Equilibrium Price Responses to Targeted Student Financial Aid*

Nano Barahona† Cauê Dobbin§ Sebastián Otero‡

May 13, 2025

Abstract

We study supply-side responses to student financial aid, focusing on how tuition responds to the targeting of aid. Our framework identifies two mechanisms: a *direct* effect, which raises tuition, and a *composition* effect, which can lower tuition if aid targets price-sensitive students. Leveraging a reform in Brazil's student loan program, we provide descriptive evidence that both mechanisms are quantitatively important. We then estimate an equilibrium model of higher education to quantify the impact of alternative targeting schemes. We find that a loan program with merit-based targeting increases tuition by 3.0%, while need-based targeting raises tuition by only 0.4%. This difference arises because low-income students—targeted in the need-based scheme—are more price sensitive. These price adjustments have a strong impact on enrollment decisions, emphasizing the importance of targeting in student financial aid policy design.

Keywords: Education, student loans, equilibrium effects, private colleges
JEL Codes: H52 , H22, I22, I23, I24

*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 ability 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 student 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

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

others, thereby weakening the composition effect.

We explore the empirical relevance of these forces in Brazil's private higher education market, which 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 the number of new loan contracts, with only about 150,000 loans issued to incoming students in 2017—a fourfold decrease from 2014 levels. Real tuition prices fell by 5.2% following the reform, while the stock prices of the four largest higher education firms dropped 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 drops 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

[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.

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 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 -2.73, 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 -3.91 for students with below-average income and -2.42 for those with above-average income. Moreover, receiving a loan reduces price sensitivity by 23.8% 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 distributed under a fixed budget equal to that of the pre-reform policy. 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, partially offsetting the upward direct effect. As a result, tuition increases by only 0.4% compared to a no-loan scenario. 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. In this case, the composition effect reinforces the direct effect, leading to a tuition increase of 3.0%.

Next, we use our model to quantify the effects of loans on college enrollment. To disentangle the impacts of loans on supply and demand, we simulate two counterfactuals for each targeting

[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.

scheme. The first counterfactual, referred to as *demand-only*, isolates demand-side effects by allowing students to respond to loan availability while holding tuition prices fixed. The second counterfactual, referred to as *equilibrium*, incorporates supply-side responses by allowing tuition prices to adjust based on changes in demand.

Under need-based targeting, enrollment is nearly identical in the demand-only and equilibrium counterfactuals, with gains of approximately 30% relative to the no-loan benchmark. This similarity reflects the modest tuition responses under this scheme. In contrast, the effects of merit-based targeting are weaker in the demand-only counterfactual—around 11.5%—because high-performing students are more likely to enroll even without loans and tend to choose more expensive programs, which reduces the number of students that can be funded under a fixed budget. In the equilibrium counterfactual, larger tuition increases under merit-based targeting offset most of the loan-induced demand gains, yielding a smaller enrollment increase of 2.9%. Similar patterns hold when restricting attention to high-quality degrees. Taken together, the results underscore the importance of targeting design in determining the effectiveness of student loan policies.

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 several strands of the literature. First, our conceptual framework builds on the literature on tax and subsidy incidence under imperfect competition. A well-established body of theoretical work provides the foundation for understanding the effects of uniform tax policies (Delipalla and Keen, 1992; Anderson et al., 2001; Weyl and Fabinger, 2013; Miravete et al., 2018; Kroft et al., 2024a,b). A key distinction between taxes and subsidies is that subsidies often target specific groups, making the composition of subsidy recipients a crucial determinant of equilibrium outcomes. Building on this idea, recent studies have examined the price effects of targeted subsidies in the context of health insurance markets. Mahoney and Weyl (2017) theoretically analyze the incidence of subsidies that vary with consumers' marginal costs. Polyakova and Ryan (2022) introduce the concept of "demographic externality"—akin to our composition effect—and estimate an equilibrium model to assess how targeted subsidies influence health insurance prices.

