in the Theory of Optimum Income Taxation," *The Review of Economic Studies*, 1971, *38* (2), 175–208.

Moffitt, Robert, "An Economic Model of Welfare Stigma," *The American Economic Review*, 1983, *73* (5), 1023–1035.

Neilson, Christopher A, "Targeted Vouchers, Competition Among Schools, and the Academic Achievement of Poor Students," *Working Paper*, 2021.

Nevo, Aviv, "Measuring Market Power in the Ready-to-Eat Cereal Industry," *Econometrica*, 2001, *69* (2), 307–342.

Ochoa, Fernando, "Targeting and Price Pass-Through in Housing Voucher Design," Working Paper 2025.

OECD, "Indicator B5: How Much Do Tertiary Students Pay and What Public Support Do They Receive?," *Education at a Glance 2014: OECD Indicators*, 2014.

Petrin, Amil, "Quantifying the Benefits of New Products: The Case of the Minivan," *Journal of Political Economy*, 2002, *110* (4), 705–729.

Piketty, Thomas and Emmanuel Saez, "Optimal Labor Income Taxation," in "Handbook of Public Economics," Vol. 5, Elsevier, 2013, pp. 391–474.

Polyakova, Maria and Stephen P Ryan, "Market Power and Redistribution: Evidence from the Affordable Care Act," *Working Paper*, 2022.

Ramsey, F. P., "A Contribution to the Theory of Taxation," *The Economic Journal*, March 1927, *37* (145), 47.

Rothschild, Michael and Joseph E. Stiglitz, "Equilibrium in Competitive Insurance Markets: An Essay on the Economics of Imperfect Information," *The Quarterly Journal of Economics*, 1976, *90* (4), 629–649.

Saez, Emmanuel, "Using Elasticities to Derive Optimal Income Tax Rates," *Review of Economic Studies*, 2001, *68* (1), 2005–229.

Sanchez, Cristian, "Equilibrium Consequences of Vouchers Under Simultaneous Extensive and Intensive Margins Competition," *Working Paper*, 2023.

Schwartzman, Simon, *Private Higher Education in Brazil*, SciELO, 2026.

Sierra, Gabriel Darío Ramírez, Alayn Alejandro González Martínez, Miguel Ángel Monroy Cruz, and Luis Gerardo Zapata Barrientos, "The Impact of Subsidies on House Prices in Mexico's Mortgage Market for Low-Income Households 2008–2019," *Journal of Housing Economics*, March 2024, *63*, 101970.

Singell, Larry D. and Joe A. Stone, "For Whom the Pell Tolls: The Response of University Tuition to Federal Grants-in-Aid," *Economics of Education Review*, June 2007, *26* (3), 285–295.

Solís, Alex, "Credit Access and College Enrollment," *Journal of Political Economy*, April 2017, *125* (2), 562–622.

Tebaldi, Pietro, "Estimating Equilibrium in Health Insurance Exchanges: Price Competition and Subsidy Design under the ACA," *Review of Economic Studies*, January 2025, *92* (1), 586–620.

Weyl, E. Glen and Michal Fabinger, "Pass-Through as an Economic Tool: Principles of Incidence under Imperfect Competition," *Journal of Political Economy*, June 2013, *121* (3), 528–583.

Supplemental Appendix for:

Equilibrium Price Responses to Targeted Student Financial Aid

Nano Barahona   Cauê Dobbin   Sebastián Otero

A Appendix figures and tables

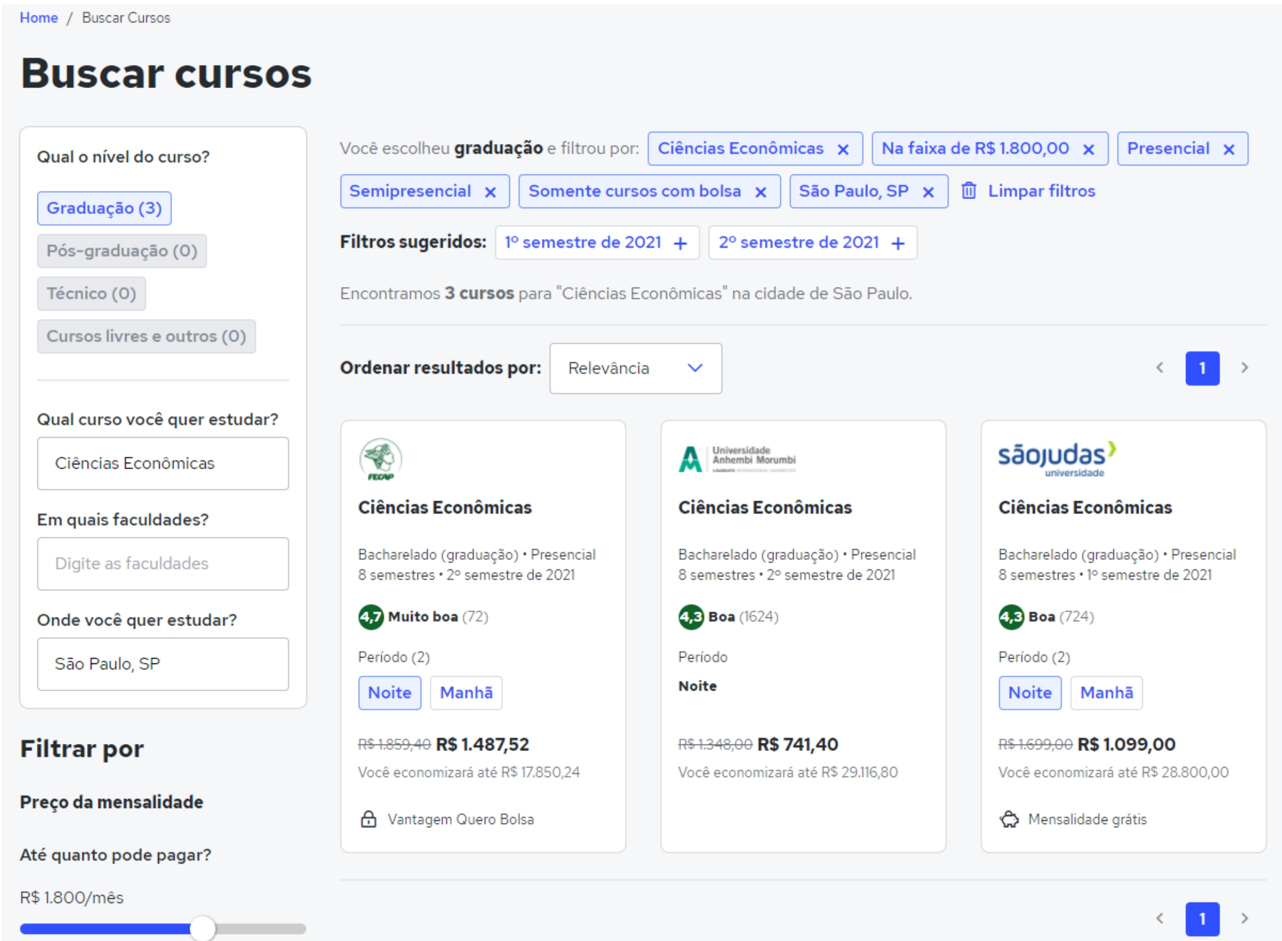

Figure A.1: Example of a Tuition Discount Marketplace Interface

Notes: This figure presents a screenshot from QueroBolsa, Brazil's largest online marketplace for tuition discounts. The platform's interface enables students to filter and select degree programs by location, field of study, and tuition price. Additional details are provided in Section 3.1.

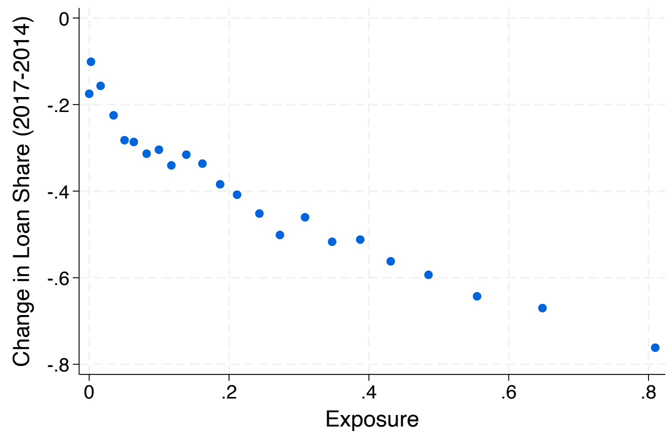

Figure A.2: Differential Exposure to the 2015 FIES Reform

Notes: This figure displays a binned scatter plot illustrating differential exposure to the 2015 FIES reform, based on degree-level data. The x-axis reports the exposure measure defined in Section 4.2, while the y-axis shows the change in the share of incoming students with loans between 2014 (pre-reform) and 2017 (post-reform).

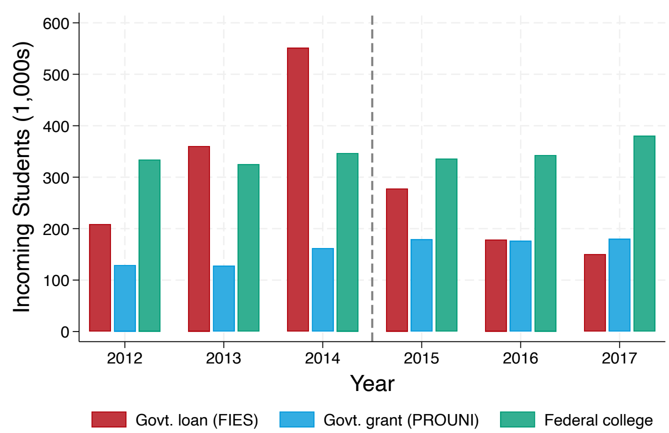

Figure A.3: Incoming Students by Funding Source, 2012–2017

Notes: The figure reports the number of incoming students (in thousands) by enrollment category and year. Government loans and government grants refer to students enrolled in private institutions using the respective federal funding instruments. Federal college refers to students enrolled in federal higher education institutions, which are tuition free. Years correspond to entry cohorts. The vertical line marks the FIES reform.