Second, we contribute to the literature on the relationship between financial aid and tuition fees. Many studies use difference-in-differences strategies similar to ours and estimate the impact of financial aid on tuition by comparing institutions with varying exposure to aid reforms (Long, 2004; Singell and Stone, 2007; Lucca et al., 2019; Baird et al., 2022). Our empirical strategy is most similar to that of Black et al. (2023), who measure exposure to the Grad PLUS reform in the United States based on how much baseline prices exceeded the pre-reform federal loan cap. De Mello and Duarte (2020) also examine the Brazilian context but use an alternative identification strategy that does not rely on the 2015 FIES reform. These studies generally find that financial aid increases tuition, consistent with the Bennett Hypothesis and our direct effect mechanism. We identify an alternative channel through which financial aid affects tuition—the composition effect—offering a more nuanced view of the Bennett Hypothesis.

Third, we build on the literature on the effects of loans on the demand for higher education. Several papers exploit discontinuities in loan availability across score thresholds as quasi-experiments to show that loans impact college enrollment on both the extensive (Solís, 2017; Londono-Velez and Rodriguez, 2020; Aguirre, 2021) and the intensive margin (Bucarey et al., 2020; Aguirre, 2021; Hampole, 2024). In addition, Avery and Hoxby (2004) and Joensen and Mattana (2021) examine how loans affect students' choices using discrete-choice frameworks. We combine these two approaches by incorporating loans in a discrete-choice model and estimating the model using discontinuities in loan availability.

Finally, to quantify the role of aid targeting and the composition effect on equilibrium outcomes, we build on the literature that examines the impact of student aid in education markets using Bertrand-Nash oligopoly models. Several studies, including Bucarey (2018); Neilson (2021); Allende (2021); Bodere (2023); and Sanchez (2023), focus on primary or secondary education and predominantly analyze voucher-based aid, where the underlying economic incentives differ substantially from those associated with student loans. Two recent studies extend this literature to higher education. Cordeiro and Cox (2023) investigate the same setting as our study, focusing on quality responses, and provide insights into the long-term effects of financial aid. Armona and Cao (2024) examine the design of financial aid in the United States and emphasize the importance of targeting colleges by quality, while holding the targeting of different students fixed.

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 $[0, 2]$.[4]

The monopolist maximizes profits by solving

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

where Q denotes the quantity purchased. We assume that the subsidy schedule $\mathcal{T}$ is such that $Q\big(p\big|\mathcal{T}\big)$ is twice-differentiable and decreasing in p, and that the profit function $Q\big(p\big|\mathcal{T}\big) \cdot p$ 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:

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

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

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

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

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

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

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

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

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

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

[4] 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.

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 ($\widetilde{p}_a > \overline{p}$).

Figure 1: Flat and targeted subsidies

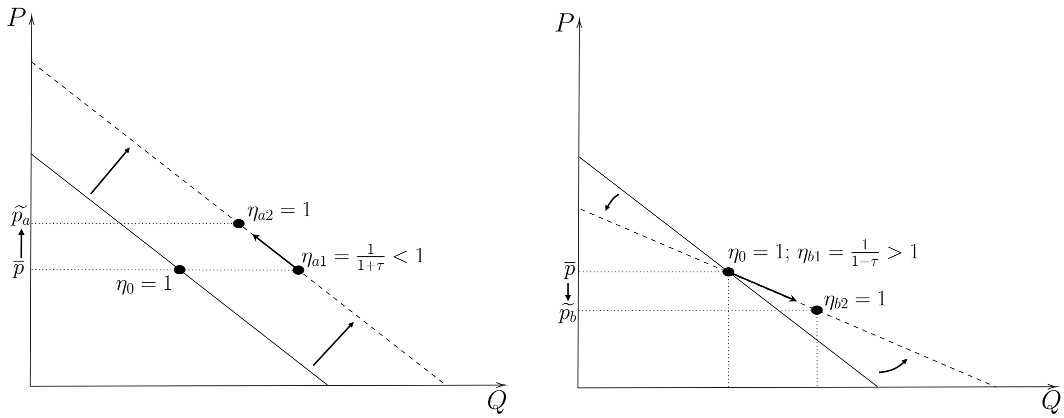

(a) Flat Subsidy (b) Targeted Subsidy

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 $\widetilde{p}_b$.

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.[5] This rotates the demand curve around the benchmark price $\overline{p}$, as shown in Panel (b) of Figure 1. The new demand curve is flatter, increasing the price elasticity at $\overline{p}$ to

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

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

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 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 formally 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 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).[6]

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

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

[6]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 $\Omega > 0$.

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 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, 76% 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 of 1,161 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, it underwent major reforms in 2010 that significantly expanded access. New loan contracts surged from fewer than 20,000 in 2009 to over 700,000 in 2014. For FIES-supported students, the government directly pays colleges, covering up to 100% of tuition costs, 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 approximately one-sixth of private colleges' revenues, comparable to student loans in the United States (Lee and Looney, 2019). In 2014, the median program enrolled roughly 25% of its students using FIES loans, with some programs relying on FIES for 100% 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 coverage rates were means-tested, with both criteria depending on family income. Under the rules in place until 2014, students were eligible if their total family income was below 20 times the federal minimum wage. Loans covered between 50% and 100% of tuition costs, depending on the ratio of family income to tuition fees.[7] In 2015, fiscal constraints led 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.[8] 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. 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 1 week, and 50% are removed 3 three weeks.[9] It is also important to note that FIES regulations explicitly prohibit colleges from denying tuition discounts to students receiving government loans.

[7]See Appendix C for details.

[8]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.

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

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 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 private firms. Appendix D.1 details how we combine these sources to compute full and discounted tuition prices for each program.[10] Overall, we recover year-specific prices for 95% of degree-years, covering 98% of total enrollment.

Degree Quality. We measure degree quality using earnings value added, estimated via a selection-on-observables approach. Specifically, we attribute earnings differences among observationally similar students attending different programs to the value added by those degrees. The earnings data come from RAIS, a matched employer-employee administrative dataset that covers the universe of formal employment in Brazil. The value added estimation incorporates a rich set of individual-level covariates, including detailed test scores and family income. Appendix D.2 provides further details on the methodology. We also consider two alternative measures of degree quality: the average income of a program's graduates and the average entrance score of its incoming students. All quality metrics are constructed using the 2010 incoming cohort and are held fixed throughout the analysis.

[10] Three sources—QueroBolsa, FNDE, and PROUNI—provide discount data, though without individual-level identifiers. We compute the average observed discount for each degree program and use this as its discounted tuition price.

3.3 Sample

We define a *degree* as a combination of major, institution, and shift.[11] 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 National Bureau of 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. 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.[12] After these restrictions, our final sample consists of over 20 million ENEM takers and 13,567 unique degree programs offered by 695 colleges, covering 88% of ENEM takers nationwide.

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 lowered the maximum per capita family income threshold from 20 to 2.5 times the federal minimum wage and introduced a minimum ENEM score of 450.[13] 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.[14]

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

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

[12] 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.

[13] 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.

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

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—a fourfold decline. Figure 2 presents the key aggregate trends surrounding the reform. Panel (a) shows this steep drop in loan volume alongside a decline in 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 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 prices in subsequent months suggest that the reform triggered a structural reassessment of the sector's financial outlook.