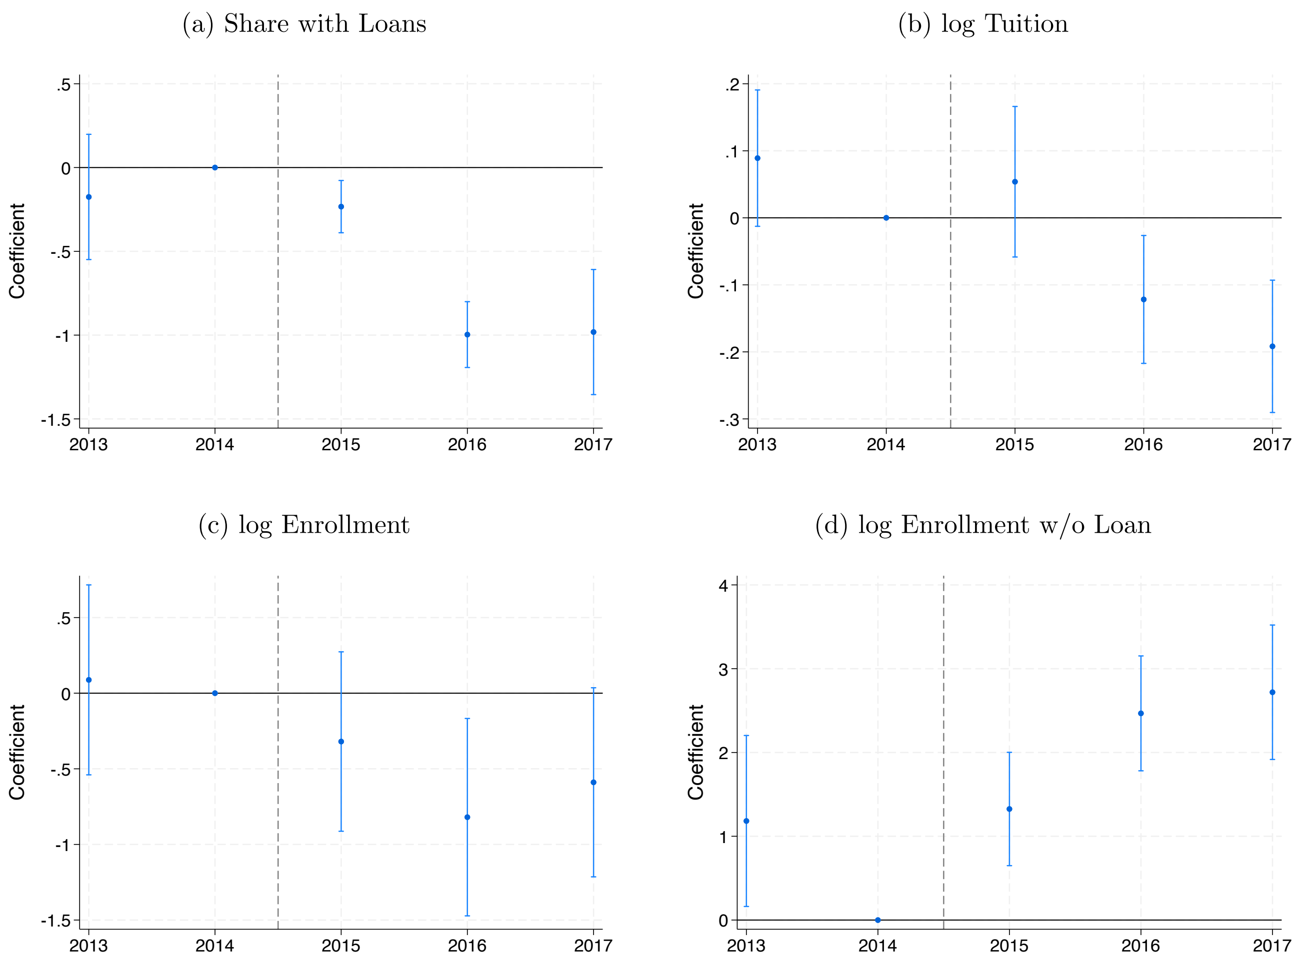

Figure A.4: Differential Exposure to the 2015 FIES Reform – Robustness

Notes: This figure presents OLS estimates of β_l from Equation (7), controlling for region-major-year fixed effects. Error bars denote 95% confidence intervals, with standard errors clustered at the college-year level. Panel A shows the share of incoming students with loans; Panel B, the log monthly tuition fee; Panel C, the log number of incoming students; and Panel D, the log number of incoming students not using a government loan. Tuition prices are calculated as enrollment-weighted averages of full and discounted prices for each degree-year and deflated to 2014 levels. The vertical line indicates the reform's announcement.

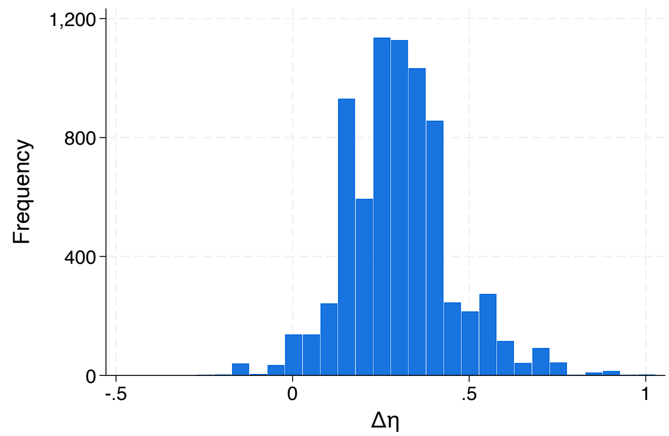

Figure A.5: Distribution of $\Delta\eta$

Notes: This figure shows the distribution of $\Delta\eta$, defined as the log difference between the average income of all students at a given college and that of FIES loan recipients at the same college, based on 2012 data. See Section 4.3 for further details.

(a) Price elasticities (b) Marginal Costs (c) Markups

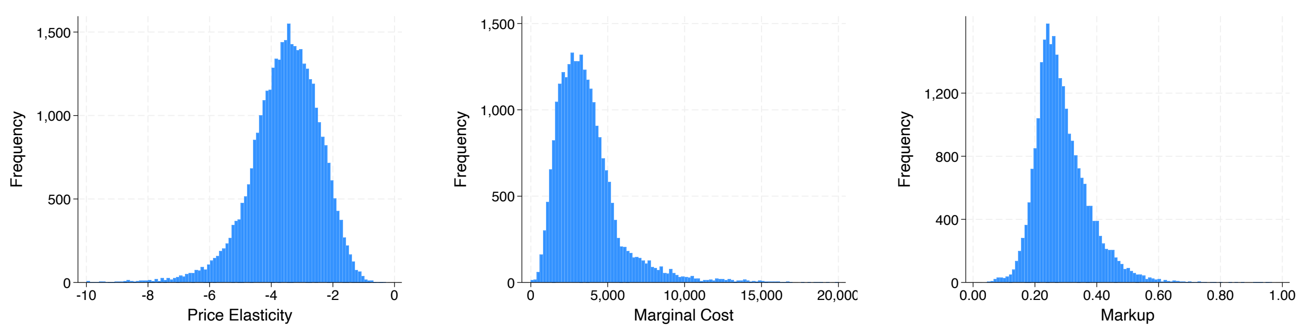

Figure A.6: Price Elasticities, Marginal Costs, and Markups

Notes: Panel A displays the distribution of price elasticities; Panel B shows marginal costs; and Panel C presents the implied markups. All quantities are derived from the model described in Section 5, using the estimated parameters from Section 6.

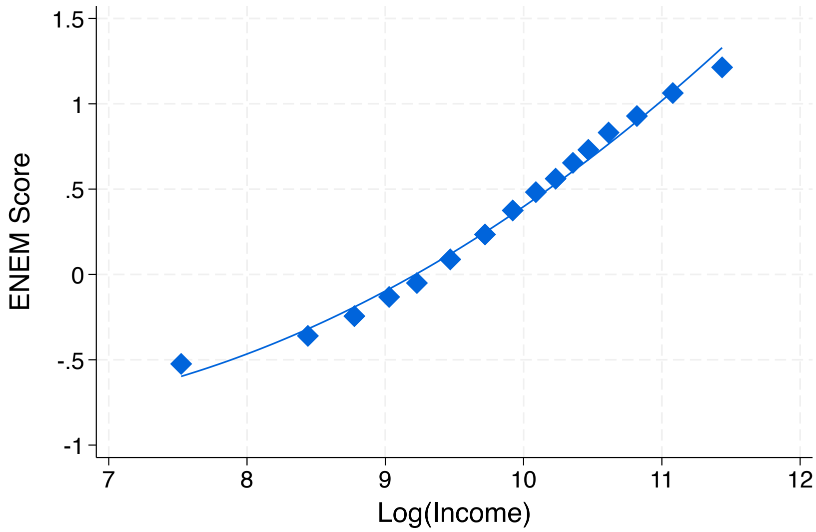

Figure A.7: Financial Aid Allocation, ENEM Scores, and Income

Notes: This figure presents the relationship between financial aid allocation, test scores, and household income. The x-axis reports log household income (in annual dollars). The y-axis shows average ENEM scores. The figure displays binned scatter plots with equally sized bins, and the lines represent quadratic fits. The sample includes two pre-reform years (2013 and 2014) and two post-reform years (2016 and 2017), and is restricted to regions with at least 5,000 ENEM takers and 1,000 incoming college students per year.

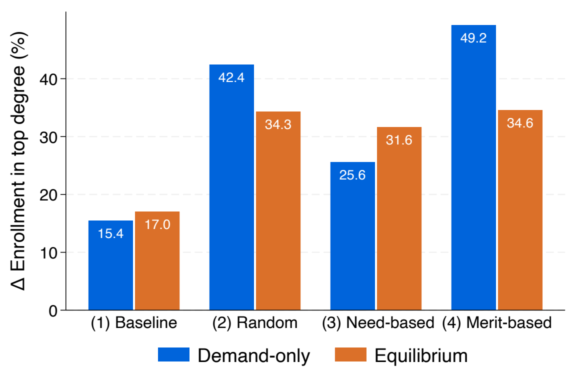

Figure A.8: Equilibrium effects of loans on top enrollment—Alternative Quality Measure

Notes: This figure reports the effects of student loans on enrollment in top-degree programs under alternative targeting schemes. Top-degree programs are defined as those in the top 10% of the quality distribution. Degree quality is measured by the average ENEM score of incoming students. Outcomes are expressed as percentage changes relative to a no-loan scenario. For each scheme, we set the coverage rate ρ so that total program spending equals the baseline budget (i.e., the counterfactuals are budget-neutral). Orange bars allow for endogenous tuition adjustments, while blue bars hold tuition fixed at observed levels. See Section 7 for details on the simulation procedure.

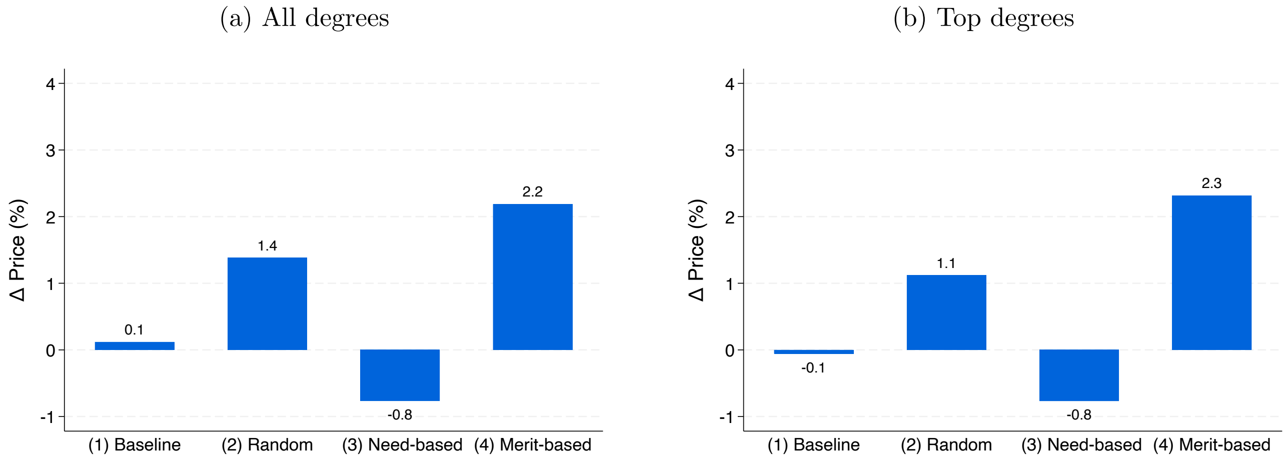

Figure A.9: Impacts of loans on prices in alternative *budget-neutral* scenarios

Notes: This figure shows the effect of student loans on tuition under alternative targeting schemes. The y-axis reports the percentage change in average tuition relative to a scenario with no loans. Average tuition is calculated as the enrollment-weighted average of full and discounted prices across degrees. For each targeting scheme, simulations are conducted at a coverage rate ρ that ensures a constant total loan program budget equal to the baseline level (*budget-neutral*). Panel A presents effects on tuition across all programs, while Panel B focuses on programs in the top 10% of the quality distribution. Degree quality is measured by the average earnings of its graduates. See Section 7 for additional details on the simulation procedure.

Table A.1: Price Effects of the 2015 Policy Change – Extended Sample

Dep var: $\log p_{jt}$	(1)	(2)	(3)	(4)	(5)
post × Exp	-0.096***	-0.093***	-0.088***	-0.090***	-0.091***
	(0.026)	(0.020)	(0.021)	(0.018)	(0.018)
post × Exp × $\Delta\eta$			0.451***	0.393***	0.413***
			(0.153)	(0.148)	(0.147)
Observations	22569	19911	22569	19911	19911
Degree FE	X	X	X	X	X
Region-Year FE	X	X	X	X	X
Major-Region-Year FE		X		X	X
Δ Inc-Year FE					X

Notes: This table reports OLS estimates of Equation (8). The dependent variable is the log of tuition for each degree-year. Tuition is calculated as the enrollment-weighted average of full and discounted prices, deflated to 2014 levels. "Exposure" measures a degree's exposure to the 2015 FIES reform; "post" indicates post-reform years; and $\Delta\eta$ is the log difference between the average income of all students and the average income of FIES loan recipients at each college in 2012. The sample goes from 2012 to 2017. Standard errors, clustered at the college-year level, are in parentheses. Asterisks indicate statistical significance: * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$.

Table A.2: Model Estimates

Panel A: Preferences	
β^h	3.440
	(0.007)
β^w	-1.066
	(0.002)
α_L^0	0.411
	(0.002)
α_{wL}^0	-0.153
	(0.002)
α^1	-5.101
	(0.001)
α_w^1	-0.417
	(0.000)
α_L^1	-0.082
	(0.003)
α_{wL}^1	-0.057
	(0.001)
Panel B: Financial Aid Targeting	
ρ_w^L	-2.2710
	(0.0115)
ρ_w^D	0.1711
	(0.2659)

Notes: This table reports GMM estimates of the model described in Section 5. The estimation sample includes two pre-reform years (2013 and 2014) and two post-reform years (2016 and 2017), and is restricted to regions with at least 5,000 ENEM takers and 1,000 incoming college students per year. Section 6 provides additional details on the sample and estimation procedure. Standard errors, clustered at the degree level, are in parentheses.

Table A.3: Model Fit: Targeted Micro-Moments

	Data	Model
E[score|enrolled]	0.532	0.54
E[income|enrolled]	9.13	9.144
E[income|enrolled, high price degree]	9.366	9.363
E[income|enrolled w/ loan]	8.79	8.785
E[income|enrolled with discount]	9.006	9.011
Enrollment discountinuity, inc 1	0.505	0.491
Enrollment discountinuity, inc 2	0.496	0.468
Enrollment discountinuity, inc 3	0.514	0.481
Enrollment discountinuity, inc 4	0.484	0.488
Tuition expenses discontinuity, inc 1	8.766	8.752
Tuition expenses discontinuity, inc 2	8.752	8.737
Tuition expenses discontinuity, inc 3	8.758	8.733
Tuition expenses discontinuity, inc 4	8.687	8.721

Notes: This table evaluates the fit of the model described in Section 5, estimated via GMM as detailed in Section 6.2. Each row corresponds to a micro-moment targeted in estimation. The column labeled "Data" reports the empirical value of each moment, while the "Model" column reports the corresponding simulated value implied by the estimated model. "Score" refers to performance on the centralized high school exit exam (ENEM), normalized to lie between 0 and 1. "Income" and "Price" are expressed in log annual dollars. "Discontinuity" captures enrollment and tuition expenditure discontinuities at the loan eligibility threshold, as described in Appendix H.1, and is reported separately by income quartile. Additional details on the construction and interpretation of these micro-moments are provided in Section 6.1.

B Derivation of Equation (5)

We begin by writing the firm's first-order condition from the profit maximization problem:

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

Solving for p:

$$p=c-\frac{S}{\frac{\partial S}{\partial p}}. \tag{B.1}$$

Step 1: Differentiate both sides of Equation (B.1) with respect to ρ

$$\frac{dp}{d\rho}=-\frac{d}{d\rho}\left(\frac{S}{\frac{\partial S}{\partial p}}\right)=-\left[\frac{\frac{\partial S}{\partial \rho}\cdot\frac{\partial S}{\partial p}+(\frac{\partial S}{\partial p})^2\frac{dp}{d\rho}-S\cdot\frac{\partial^2 S}{\partial p\partial\rho}-\frac{\partial^2 S}{\partial p^2}\cdot\frac{dp}{d\rho}\cdot S}{\left(\frac{\partial S}{\partial p}\right)^2}\right].$$

Now isolate $\frac{dp}{d\rho}$:

$$\frac{dp}{d\rho}=\frac{-\frac{\partial S}{\partial \rho}\cdot\frac{\partial S}{\partial p}+S\cdot\frac{\partial^2 S}{\partial p\partial\rho}}{\left(\frac{\partial S}{\partial p}\right)^2\cdot\left(2-\frac{S}{\left(\frac{\partial S}{\partial p}\right)^2}\cdot\frac{\partial^2 S}{\partial p^2}\right)}.$$

Therefore, $\frac{d\log p}{d\rho}$ can be written as

$$\frac{d\log p}{d\rho}=\frac{1}{p}\cdot\frac{dp}{d\rho}=\frac{1}{p}\cdot\frac{-\frac{\partial S}{\partial \rho}\cdot\frac{\partial S}{\partial p}+S\cdot\frac{\partial^2 S}{\partial p\partial\rho}}{\left(\frac{\partial S}{\partial p}\right)^2\left(2-\frac{S}{\left(\frac{\partial S}{\partial p}\right)^2}\cdot\frac{\partial^2 S}{\partial p^2}\right)}. \tag{B.2}$$

Step 2: Formulas for $\frac{\partial S}{\partial \rho}$ and $\frac{\partial^2 S}{\partial p\partial\rho}$

Recall Equation (4):

$$S=\int_0^\rho s_i^L\,di+\int_\rho^1 s_i^{NL}\,di,$$

where s_i^L and s_i^{NL} denote the probabilities that student i enrolls in college with and without a loan, respectively.

Differentiating S with respect to ρ and applying the Leibniz rule yields

$$\frac{\partial S}{\partial \rho}=s_\rho^L-s_\rho^{NL}, \tag{B.3}$$

where the subscript ρ indicates that the enrollment probabilities (s_ρ^L and s_ρ^{NL}) are computed for the student in the margin between receiving a loan or not.

Now differentiate S with respect to p:

$$\frac{\partial S}{\partial p} = \int_0^\rho \frac{\partial s_i^L}{\partial p}\, di + \int_\rho^1 \frac{\partial s_i^{NL}}{\partial p}\, di.$$

Differentiate $\frac{\partial S}{\partial p}$ with respect to ρ and apply the Leibniz rule again

$$\frac{\partial^2 S}{\partial p \partial \rho} = \frac{\partial s_\rho^L}{\partial p} - \frac{\partial s_\rho^{NL}}{\partial p}.$$

Step 3: Simplifying the numerator of Equation (B.2)

Define the following elasticities:

$$\eta = -\frac{p}{S} \cdot \frac{\partial S}{\partial p},$$
$$\eta_\rho^L = -\frac{p}{s_\rho^L} \cdot \frac{\partial s_\rho^L}{\partial p},$$
$$\eta_\rho^{NL} = -\frac{p}{s_\rho^{NL}} \cdot \frac{\partial s_\rho^{NL}}{\partial p}.$$

Using these elasticities, we can write $\frac{\partial S}{\partial \rho}$ and $\frac{\partial^2 S}{\partial p \partial \rho}$ as

$$\begin{aligned} \frac{\partial S}{\partial p} &= -\frac{\eta S}{p}, \\ \frac{\partial^2 S}{\partial p \partial \rho} &= -\eta_\rho^L \cdot \frac{s_\rho^L}{p} + \eta_\rho^{NL} \cdot \frac{s_\rho^{NL}}{p}. \end{aligned} \tag{B.4}$$

Substitute Equation (B.3) into the numerator of Equation (B.2):

$$-\frac{\partial S}{\partial \rho} \cdot \frac{\partial S}{\partial p} + S \cdot \frac{\partial^2 S}{\partial p \partial \rho} = -(s_\rho^L - s_\rho^{NL}) \cdot \frac{\partial S}{\partial p} + S \cdot \frac{\partial^2 S}{\partial p \partial \rho}. \tag{B.5}$$

Now substitute Equation (B.4) into Equation (B.5):

$$-\frac{\partial S}{\partial \rho} \cdot \frac{\partial S}{\partial p} + S \cdot \frac{\partial^2 S}{\partial p \partial \rho} = -(s_\rho^L - s_\rho^{NL}) \cdot \left(-\frac{\eta S}{p}\right) + S \cdot \left(-\eta_\rho^L \cdot \frac{s_\rho^L}{p} + \eta_\rho^{NL} \cdot \frac{s_\rho^{NL}}{p}\right). \tag{B.6}$$

Simplifying Equation (B.6):

$$-\frac{\partial S}{\partial \rho} \cdot \frac{\partial S}{\partial p} + S \cdot \frac{\partial^2 S}{\partial p \partial \rho} = \frac{S}{p}\left[\eta(s_\rho^L - s_\rho^{NL}) - \eta_\rho^L s_\rho^L + \eta_\rho^{NL} s_\rho^{NL}\right]. \tag{B.7}$$

Step 4: Simplifying the denominator of Equation (B.2)

Define curvature as

$$\lambda \equiv \frac{S \cdot \frac{\partial^2 S}{\partial p^2}}{\left(\frac{\partial S}{\partial p}\right)^2}. \tag{B.8}$$

Substitute Equations (B.4) and (B.8) into denominator of Equation (B.2):

$$\left(\frac{\partial S}{\partial p}\right)^2 \left(2 - \frac{S}{\left(\frac{\partial S}{\partial p}\right)^2} \cdot \frac{\partial^2 S}{\partial p^2}\right) = \left(\frac{\eta S}{p}\right)^2 (2 - \lambda). \tag{B.9}$$

Step 5: Inserting Equations (B.7) and (B.9) into (B.2)

Inserting Equations (B.7) and (B.9) into (B.2):

$$\begin{aligned}\frac{d \log p}{d\rho} &= \frac{1}{p} \cdot \frac{\frac{S}{p}\left[\eta(s_\rho^L - s_\rho^{NL}) - \eta_\rho^L s_\rho^L + \eta_\rho^{NL} s_\rho^{NL}\right]}{\left(\frac{\eta S}{p}\right)^2 (2-\lambda)} \\ &= \frac{1}{\eta^2} \frac{1}{2-\lambda} \frac{s_\rho^{NL}}{S} \cdot \left[\left(\eta_\rho^{NL} - \eta_\rho^L\right) - \left(\eta_\rho^L - \eta\right) \cdot \frac{s_\rho^L - s_\rho^{NL}}{s_\rho^{NL}}\right].\end{aligned}$$

Finally, defining $\Omega \equiv \frac{1}{\eta^2} \cdot \frac{1}{2-\lambda} \cdot \frac{s_\rho^{NL}}{S}$, we arrive at the expression for Equation (5) presented in the main text:

$$\frac{d \log p}{d\rho} = \Omega \cdot \left[\left(\eta_\rho^{NL} - \eta_\rho^L\right) - \left(\eta_\rho^L - \eta\right) \cdot \frac{s_\rho^L - s_\rho^{NL}}{s_\rho^{NL}}\right]. \quad \square$$

C The Brazilian Federal Student Loan Program (FIES): details

C.1 Share of tuition covered by aid

C.1.1 Pre-reform

Between 2010 and 2014, following *Portaria Normativa MEC 10, 2010*, the percentage of educational charges financed through FIES was based on the student's tuition burden as a share of their per capita gross family income. The percentage of income commitment is calculated using the formula established in Article 7:

$$\left(\frac{\text{VS}}{6} \div \text{RF}\right) \cdot 100,$$

where:

- VS is the total semester tuition (valor da semestralidade) charged to the student by the institution of higher education, inclusive of all collective and regular discounts offered by the institution, including prompt payment reductions, regardless of course frequency or modality;
- RF is the student's gross monthly per capita family income, computed by dividing the family's gross monthly income by the number of family members as defined in Article 8 of the regulation.