Figure 2: The 2015 FIES Reform: Aggregate Trends

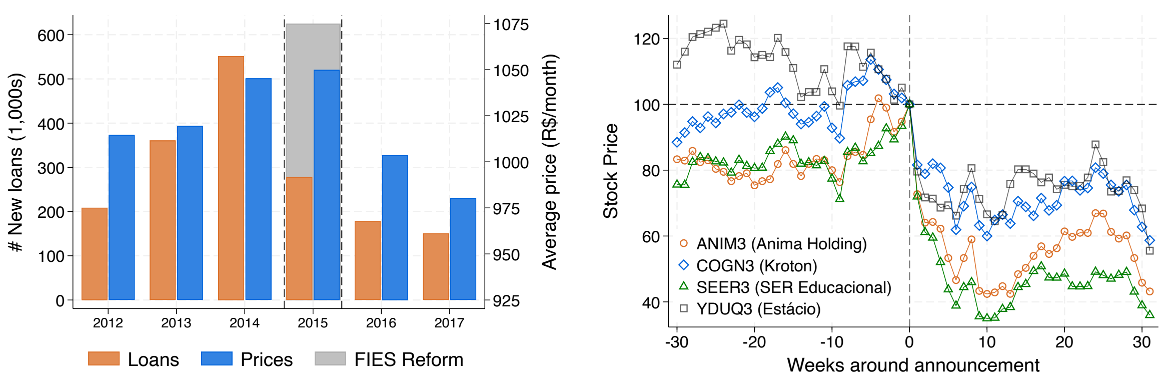

(a) Number of loans and tuition costs

(b) Stock prices

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 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 before the reform. 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 enrollment in the same year, and $\bar{N}_j$ the degree-specific loan cap imposed by the policy.[15] 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,

[15] 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.

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.