Article 6 defines the financing tier based on this income burden:

- 100% financing: Available when the income burden is greater than or equal to 60%.
- 75% financing: Available when the income burden is greater than or equal to 40% and less than 60%.
- 50% financing: Available when the income burden is greater than or equal to 20% and less than 40%.

C.1.2 Post-reform

Between 2015 and 2017, in accordance with Article 6 of Portaria Normativa MEC 10, 2015, the percentage of educational charges financed through FIES was based on the student's family income, as defined in Annex V of the regulation. Specifically, the financing percentage f is computed using the formula

$$f = \left[1 - \left(\frac{k_i^m \cdot R_i - d_i}{m}\right)\right] \cdot 100,$$

where:

- f: percentage of tuition covered by FIES,

- k_i^m: marginal income commitment rate for income bracket i,
- R_i: gross monthly per capita family income in reais,
- d_i: deductible amount in reais for income bracket i,
- m: tuition charged by the higher education institution in reais.

The parameters k_i^m and d_i vary by income bracket, as specified in the table below. To ensure that students participate in the cost of education, the student's minimum out-of-pocket contribution is subject to a floor known as the minimum participation value (VMP). If the implied contribution $(k_i^m \cdot R_i - d_i)$ falls below the VMP, the student must instead contribute the VMP.

Table C.1: Parameters for FIES Financing Formula

Income Bracket	**Marginal** k_i^m	**Deduction** d_i	**VMP**	**Effective Rate**
Up to 0.5 minimum wage	15.00%	0.00%	50.00%	15.00%
0.5 to 1.0 minimum wage	26.50%	45.31%	50.00%	20.75%
1.0 to 1.5 minimum wage	38.00%	135.93%	50.00%	26.50%
1.5 to 2.0 minimum wage	49.50%	271.86%	50.00%	32.25%
2.0 to 2.5 minimum wages	61.00%	453.10%	50.00%	38.00%

Notes: This table is taken from Annex V of Portaria Normativa MEC 10, 2015. The income brackets reported in the first column represent household income and are expressed in terms of minimum wages per capita. The column labeled "Effective Rate" shows the average student contribution as a percentage of income within each bracket.

C.2 Loan allocation

Starting in 2015, FIES introduced a centralized allocation mechanism that assigned subsidized loan contracts to eligible applicants based on ENEM scores. The allocation process had two stages: degree-level seat allocation and student-level assignment.

Seat Allocation to Institutions. Each semester, the Ministry of Education (MEC) sets a national cap on the number of FIES contracts available. The MEC allocates this cap across degrees in three stages: first by geographic location, then by field of study, and finally by degree quality (as measured by SINAES[44] scores). In the first stage, seats are allocated across locations proportionally to estimated demand. Demand is measured using the number of ENEM takers and the number of students who applied for FIES in 2015. This demand measure is then multiplied by an adjustment factor ranging from 1.3 to 0.7, with higher factors assigned to locations with lower Human Development Index. In the second stage, within each location, 70% of seats are reserved for priority fields—health, engineering, and teacher education—while the remaining 30% are allocated to other disciplines. In the final stage, within each field-location cell, seats are distributed across degrees based on their quality scores. Importantly, these quality scores are essentially uncorrelated with pre-reform FIES shares (correlation = 0.011).

Student application and ranking. Seats approved by the MEC were published on the FIES Seleção platform. Eligible students—those with per capita family income below 2.5 times the minimum wage and a minimum ENEM score of 450 (plus a nonzero essay score)—could apply during a 4-day window. In 2015, students could select only one degree. From 2016 onward, applicants could select up to three degree options. Students were ranked by their ENEM average score, with tie-breakers applied sequentially: essay score, linguistics, math, natural sciences, and humanities. Spots are allocated through an iterative deferred acceptance mechanism, in which students are sequentially asked to submit rank-ordered lists over the course of several trial days. At the end of each day, the system produces a cutoff grade that is the lowest grade necessary to be accepted by a specific program. Students can change their choices freely during the trial period. Only the choices made on the last day are used to allocate loans. See Bo and Hakimov (2019) for a comprehensive discussion of the formal properties of the iterative deferred acceptance mechanism.

References. The primary rules were set by *Portaria Normativa MEC 8/2015* (for 2015) and *Portaria Normativa MEC 13/2015* (for 2016 onward).

[44] SINAES (Sistema Nacional de Avaliação da Educação Superior) is Brazil's federal higher education quality assessment system. It combines student performance on a national college exit exam (ENADE) with evaluations of institutions and degree programs, including factors such as infrastructure and faculty qualifications.

D Tuition data

We construct degree-level tuition prices by integrating data from four sources. The first two sources are administrative records from Brazil's government fellowship and loan programs, PROUNI and FIES, obtained from the National Education Fund (FNDE). These records track government payments, which enable us to estimate tuition fees at participating institutions. The third source is a nationally representative survey conducted by Hoper, a consultancy specializing in higher education. The fourth source is administrative data from QueroBolsa, Brazil's largest degree search platform.

All four datasets report posted tuition fees, which we define as full prices, while all but Hoper also provide discounted prices. PROUNI reports the average discounted price paid in each degree-year. FIES provides individual-level tuition data but lacks identifiers to link students to other sources. Thus, we compute the degree-year discounted price as the mean discount among all students receiving aid in that year. QueroBolsa reports all discounts offered through its platform, which we aggregate to obtain a degree-year average. Appendix E.1 shows that approximating each degree with only two prices—full and discounted—is a reasonable representation of the data.

To estimate full prices, we regress log-prices on degree-year and source fixed effects and use the degree-year coefficients as estimated prices. This approach accounts for systematic differences across sources. In cases in which information for a certain year is missing, we run a regression of the predicted price on degree and year fixed effects, imputing the missing values based on the sum of the coefficients.

To estimate discounted prices, we first compute discount rates as the difference between log full and log discounted prices. We then apply the same procedure to combine datasets and impute missing values, incorporating additional steps when necessary. If the discount rate is missing for a degree in all years, we regress it on major, college, and year fixed effects. If no data are available for an entire college, we impute values using major and year fixed effects.

Finally, we reconstruct discounted prices by applying the estimated discount rates to full prices. Note that we observe an individual-level indicator of discount usage for all enrolled students in the Census of Higher Education. The procedure described here serves only to estimate discount magnitudes.

Following this methodology, we recover full and discounted prices for 95% of degree-years, covering 98% of total enrollment. All prices are deflated to 2014 price levels using the Brazilian consumer price index (IPCA).

E Price discrimination: details

This appendix provides further details on how colleges price discriminate in our setting, with the goal of validating the modeling assumptions underlying the framework in Section 5. We examine three elements of how we model colleges' pricing decisions.

First, we present empirical evidence supporting our two-price representation of the data, which aggregates all students receiving a discount within a given degree into a single "discounted price." Second, we discuss the assumption that colleges cannot choose which students receive discounts. Third, we examine the assumption that the structural errors governing discount behavior in the model remain invariant to changes in loan policy.

E.1 Assessing within–degree-year price variation

For simplicity, our model assumes that colleges can charge only two prices per degree. In practice, this requires collapsing multiple discount levels into a single representative "discounted price."

To evaluate how restrictive this assumption is, we use data from QueroBolsa, which reports the range of discounts offered for a subset of degrees. The dataset includes 22,083 unique degree–year observations. On average, the largest discounted price in a given degree–year is 3.4% above the mean discounted price, while the smallest is 3.6% below, with both differences expressed as a share of full tuition. This modest dispersion contrasts with the much larger 35% average gap between discounted and full prices.[45] These results indicate that representing each degree with only two prices provides a reasonable approximation of the underlying price variation.

E.2 Discount targeting

In our model, colleges cannot choose which students receive discounts within a degree. We validate this assumption in two steps. First, Appendix E.2.1 examines the characteristics of students who receive discounts and shows that they are not consistent with alternative models in which colleges optimally target discounts. Second, Appendix E.2.2 demonstrates that discount targeting is uncorrelated with loan targeting.[46]

E.2.1 Characteristics of Students Receiving Discounts

We begin by comparing the characteristics of students who receive tuition discounts with those who pay the full price. If colleges were optimally targeting discounts, we would expect discounted status to be strongly correlated with certain student characteristics. Consider a few examples. If colleges

[45] The 33% average discount rate reported in Section 3.1 is computed across all degrees for which we have price data. The 35% figure here is computed using the QueroBolsa subsample used for the within-degree price-dispersion analysis in this subsection of Appendix E.

[46] All results in Appendix E.2 use data from 2014, one year before the FIES reform. Results are similar for other years.

were maximizing profits by targeting students with lower willingness to pay—as in standard models of price discrimination—discounted students would have lower income. If they were targeting high-ability students, discounted students would have higher exam scores. If they were targeting students with stronger social networks, discounted students would come from more educated families.

Panel A of Figure E.1 plots the share of enrolled students receiving a tuition discount as a function of household income (blue diamonds). For comparison, it also shows the share of students taking up a government loan (orange dots).

The figure shows that discount receipt is only weakly correlated with income, suggesting that colleges do not successfully target discounts toward students with the lowest willingness to pay. There are, however, at least two caveats to this conclusion. First, income is self-reported and may contain significant measurement error. However, the strong relationship between income and loan take-up (also shown in Panel A) indicates that the income measure is informative. Second, colleges may target discounts based on characteristics other than income that are more predictive of willingness to pay. To explore this, Panel B of Figure E.1 plots the relationship between income and the sticker price of the degrees in which students enroll. High-income students are substantially more likely to enroll in higher-priced degrees. Thus, if discounts were highly correlated with willingness to pay, it would be unlikely that their allocation would be so weakly correlated with income.