Figure 3: Differential Exposure to the 2015 FIES Reform

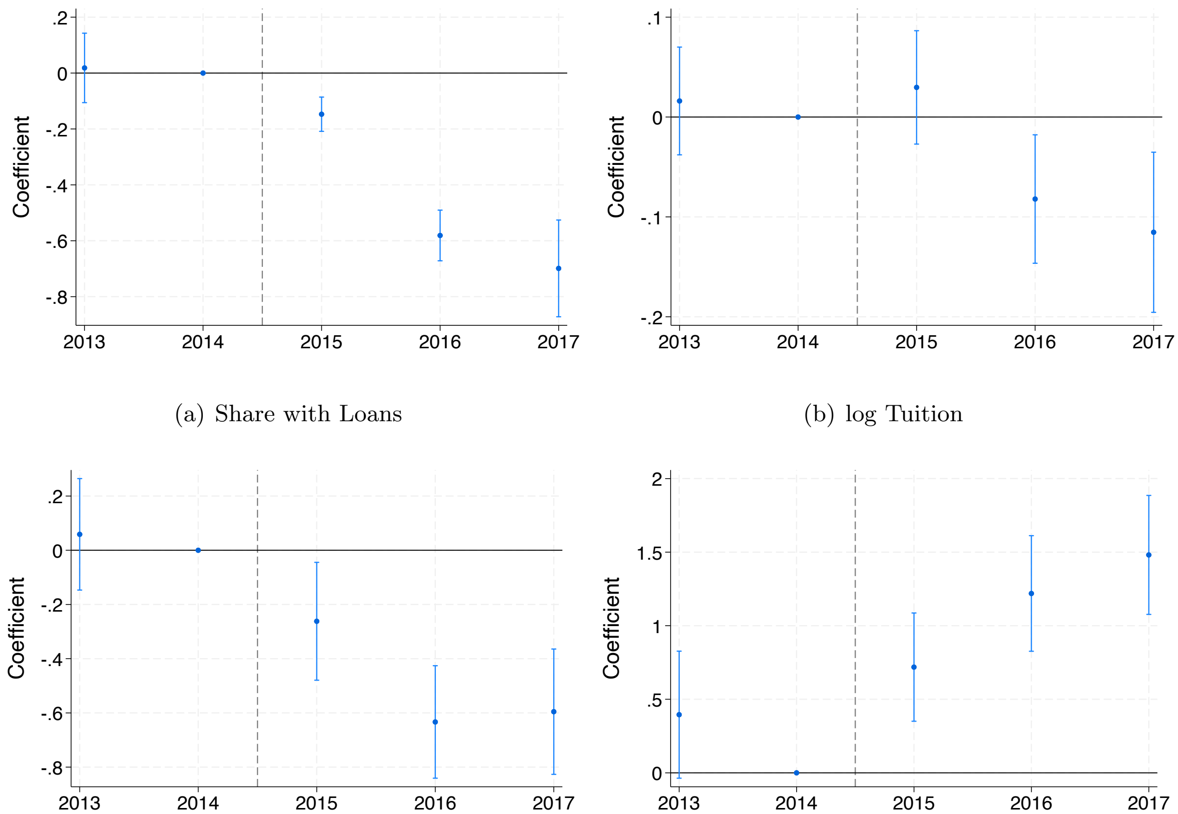

(a) Share with Loans

(b) log Tuition

(c) log Enrollment

(d) log Enrollment w/o Loan

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 number of incoming students; in Panel (c), the log monthly tuition fee; and, Panel (d), 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 parallel with the FIES reform, the federal government implemented several fiscal adjustments in 2015, leading to broader economic shifts. While the FIES reform was the most direct policy change affecting higher education, macroeconomic conditions could also have influenced the estimates shown in Figure 3. For example, degree programs more exposed to the reform might be tied to fields of study and occupations that were disproportionately impacted by other policy changes. To assess whether our results are driven by such concurrent factors, we augment Equation (7) with region-major-year fixed effects, which control for region-specific trends across fields of study. The results—shown in Figure A.3—are consistent with our baseline estimates. This robustness suggests that the observed effects are not confounded by other simultaneous policy shifts, and that degrees more exposed to the FIES reform were not disproportionately affected by broader economic changes.

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 in 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, $\eta_\rho^L - \eta$.[16]

This measure provides meaningful insight into differences in price elasticity under two key assumptions. First, income must serve as a valid proxy for price sensitivity.[17] 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 (7) to allow for heterogeneity in tuition responses driven by differences in $\Delta\eta_j$. Specif-

[16] Figure A.4 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.

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

ically, we estimate:

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

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-policy years (2013 and 2014) and two post-policy years (2016 and 2017). As a robustness check, 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.[18]

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.

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 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

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

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 Prices 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	13,447	11,695	13,447	11,695	11,695
Degree FE	Yes	Yes	Yes	Yes	Yes
Region-Year FE	Yes	Yes	Yes	Yes	Yes
Major-Region-Year FE	No	Yes	No	Yes	Yes
Δ Inc-Year FE	No	No	No	No	Yes

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.

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

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

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

$$\delta_{jt} = \delta_j + \delta_{r(j)t} + \xi_{jt}, \tag{10}$$

where δ_j captures degree fixed effects and absorbs time-invariant characteristics of the program. The term $\delta_{r(j)t}$ accounts for region- and time-specific market conditions that influence demand. Lastly, ξ_{jt} represents unobserved, time-varying demand shocks that affect the degree's overall attractiveness.

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

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

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 $U_{i0t} = \epsilon_{i0t}$.[19] Finally, ϵ_{ijt} is a student-specific demand shock, which we assume follows a generalized extreme value distribution and allows for unobserved heterogeneity in individual preferences.

The third component, $\alpha^0_{ij} + \alpha^1_{ij} p_{ijt}$, captures the financial costs associated with enrolling in a given program. We allow this component to depend on tuition through α^1_{ij}, which reflects the student's price sensitivity. The term α^0_{ij} accounts for financing-related factors that influence enrollment but are independent of tuition levels. For instance, student who face severe liquidity constraints may be unable to enroll in any program, even the least expensive ones. This means

[19] 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.

that marginal tuition changes would have little impact on their enrollment decisions.