Next, we examine whether other student characteristics are associated with receiving discounts. We estimate the following linear probability model:

$$\text{Aid}_i = \beta_0 + \beta_x x_i + \varepsilon_i, \quad \text{(E.1)}$$

where i indexes students, Aid_i is an indicator for receiving financial aid (either a discount or a loan), and x_i is a vector of student characteristics. The vector x_i includes log household income, standardized exam scores, an indicator for attending a public high school, and an indicator for whether either parent attended college. We estimate separate regressions using either discount receipt or loan take-up as the dependent variable.

The results, shown in columns (1) and (4) of Table E.1, confirm the patterns observed in Figure E.1. Discount receipt is weakly correlated with student characteristics: coefficients are small and statistically insignificant, and the R^2 is below 1%. In contrast, loan take-up is strongly associated with low-SES characteristics.

One concern with these results is that they combine variation both within and across degrees, which could be misleading due to student sorting. For example, suppose colleges indeed target low-income students with their discounts, but more expensive degrees—which enroll fewer low-income students—offer more discounts. In that case, we might find no overall correlation between income and discount status, mistakenly concluding that colleges are not targeting discounts. However, Columns (2)–(3) and (5)–(6) include firm and degree fixed effects, respectively, and the conclusions remain unchanged.

This suggests that student sorting across degrees is not driving our results.

Overall, these results indicate that discounts are poorly targeted in our setting, consistent with our assumption that colleges cannot choose which students receive discounts. This pattern is expected for two reasons. First, the government does not share information about students' financial aid applications with colleges, unlike in the U.S. context (Fillmore, 2023). Second, a large share of discounts are obtained through platforms such as QueroBolsa, which impose very few conditions for receiving a discount.

(a) Financial Aid and Income　　　(b) Tuition Prices and Income

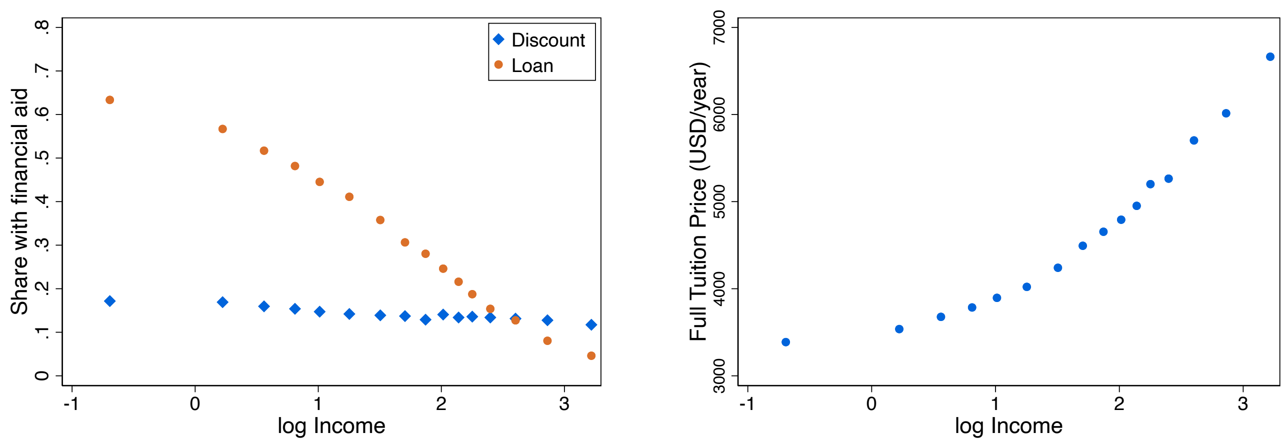

Figure E.1: Financial Aid, Tuition Prices, and Income

Notes: This figure plots the relationship between financial aid allocation and log annual household income. Panel A shows the share of enrolled students receiving tuition discounts (blue diamonds) and government loans (orange dots). Panel B shows the average full tuition price of the degree in which students are enrolled. Each panel presents a binned scatter plot with equal-size bins. The data are from 2014.

E.2.2 Discount versus Loan Targeting

For the purposes of this paper, our interest in discounts lies in their potential to attenuate the composition effect. Therefore, the key question is not only how well discounts are targeted overall, but also how closely discount and loan targeting are correlated. It is possible that colleges have limited capacity to target discounts—as the evidence in Appendix E.2.1 suggests—but use whatever targeting capacity they possess to prevent students with loans from also receiving discounts.

To investigate this issue, we first estimate Equation (E.1) using loan usage as the dependent variable and use the estimated coefficients to compute the predicted probability of loan usage for each student. We then relate this predicted probability to the actual usage of both loans and discounts.

Figure E.2 shows the results. As expected, students with a higher estimated predicted probability of using a loan are indeed more likely to use one. In contrast, there is no relationship between the predicted probability of loan usage and the likelihood of receiving a discount.

These results indicate that colleges do not target discounts to students who are more or less likely

Table E.1: Correlates of Financial Aid Usage

	Discount			Loan		
	(1)	(2)	(3)	(4)	(5)	(6)
log(income)	-0.017	-0.009	-0.001	-0.109***	-0.090***	-0.078***
	(0.022)	(0.007)	(0.004)	(0.013)	(0.007)	(0.004)
Public High-School	-0.003	0.017***	0.006	0.097***	0.099***	0.114***
	(0.024)	(0.007)	(0.006)	(0.021)	(0.016)	(0.008)
College-educated Parents	-0.002	-0.005	0.004	-0.045***	-0.046***	-0.053***
	(0.011)	(0.007)	(0.003)	(0.009)	(0.007)	(0.005)
Score	0.003	-0.002	0.007***	-0.029***	-0.015***	-0.018***
	(0.007)	(0.004)	(0.002)	(0.008)	(0.004)	(0.002)
Observations	233095	233095	233095	233095	233095	233095
R^2 (within)	0.00	0.00	0.00	0.10	0.07	0.07
Firm FE		X	X		X	X
Degree FE			X			X

Notes: This table reports estimates from linear probability models of financial aid usage. The dependent variable is an indicator for receiving a tuition discount (columns 1–3) or a government loan (columns 4–6). All regressions include log income, standardized exam scores, and indicators for public high-school attendance and parental college education. Columns (2) and (5) include firm fixed effects; columns (3) and (6) include degree fixed effects. The data are from 2014. Standard errors are clustered at the firm level. Asterisks denote statistical significance: * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$.

to use government loans. This pattern is consistent with two features of our context. First, as shown above in Appendix E.2.1, colleges do not appear to target discounts effectively. Second, regulations prohibit colleges from conditioning discounts on whether students receive government financial aid.

E.3 Testing the exogeneity of structural residuals

This appendix validates our assumption that the structural residuals governing colleges' discount behavior in our model—κ and ρ^D—are invariant to changes in loan policy. To test this, we examine whether these residuals were affected by the 2015 FIES reform.

As described in Section 5, κ represents residual variation in the magnitude of discounts that is not explained by endogenous responses to demand differences between full-price and discount students. Formally, colleges behave as if the marginal cost of enrolling a discount student were $c^D = c - \kappa$, where c is the marginal cost for full-price students. Similarly, ρ^D captures residual variation in the number of students receiving discounts across degrees, reflecting heterogeneity in discount availability that is not driven by student selection.

To examine whether these structural residuals responded to the reform, we estimate the following difference-in-differences regression, which mirrors the specification in Section 4:

$$\text{Outcome}_{jt} = \gamma_j + \gamma_{rt} + \beta^{Outcome} \cdot (\text{post}_t \times \text{Exp}_j) + \varepsilon_{jt}, \tag{E.2}$$

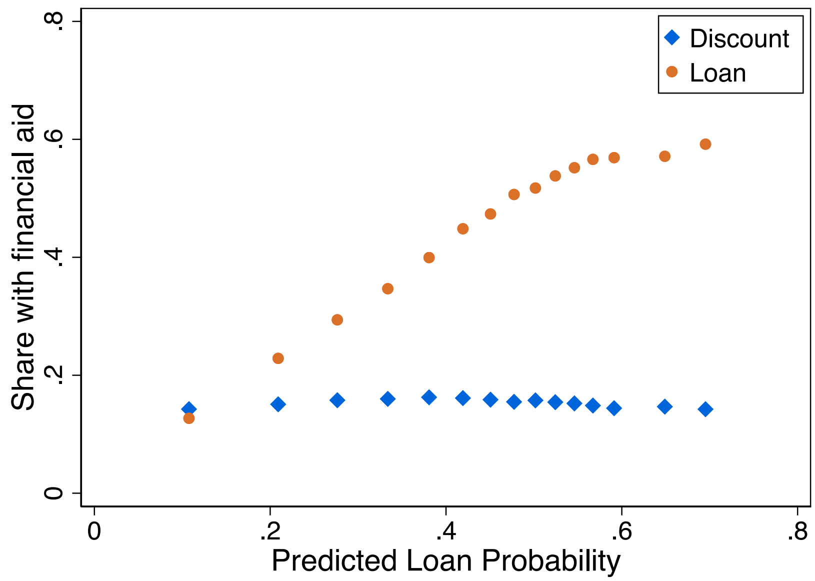

Figure E.2: Financial Aid by Predicted Probability of Loan Usage

Notes: This figure plots the relationship between the predicted probability of government loan receipt and the actual receipt of loans (orange squares) or tuition discounts (blue dots). Predicted probabilities are obtained from a linear probability model that includes log income, exam scores, and socioeconomic indicators. The figure presents binned scatter plots with equal-size bins and absorbs degree fixed effects. The data are from 2014.

where j indexes degrees, t years, and Outcome_{jt} is one of three variables: the marginal cost for discount students (c^D), markup for discount students (μ^D), or the structural residual determining discount availability (ρ^D). The coefficient of interest, $\beta^{Outcome}$, measures how each outcome responded to exposure to the reform. The specification includes degree fixed effects (γ_j) and region-by-year fixed effects (γ_{rt}). Exposure to the reform, Exp_j, is defined as in Section 4, and post_t equals one for post-reform years.

Equation (E.2) permits the following decomposition:

$$\beta^{p^D} = \beta^{c^D} + \beta^{\mu^D},$$

so that any change in discounted prices, p^D, associated with exposure to the reform can be attributed to either changes in costs or markups. If κ is truly invariant, only markups should respond to the reform.

Table E.2 presents the results. Degrees more exposed to the reform experienced sizable reductions in markups, while the estimated effects on marginal costs are small and statistically insignificant. Because c^D is defined as $c^D = c - \kappa$, these results indicate that the structural error governing the magnitude of discounts—captured by κ—did not change in response to the reform.

Next, we examine the structural error governing the prevalence of discounts, ρ^D. We find no statistically significant effect of the reform on ρ^D. Moreover, the estimated coefficient is economically small: a one–standard-deviation increase in exposure is associated with only a 0.1% change in discount availability.

Together, these results indicate that both κ and ρ^D remained stable after the 2015 FIES reform, supporting their treatment as invariant structural residuals in the model estimation and counterfactual simulations.

Table E.2: Effects of the 2015 FIES Reform on Discounted Prices, Markups, and Structural Residuals

	Markup	Marginal Cost	ρ^D
	(1)	(2)	(3)
post × Exp	-15.962***	-0.779	-0.040
	(5.506)	(11.172)	(0.040)
Observations	9898	9898	9898
Degree FE	Yes	Yes	Yes
Region–Year FE	Yes	Yes	Yes

Notes: This table reports regressions of degree-level outcomes on the interaction between exposure to the 2015 FIES reform and a post-reform indicator, controlling for degree and region-by-year fixed effects. Outcomes include the marginal cost of discount students (c^D), the markup (μ^D), and the residual determining the number of discounts (ρ^D). The sample includes two pre-reform years (2013–2014) and two post-reform years (2016–2017), and includes only degrees that have at least one student with a tuition discount. Standard errors, clustered at the college-year level, are reported in parentheses. Asterisks denote statistical significance: * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$.

F Online education

In this appendix, we discuss the expansion of online education in Brazilian higher education and show that our results are robust to accounting for this expansion.

F.1 Measuring Exposure to Online Competition

Online education expanded substantially in Brazil during the period we study, 2012 to 2017. The share of incoming students enrolled in online degrees increased from 19.7% in 2012 to 33.3% in 2017.

We measure the exposure of in-person degrees to competition from online alternatives following Barahona et al. (2025), who highlight an important institutional feature of the online higher education market in Brazil. Although instruction in online degrees is fully remote, online programs require certain in-person components, such as exams and laboratory work. These activities are conducted either at the institution's main campus or at designated local hubs. As a result, the geographic reach of online programs is limited by the location of these hubs.

We therefore define the exposure of in-person degree j to online competition in year t, $ExpOnline_{jt}$, as the number of online degrees in the same major as degree j that operate a local hub in degree j's region[47] in year t. There is substantial variation in exposure. In 2017, the median in-person degree in our sample faces 6 online competitors, while 18% face 100 or more and 3% face 1,000 or more.

F.2 Assessing Robustness to Online Competition

To evaluate whether the expansion of online education is driving our results, we augment Equation (8) with a control for online competition:

$$\log p_{jt} = \gamma_j + \gamma_{rt} + \beta_0 \cdot \text{post}_t \cdot \text{Exp}_j + \beta_C \cdot \text{post}_t \cdot \text{Exp}_j \cdot \Delta\eta_j + \delta \cdot \text{ExpOnline}_{jt} + \epsilon_{jt}, \tag{F.1}$$

where $\log p_{jt}$ denotes the log of the average tuition for degree j in year t, post_t is an indicator for the years following the 2015 reform, γ_j denotes degree fixed effects, and γ_{rt} denotes region-year fixed effects. The variable Exp_j, defined in Section 4.2, measures degree j's exposure to the 2015 reform.

OLS estimates of Equation (F.1) are reported in Table F.1. Columns (1) and (2) reproduce our baseline specifications from Table 1 in the main text, which do not control for online exposure. Columns (3) and (4) add the control for online competition.

Two patterns emerge. First, greater exposure to online programs is associated with lower tuition in in-person degrees, consistent with increased competitive pressure from online alternatives. Second, the estimated effect of the 2015 FIES reform on tuition is essentially unchanged once we control for online

[47] As in the rest of the paper, we define *regions* using the Brazilian Institute of Geography and Statistics (IBGE) meso-region classification, which groups Brazil's 5,568 municipalities into 137 meso-regions based on geographic proximity and shared socioeconomic characteristics.

exposure. This indicates that the expansion of online education is not driving our main results.

Table F.1: Effects of the 2015 FIES Reform on Tuition – Robustness to Online Competition

Dep var: $\log p_{jt}$	(1)	(2)	(3)	(4)
post × Exp	-0.106***	-0.108***	-0.099***	-0.100***
	(0.028)	(0.028)	(0.022)	(0.022)
post × Exp × $\Delta\eta$			0.503***	0.494***
			(0.151)	(0.151)
ExpOnline		-0.033***		-0.032***
		(0.008)		(0.008)
Observations	13447	13447	13447	13447
Degree FE	X	X	X	X
Region-Year FE	X	X	X	X

Notes: This table reports OLS estimates of Equation (F.1). The dependent variable is the log of tuition for each degree-year observation. Tuition is calculated as the enrollment-weighted average of full and discounted prices and is deflated to 2014 reais. "Exposure" measures a degree's exposure to the 2015 FIES reform; "post" indicates post-reform years; and $\Delta\eta$ is the log difference between the average income of all students and the average income of FIES loan recipients at each college in 2012. $ExpOnline_{jt}$ measures exposure to online competition and is defined as the number of online degrees in the same major operating a hub in the degree's region in year t. This variable is divided by 1,000 in the regressions to improve readability of the reported coefficients. The sample includes two pre-reform years (2013 and 2014) and two post-reform years (2016 and 2017). Standard errors, clustered at the college-year level, are reported in parentheses. Asterisks denote statistical significance: * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$.

G Public colleges and the outside option

This appendix discusses the modeling choice made in Section 5 to incorporate public colleges into the outside option. We proceed in three steps. First, Appendix G.1 shows that the model presented in Section 5 is equivalent to a model that explicitly includes public colleges. The key assumption underlying this result is that admission cutoffs for public degrees are not affected by changes in the availability of student loans for private degrees. Second, Appendix G.2 examines the empirical plausibility of this assumption by documenting large differences in the types of students who attend private versus public colleges. Nevertheless, the two markets are not completely segmented, and the assumption of fixed admission cutoffs should therefore be viewed as an approximation rather than an exact description. Third, Appendix G.3 evaluates the sensitivity of our results to this approximation by reestimating the model after excluding students who enroll in public colleges from the estimation sample. We find that the results are robust, indicating that the treatment of public colleges as part of the outside option is not driving our findings.

G.1 Extended framework with public colleges

This appendix shows that a model that explicitly includes tuition-free but selective public colleges is equivalent to the model in Section 5 in which public options are absorbed into a score-dependent outside option. For clarity, we focus on a single market throughout this derivation and omit the region r and time t subscripts.

Environment

Let $\mathcal{J}$ denote the set of private degrees, as in Section 5, and let $\mathcal{G}$ denote the set of public degrees.

Each public degree $g \in \mathcal{G}$ satisfies:

1. *Zero tuition:* $p_{ig}^{pub} = 0$ for all i.
2. *Selectivity:* exogenous cutoffs $\widetilde{h}_g$ such that student i can enroll in g only if $h_i \geq \widetilde{h}_g$.

Define the set of public degrees available to student i as

$$\mathcal{G}(h_i) \equiv \{g \in \mathcal{G} : h_i \geq \widetilde{h}_g\}. \tag{G.1}$$

Utility from private degree $j \in \mathcal{J}$ is specified exactly as in Section 5. The no-college outside option (not enrolling in any degree) yields utility

$$U_{i0}^{base} = V_0^{base}(w_i, h_i) + \varepsilon_{i0}^{base}, \tag{G.2}$$

while each public degree $g \in \mathcal{G}(h_i)$ yields

$$U_{ig}^{pub} = V_g^{pub}(w_i, h_i) + \varepsilon_{ig}^{pub}. \tag{G.3}$$

The functions $V_0^{base}(\cdot)$ and $V_g^{pub}(\cdot)$ capture the mean utility of the no-college outside option and of public degrees, respectively, and are allowed to depend on student income w_i and score h_i. All idiosyncratic shocks

$$\{\varepsilon_{ij}\}_{j \in \mathcal{J}}, \quad \varepsilon_{i0}^{base}, \quad \{\varepsilon_{ig}^{pub}\}_{g \in \mathcal{G}(h_i)}$$

are assumed to be i.i.d. Type-I extreme value with a common scale parameter normalized to one, as is standard in multinomial logit choice models.

Collapsing public colleges into the outside option

Let the set of non-private alternatives for student i be

$$\mathcal{K}(h_i) \equiv \{\text{no-college outside}\} \cup \mathcal{G}(h_i).$$

Define the *augmented outside option* as the maximum utility over all non-private alternatives:

$$U_{i0}^{aug} \equiv \max_{k \in \mathcal{K}(h_i)} U_{ik}. \tag{G.4}$$

Note that a model in which students choose among private degrees and this augmented outside option is *exactly equivalent* to a model in which students choose among private degrees, all public degrees available to them, and the no-college outside option. The two formulations are equivalent in the sense that each private degree faces exactly the same demand curve in both models. As a result, equilibrium prices are identical, and so are equilibrium enrollment shares.

Moreover, the structure of the augmented outside option is known and admits a closed-form representation:

$$U_{i0}^{aug}(w_i, h_i) \overset{d}{=} V^{aug}(w_i, h_i) + \tilde{\varepsilon}_{i0}, \tag{G.5}$$

where

$$V^{aug}(w_i, h_i) \equiv \log\left(\exp\big(V_0^{base}(w_i, h_i)\big) + \sum_{g \in \mathcal{G}(h_i)} \exp\big(V_g^{pub}(w_i, h_i)\big) \right), \tag{G.6}$$

and $\tilde{\varepsilon}_{i0}$ is independent of the idiosyncratic shocks associated with private degrees and follows a Type-I extreme value distribution with the same scale parameter as all other shocks.

The systematic component of the augmented outside option, $V^{aug}(w_i, h_i)$, is in general a potentially complex function of student income w_i and score h_i. For computational tractability, we approximate this function linearly:

$$V^{aug}(w_i, h_i) \approx \gamma_0 + \gamma_h h_i + \gamma_w w_i, \tag{G.7}$$

where γ_0 is a constant and γ_h and γ_w govern the dependence of the outside option on score and income.

The final step is to note that V^{aug} can be normalized to zero without loss of generality. In multinomial choice models, the level of mean utility is not identified; only differences in utility across alternatives are identified. In particular, what is identified is utility relative to the outside option.

To see this formally, suppose that the "true" mean and individual-specific components of utility for private option j are given by

$$\delta_j = \delta_j^{FE} + \widetilde{\delta} + \xi_j, \tag{G.8}$$

$$\mu_{ij} = \widetilde{\beta}^h h_i + \widetilde{\beta}^w w_i + \epsilon_{ij}. \tag{G.9}$$

The parameters $\widetilde{\delta}$, $\widetilde{\beta}^h$, and $\widetilde{\beta}^w$, as well as the parameters defining V^{aug} (γ_0, γ_h, and γ_w) are not separately identified. Only their differences are identified.

Normalizing V^{aug} to zero therefore amounts to a reparameterization. The parameters that are identified and estimated in the model presented in Section 5 are

$$\delta = \widetilde{\delta} - \gamma_0,$$
$$\beta^h = \widetilde{\beta}^h - \gamma_h,$$
$$\beta^w = \widetilde{\beta}^w - \gamma_w,$$

which correspond to utilities measured relative to the augmented outside option.

Interpretation

The equivalence established above is subject to two important caveats. First, it relies a linear approximation to represent the systematic component of the augmented outside option. Second, it assumes that the admission cutoffs $\widetilde{h}_g$ for public degrees are exogenous and fixed.

The latter assumption is innocuous for estimation, but it has implications for the counterfactual exercises in Section 7. In particular, it implies that admission cutoffs at public institutions do not adjust in response to changes in the federal loan program. This is a reasonable approximation insofar as the students who rely on loans to attend private degrees differ systematically from those who typically enroll in public universities. Next, we examine the empirical plausibility of this assumption.

G.2 Differences between public and private higher-education markets in Brazil

This section documents key differences between the markets for public and private colleges in Brazil. Public colleges are tuition-free, highly selective, and widely perceived as more prestigious. Private colleges charge tuition and are typically viewed as lower quality, but offer substantially greater flexibility in program structure. Two dimensions capture these differences most clearly: student academic preparation, measured by ENEM scores, and program shifts.