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

$$p_{ijt} = p_{jt}(D_{ij}) = \begin{cases} p^F_{jt}, & \text{if } D_{ij} = 0, \\ p^D_{jt}, & \text{if } D_{ij} = 1, \end{cases}$$

where p^F_{jt} and p^D_{jt} are the full and discounted tuition rates for degree j in year t, and D_{ij} indicates whether student i has a discount in degree j.

We allow student loans to influence the financial-cost component of utility through two channels: by altering price sensitivity, α^1_{ij}, and by affecting the baseline utility cost of education financing, α^0_{ij}. 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_{ij} and α^1_{ij}, to flexibly depend on the student's loan status and student characteristics. Specifically, we define:

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

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

where L_{ij} indicates whether student i would use a loan if they enroll in degree j. 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.

5.2 Financial aid allocation

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

$$L_{ij} = \mathbb{1}\{\rho^L_w \cdot w_i + \rho^L_{f(j)t} \geq \vartheta^L_{ij}\} \cdot \mathbb{1}\{h_i \geq \overline{H}_{jt}\} \tag{14}$$

$$D_{ij} = \mathbb{1}\{\rho^D_w \cdot w_i + \rho^D_{f(j)t} \geq \vartheta^D_{ij}\}, \tag{15}$$

where ϑ^L_{ij} and ϑ^D_{ij} are idiosyncratic error terms that follow a logistic distribution and capture unobserved individual disutility in searching for a loan or discount for a given degree program. Students' propensity to receive loans and discounts also depends on their income through the parameters ρ^L_w and ρ^D_w, respectively, which reflects the fact that both eligibility and the likelihood of seeking financial aid might depend on household income. We also assume that the government sets degree-specific loan score thresholds, as defined by $\overline{H}_{jt}$.

The parameters $\rho^L_{f(j)t}$ and $\rho^D_{f(j)t}$ capture systematic unexplained variation in the share of students with loans and discounts in a given college f. This is motivated by the fact that we observe substantial variation in financial aid usage across degree programs, even after controlling for student characteristics. This variation is attributed to factors not explicitly modeled, such as the social stigma associated with such aid (Moffitt, 1983; Finkelstein and Notowidigdo, 2019) and differential access to information about financial aid (Castleman and Page, 2015).[20]

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

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

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

$$s_{jt}(l,d) = \mathbb{E}\left[\mathbb{1}\big\{L_{ij} = l\big\} \cdot \mathbb{1}\big\{D_{ij} = d\big\} \cdot q_{ij} \big| i \in \mathcal{I}_t\right], \tag{16}$$

where $\mathbb{E}$ denotes the expectation operator over preference shocks, ϵ_{ij}, and financial aid shocks, $\vartheta^L_{ij}, \vartheta^D_{ij}$. The total market share for degree j in market t is given by

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

5.3 Supply

Each college f offers a set of degrees $\mathcal{J}_f$ and each degree is associated with a time-varying marginal cost c_{jt}. Colleges choose both full and discounted prices for each of their degrees to maximize expected profits:

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

where $s^F_{jt} \equiv \sum_{l\in\{0,1\}} s_{jt}(l,0)$ and $s^D_{jt} \equiv \sum_{l\in\{0,1\}} s_{jt}(l,1)$ represent the enrollment of students paying full and discounted prices, respectively. We decompose marginal cost as follows:

$$c_{jt} = c_j + c_{r(j)t} + \omega_{jt}, \tag{18}$$

where c_j are degree fixed effects, $c_{r(j)t}$ are market fixed effects, and ω_{jt} is a degree-year-specific cost shock.

[20]Stigma is 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 also matters: as discussed in Section 3, many students rely on online platforms like QueroBolsa to learn about aid. Colleges with lower name recognition tend to attract more students through such platforms, leading to higher aid usage. Consistent with this, 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%.

The difference between full and discounted prices is partly driven by differences in the demand curves of students paying full versus discounted prices. For instance, if discounted price demand is more elastic, discounted price markups will be lower than full price ones, leading to larger discounts. The structural error term κ_{jt} captures residual variation in the difference between full and discounted prices that cannot be explained by differences in demand. It accounts for unmodeled factors, such as reputational concerns, that may influence a college's discount policies. For example, institutions might offer larger discounts to cultivate a socially responsible image or, conversely, restrict discounts to signal exclusivity.

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.

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 number of students enrolled in each degree in market rt, with and without discounts:

$$\vec{s}_{rt} \equiv (s^D_{1rt}, ..., s^D_{\mathcal{J}_{rt}rt}, s^F_{1rt}, ..., s^F_{\mathcal{J}_{rt}rt})'.$$

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_{jt}'s of degrees in market rt, and are ordered as $\vec{s}_{rt}$. The elements of these vectors are indexed by ι, and each element $d_{\iota rt}$ of $\vec{d}_{rt}$ is a dummy indicating whether element ι refers to the discounted version of a degree.

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

$$p_{\iota rt} = c_{\iota rt} + \Delta^{-1}_{rt[\iota,\cdot]}\vec{s}_{rt} - \kappa_{\iota rt}d_{\iota rt}, \tag{19}$$

where $\Delta^{-1}_{rt[\iota,\cdot]}$ is the ιth row of the inverse of Δ_{rt}. The (ι, υ) element of the matrix Δ_{rt} is defined as

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

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, $\Delta^{-1}_{rt[\iota,\cdot]}\vec{s}_{rt}$, 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 typical framework and represents a structural error that captures unexplained variation in the difference between full and discounted prices, and reflects factors such as reputation or branding strategies.

5.4 Discussion

The role of price discrimination. The stylized model in Section 2.2 assumes uniform pricing, whereas the full model incorporates price discrimination, which alters the mechanics of the composition effect. In a setting without price discrimination, the composition effect hinges on the difference between the overall price elasticity of demand and the elasticity specific to loan holders. When colleges practice price discrimination, however, the same degree with and without discounts effectively represents two different products from the perspective of pricing decisions, as each is associated with a distinct demand curve. This differentiation can weaken the composition effect.

To illustrate, consider a scenario where loan holders have much lower incomes than the average student, leading to a more elastic demand curve. Absent price discrimination, this would imply a strong composition effect. However, if all loan holders receive tuition discounts and these discounts are also targeted at other low-income students, the composition effect diminishes. In this case, the elasticity of loan holders becomes comparable to that of other discounted students, highlighting how the relationship between discount allocation and loan assignment shapes the composition effect.

Limitations of the model. Even though our model is flexible enough to capture several key features of this market, it has three important limitations. First, colleges take demand as given. This assumption simplifies the analysis but overlooks the potential for colleges to change demand through investments in educational quality, infrastructure, or marketing.

Second, the model assumes that colleges do not control which students receive loans and discounts. For government loans, this assumption reflects legal constraints that prevent colleges from rejecting students based on loan status. For discounts, it reflects the fact that students self-select through financial aid platforms such ah QueroBolsa, which leaves colleges with limited control over which students receive discounts. However, colleges may still influence financial aid uptake by, for example, establishing offices to assist students in navigating the loan application process (Bettinger et al., 2012). Also, colleges can increase the duration for which their discounts appear on financial aid aggregator websites, which expands the share of students who find and use these discounts. For simplicity, we abstract from these issues and assume that colleges take financial aid allocation as given.

Third, our latent variable model—presented in Equations (14) and (15)—determines the *allocation* of financial aid, which in practice depends on both *eligibility* and *take-up*. We do not model eligibility and take-up separately. This limitation may affect our counterfactual simulation results, as it precludes a connection between tuition prices and aid usage. For instance, in reality, a student might avoid taking up a government loan due to stigma when prices are low but may decide to take it when prices increase—a behavior not captured in our simulations.