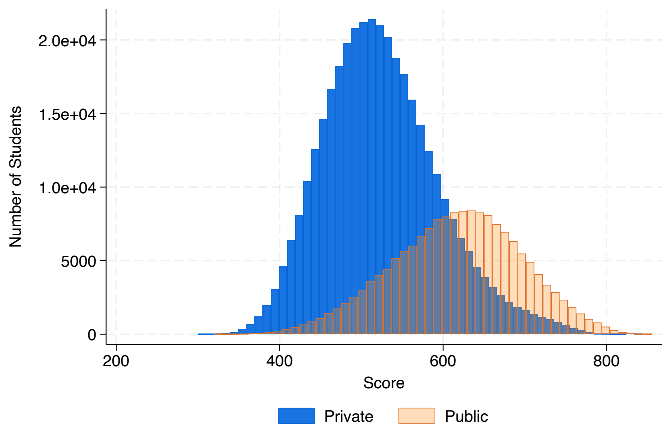

Figure G.1: Entry scores in public and private colleges

Notes: This figure shows the distribution of ENEM scores for students enrolled in public and private colleges. The data include all students who enrolled for the first time in 2014.

Figure G.1 illustrates the differences in academic preparation across sectors. Students enrolled in public colleges have substantially higher entry scores. Although there is some overlap, private-college students in general score below their public-college counterparts. This pattern reflects the selectivity of public universities and implies that, for many students currently enrolled in private institutions, admission to the public sector is not a realistic option.

Table G.1: Distribution of students by shift and sector

Turno	Private (%)	Public (%)
Morning	27.0	15.6
Afternoon	3.2	7.3
Night	63.3	32.7
Full-day	6.5	44.4

Notes: The table reports the percentage of students enrolled in each shift (morning, afternoon, night, and full-day), separately for public and private colleges. The data include all students who enrolled for the first time in 2014.

Table G.1 highlights a second key difference between sectors: program shifts. Private colleges offer substantially more flexible schedules, with a large share of students enrolled in night programs. In contrast, public colleges are concentrated in full-day and daytime shifts. This difference is important for students who work while studying or face other time constraints, for whom public colleges may be impractical even if admission were possible.

For many students enrolled in private colleges, there is no comparable alternative in the public sector. 45.1 % of students in private universities do not have, in the same region where they study, an option in a public college offering the same major and shift. In contrast, only 9.0 % of students in private colleges lack an alternative in another private institution offering the same major and shift.

Taken together, these results indicate that public and private higher-education markets in Brazil are fundamentally different. Public colleges serve a more academically prepared student population and offer less flexible program structures, whereas private colleges enroll students with lower entry scores and provide schedules that better accommodate work and other constraints. At the same time, the overlap in score distributions and the presence of students with feasible cross-sector alternatives imply that the two markets are not fully segmented: some students in each sector could plausibly consider options in the other. Accordingly, the assumption of fixed admission cutoffs should be viewed as an approximation rather than an exact description. We next assess the sensitivity of our results to this approximation.

G.3 Robustness: Alternative sample definition

We now assess the robustness of our results to an alternative treatment of students who enroll in public colleges. Rather than incorporating public-college enrollment into the outside option, we exclude these students from the estimation sample altogether. Apart from this change in sample construction, the estimation procedure follows Section 6, and the counterfactual simulations are conducted as in Section 7.

For brevity, we do not compare the full set of estimated coefficients across specifications. Instead, we focus on the main implications of the model: the effects of loans under alternative targeting schemes on equilibrium prices, and the resulting consequences for enrollment decisions.

Figure G.2 reports the estimated price effects of the loan program under different targeting schemes when students enrolled in public colleges are excluded from the sample. These results are directly comparable to Figure 5 in the main text. We find very similar patterns. In particular, tuition levels are substantially higher under merit-based targeting than under need-based targeting, underscoring the central role of the composition effect in determining equilibrium prices.

We next examine the implications of these price responses for enrollment decisions. In the main text, these effects are shown in Figure 6, which already focuses on enrollment in the private sector. However, students observed enrolling in public colleges in the data may still influence those results since they might choose private colleges under certain counterfactual scenarios. By contrast, Figure G.3 excludes these students entirely.

We find that enrollment responses are very similar across the two approaches. In particular, there remain large differences between demand-only and equilibrium counterfactuals, highlighting the importance of endogenous price responses in shaping the aggregate effects of student loans.

Overall, the results are robust to excluding public-college students from the sample, indicating that our treatment of public colleges as part of the outside option is not driving our findings.

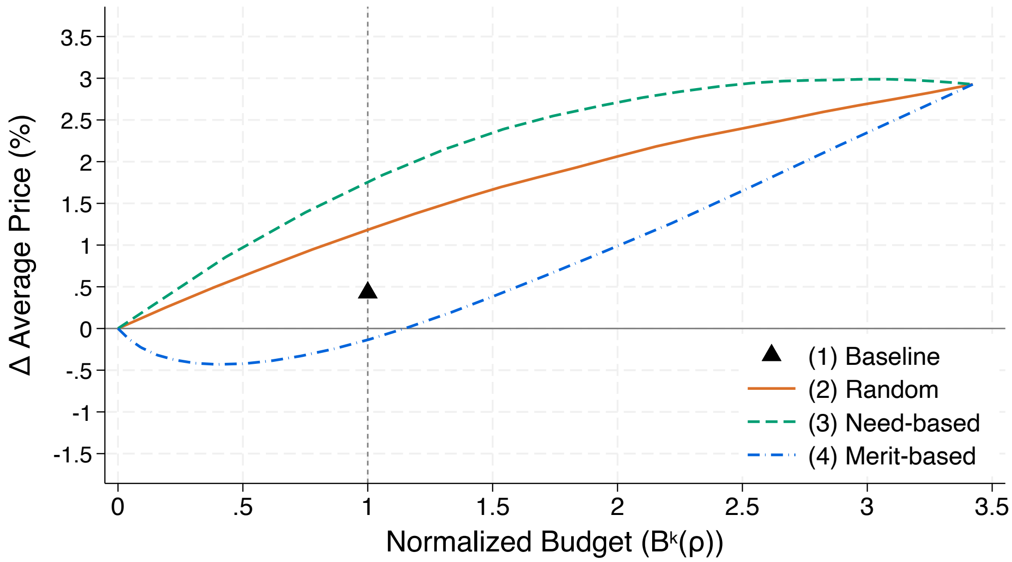

Figure G.2: Impacts of loans on prices: Robustness

Notes: This figure shows the effect of student loans on tuition under alternative targeting schemes. The y-axis reports the percentage change in average tuition relative to a scenario with no loans. Average tuition is computed as the enrollment-weighted average of full and discounted prices across degrees. Results are based on the baseline estimated parameters (black triangle) and counterfactual simulations under three targeting schemes: random allocation (orange solid line), need-based targeting (blue dashed line), and merit-based targeting (green dash-dot line). The vertical dashed line indicates budget-neutral loan programs (i.e., with the same total budget as the baseline). The x-axis reports the budget of the loan program under each counterfactual, normalized by the baseline budget, as defined in Equation (22). The sample excludes students enrolled in public colleges.

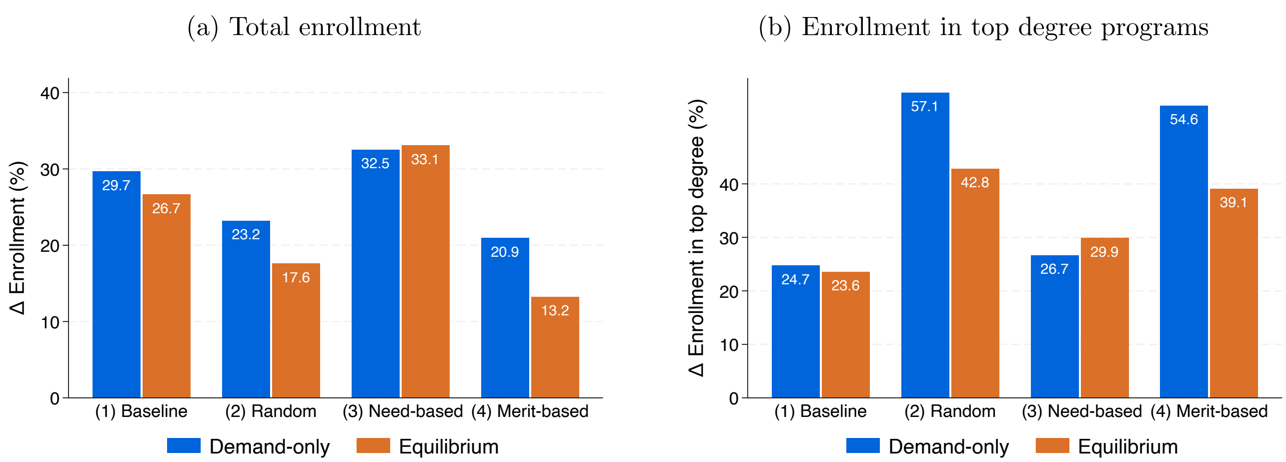

Figure G.3: Equilibrium effects of loans on enrollment: Robustness

Notes: This figure shows the effects of student loans on enrollment under alternative targeting schemes. Outcomes are reported as percentage changes relative to a scenario with no loans. For each targeting scheme, simulations are conducted at a coverage rate ρ that ensures a constant total loan program budget equal to the baseline level (budget-neutral). Panel A presents effects on total enrollment, while Panel B focuses on enrollment in programs in the top 10% of the quality distribution. Degree quality is measured by the average earnings of its graduates. Orange bars allow for supply-side responses by letting tuition adjust endogenously; blue bars hold prices fixed at their observed levels. The sample excludes students enrolled in public colleges.

H Model estimation: details

H.1 Loan cutoffs

In this appendix, we detail how we build micro-moments for model estimation using loan eligibility cutoffs. First, we highlight the relevant features of the setting. Second, we formally define the micro-moments. Third, we describe the underlying variation in the data used to compute these moments.

H.1.1 Setting

As discussed in Section 4.1, since 2015 the number of FIES-funded students per degree has been capped. In most programs, demand for loans exceeded this cap, which led to loan allocation through an iterative deferred acceptance mechanism based on scores on a centralized exam (ENEM), which generated degree-specific cutoffs for loan eligibility.

We impose two additional sample restrictions throughout this appendix. First, since cutoffs did not exist before 2015, we restrict the sample to post-2015 years. Second, the cutoff score for each degree is endogenously determined by the score of the last student assigned a loan. By definition, this student enrolls in the assigned program; otherwise, the loan would have been offered to the next student. Following Chaisemartin and Behaghel (2020), we exclude these students from the analysis.[48] The sample otherwise follows the definition in Section 6.2 and is consistent with the rest of the paper.

H.1.2 Cutoff-based micromoments for model estimation

We incorporate loan eligibility discontinuities in the estimation by matching the log difference in outcomes between students just above and just below the cutoff. We focus on two key outcomes: enrollment and tuition expenditures. To account for heterogeneity, we match these moments separately for each quartile of household income. Formally, we include the following micro-moment in the estimation:

$$\log\left(\mathbb{E}[y_{ijrt} \mid \underline{B}_{ijrt},\, w_i \in \overline{w}_{\upsilon}]\right) - \log\left(\mathbb{E}[y_{ijrt} \mid \overline{B}_{ijrt},\, w_i \in \overline{w}_{\upsilon}]\right), \quad \forall \upsilon \in \{1,2,3,4\}, \tag{H.1}$$

where $\underline{B}_{ijrt} \equiv \{\overline{h}_{jrt} \leq h_i < \overline{h}_{jrt} + bw\}$ and $\overline{B}_{ijrt} \equiv \{\overline{h}_{jrt} - bw < h_i < \overline{h}_{jrt}\}$ indicate whether student i falls within the bandwidth bw above or below the cutoff $\overline{h}_{jrt}$ for degree j in region r and year t. Household income is denoted by w_i, and $\overline{w}_{\upsilon}$ represents the set of incomes in quartile υ. The outcome variable y_{ijrt} is either q_{ijrt}, an indicator for enrollment, or $p_{ijrt}q_{ijrt}$, which captures tuition expenditures.

[48] Although the setting in Chaisemartin and Behaghel (2020) does not map exactly to ours—since they consider a randomized waitlist—it provides the most closely related econometric framework, to the best of our knowledge.

H.1.3 Reduced-form results

We provide reduced-form estimates of the effect of loans by structuring the problem as a standard regression discontinuity (RD) design. Our goal is to provide clear visualization of the data underlying the estimation of Equation (H.1) and test for manipulation. Specifically, we estimate the following local linear regression:

$$y_{ijrt} = \beta_0 + \beta \cdot \text{Above}_{ijrt} + \beta_L \cdot (h_i - \overline{h}_{jrt}) + \beta_H \cdot \text{Above}_{ijrt} \cdot (h_i - \overline{h}_{jrt}) + \epsilon_{ijrt}, \tag{H.2}$$

where $\text{Above}_{ijrt} \equiv \mathbb{1}\{h_i \geq \overline{h}_{jrt}\}$ is an indicator for whether student i's ENEM score h_i is above the loan eligibility cutoff $\overline{h}_{jrt}$ for degree j in region r and year t. The outcome y_{ijrt} is either q_{ijrt} or $p_{ijrt}q_{ijrt}$, as defined in Section H.1.2. Our parameter of interest is β, which captures discontinuous change in the outcome at the cutoff. The coefficient β_0 is a constant term, β_L and β_H control for piecewise linear trends, and ϵ_{ijrt} represents residual variation.

The dataset is structured at student-degree level, which means that each student appears multiple times, once for each degree in their market. Students can only enroll in one degree. We follow Calonico et al. (2014) to determine the optimal bandwidth to estimate Equation (H.2).

Results are shown in Figure H.1. Panel A presents the distribution of the running variable, the relative score $h_i - \overline{h}_{jrt}$. The distribution is smooth, with no evidence of manipulation at the cutoff. We formally test for manipulation using the approach of McCrary (2008), as implemented by Cattaneo et al. (2018), and fail to reject the null of no manipulation (t-statistic = 0.16, p-value = 0.87).

Panel B in Figure H.1 plots the relationship between relative scores and enrollment. There is a clear discontinuity in enrollment probability at the loan eligibility cutoff, which indicates that students are more likely to enroll in a degree program when they qualify for a loan in that degree. The effect is substantial: Estimates from Equation (H.2), reported in the figure, indicate that students are 0.076 percentage points more likely to enroll in a degree when they are eligible for a loan.[49] Panel C plots the relationship between relative scores and tuition expenditures. The results show that tuition expenditures are also higher just above the cutoff compared with just below.

To assess whether students just above and just below the cutoff differ in important ways, we estimate Equation (H.2) using a series of predetermined student characteristics as outcomes. Figure H.2 presents the results for household income and indicators for whether at least one parent has a college degree, whether the student attended a public high school, and whether they reported an intention to apply for a FIES loan when taking the ENEM exam. All variables are balanced across the cutoff.

[49] The enrollment probability for students just below the cutoff is 0.36 percentage points. This number is low, because it reflects the probability of enrolling in one specific degree out of all available options. The average number of degrees in a given market is 88.

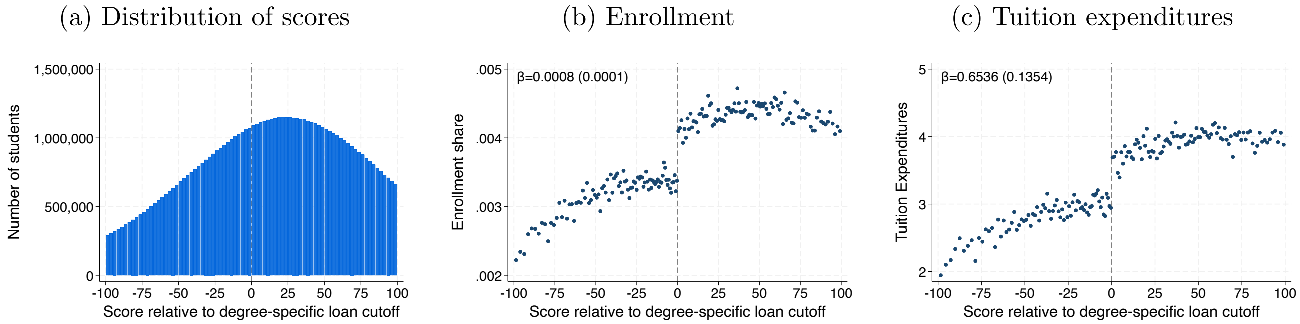

Figure H.1: Loan Eligibility: Regression Discontinuity Estimates

Notes: This figure presents results from a regression discontinuity analysis of loan eligibility effects. The sample is at student-degree level, so each student appears once for each degree in their market. The running variable is the relative score, defined as the difference between a student's ENEM score and the FIES loan eligibility cutoff for that degree. Panel A shows the distribution of relative scores. Panels B and C present binned scatter plots with relative scores on the x-axis. The outcome in Panel B is an indicator for enrollment; in Panel C, tuition expenditures. OLS estimates of Equation (H.2) are reported in both panels. Standard errors, clustered at degree level, are in parentheses.

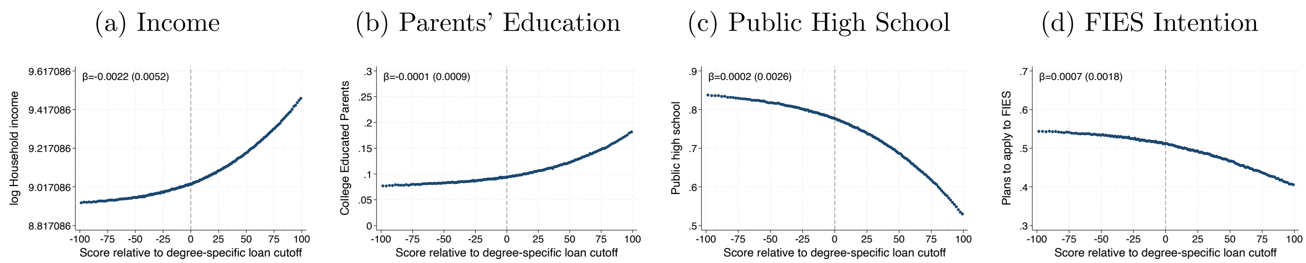

Figure H.2: Loan eligibility: Regression discontinuity balance

Notes: This figure presents regression discontinuity balance tests, assessing whether predetermined student characteristics change discontinuously at the loan eligibility threshold. The sample is at student-degree level, so each student appears once for each degree in their market. The running variable is relative score, defined as the difference between a student's ENEM score and the FIES loan eligibility cutoff for that degree. Each panel reports OLS estimates of Equation (H.2) using a different predetermined characteristic as the outcome variable: (a) log household income (in annual dollars); (b) an indicator for whether at least one parent attended college; (c) an indicator for whether the student attended a public high school; and (d) an indicator for whether the student reported an intention to apply for a FIES loan when taking the ENEM exam. Standard errors, clustered at degree level, are in parentheses.

H.2 Other micro-moments

Other micro-moments used in the model estimation are defined as follows:

$$\text{Average score of enrolled students:} \quad \frac{\mathbb{E}\left[h_i q_{ijrt}\right]}{\mathbb{E}\left[q_{ijrt}\right]}$$

$$\text{Average income of enrolled students:} \quad \frac{\mathbb{E}\left[w_i q_{ijrt}\right]}{\mathbb{E}\left[q_{ijrt}\right]}$$

$$\text{Average income of students in high-price degrees:} \quad \frac{\mathbb{E}\left[w_i q_{ijrt} \mid p^F_{jrt} > \text{median}(p^F_{jrt})\right]}{\mathbb{E}\left[q_{ijrt} \mid p^F_{jrt} > \text{median}(p^F_{jrt})\right]}$$

$$\text{Average income of enrolled students with loans:} \quad \frac{\mathbb{E}\left[w_i q_{ijrt} \mid L_{ijrt} = 1\right]}{\mathbb{E}\left[q_{ijrt} \mid L_{ijrt} = 1\right]}$$

$$\text{Average income of enrolled students with discounts:} \quad \frac{\mathbb{E}\left[w_i q_{ijrt} \mid D_{ijrt} = 1\right]}{\mathbb{E}\left[q_{ijrt} \mid D_{ijrt} = 1\right]}$$

H.3 Estimator

Our estimation approach builds on the methodology of BLP, while extending it to incorporate additional high-dimensional parameters specific to our model. Formally, our estimator minimizes $Q(\theta)$ subject to the constraint whereby five model-predicted quantities match their empirical counterparts:

$$\begin{aligned}
\theta^* = {} & \arg\min_{\theta} Q(\theta) \\
& \text{subject to} \\
& \mathbb{E}\left[q_{ijrt} \mid j, t\right] = \widehat{s}_{jrt}, \quad \forall j, t, \\
& \frac{\mathbb{E}\left[\sum_{j \in \mathcal{J}_{ft}} q_{ijrt} L_{ijrt} \,\middle|\, f, t\right]}{\mathbb{E}\left[\sum_{j \in \mathcal{J}_{ft}} q_{ijrt} \,\middle|\, f, t\right]} = \widehat{s}^L_{ft}, \quad \forall f, t, \\
& \frac{\mathbb{E}\left[\sum_{j \in \mathcal{J}_{ft}} q_{ijrt} D_{ijrt} \,\middle|\, f, t\right]}{\mathbb{E}\left[\sum_{j \in \mathcal{J}_{ft}} q_{ijrt} \,\middle|\, f, t\right]} = \widehat{s}^D_{ft}, \quad \forall f, t, \\
& p_{\iota rt} = c_{\iota rt} + \left(\Delta_{rt}^{-1} \vec{s}_{rt}\right)_{\iota}, \quad \forall \iota \text{ s.t. } d_{\iota rt} = 0, \\
& p_{\iota rt} = c_{\iota rt} + \left(\Delta_{rt}^{-1} \vec{s}_{rt}\right)_{\iota} - \kappa_{\iota rt}, \quad \forall \iota \text{ s.t. } d_{\iota rt} = 1,
\end{aligned} \tag{H.3}$$

where Q is the GMM objective function, defined in Equation (21). The empirical targets $\widehat{s}_{jrt}$, $\widehat{s}^L_{ft}$, and $\widehat{s}^D_{ft}$ denote the observed market share of degree j in year t, the observed share of students with loans among those enrolled in college f in year t, and the observed share of students with discounts among those enrolled in college f in year t, respectively.