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_{jt}, 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:

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

where $\text{post}_t \equiv \mathbb{1}\{t > 2015\}$ is an indicator for post-reform years.

We use this instrument to estimate price elasticities, assuming that Z_{jt} is uncorrelated with unobserved demand and supply shocks, conditional on region-year and degree fixed effects. Formally, we impose the moment conditions $\mathbb{E}[Z_{jt} \cdot \xi_{jt}] = 0$ and $\mathbb{E}[Z_{jt} \cdot \omega_{jt}] = 0$. 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_{jt}.[21] 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 the eligibility score threshold. As described in Section 3.1,

[21] 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.

students apply for loans through a centralized admission system that produces loan-eligibility score cutoffs for each degree in the system (see Appendix C for a detailed description of the allocation system).

Section 5.4 describes 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 captures the log difference in enrollment rates for students just above and just below the eligibility threshold; the second captures the corresponding 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. Appendix E.1.2 formally defines these moments, and Figure E.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.

First, 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. Appendix E.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 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 Appendix E.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 (14) and (15); 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_{jt}, and marginal costs, c_{jt}. These parameters are recovered as follows. First, δ_{jt} 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_{jt} is recovered by inverting the firms' first-order conditions (FOCs). From δ_{jt} and c_{jt} one can recover the model residuals, ξ_{jt} and ω_{jt}, 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_{jt}; loan usage propensities, $\rho^L_{f(j)t}$; discount usage propensities, $\rho^D_{f(j)t}$; marginal costs, c_{jt}; and the structural errors that govern discount magnitudes, κ_{jt}. 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_{jt}, while matching the share of students with loans and discounts in each college provides $\rho^L_{f(j)t}$ and $\rho^D_{f(j)t}$, respectivelly. We recover the remaining two sets of parameters, c_{jt} and κ_{jt}, by inverting the firms' FOCs, leveraging the fact that firms optimize two choice variables (p^F_{jt} and p^D_{jt}). 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 Appendix E.3, we formally define the estimator.

6.3 Model estimates

We report the estimated preference and targeting parameters in Table A.2 and the distributions of price elasticities, marginal costs, and markups in Figure A.5. The median price elasticity is -2.73, 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 –3.91, whereas for those with above-average income it is –2.42. Also, taking a loan reduces individual price sensitivity, α_i, by 23.8% on average.

The median tuition price and marginal cost are \$4,152 and \$2,620, respectively. The implied median markup is 0.37, which indicates that colleges exert substantial market power. The median structural error that governs unexplained differences between full and discounted prices, κ_t, is \$788, 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.[22]

Our estimates of the financial aid targeting parameters, presented in Panel B of Table A.2, suggest that loans are effectively directed toward low-income students, while tuition discounts exhibit significantly weaker targeting. A one standard deviation increase in family income decreases the probability of receiving a loan by 7.1 percentage points. In contrast, the same increase in income reduces the likelihood of receiving a tuition discount by only 0.7 percentage point. These findings are consistent with descriptive patterns in the data: Loan receipt is substantially more correlated with family income than tuition discounts, as shown in Figure A.6.

Interpreting the estimates: Direct and composition effects. We now interpret 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 there is a large difference in price elasticity between students with and without loans. 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 (14), we first compute each student's probability of receiving a loan and organize students into groups based on these probabilities.[23] 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. Comparing these scenarios allows us to isolate the direct effect of loan receipt on demand, holding fixed the composition of students within each group. Throughout, tuition prices are held fixed at their

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

[23] 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.

observed levels in the data.

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 difference between loan and no-loan scenarios. This suggests that the composition effect plays a quantitatively important role.

We next assess the role of price discrimination. As discussed in Section 5.4, 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 effectively zero (-0.0001) 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.

Figure 4: Price elasticity and loan targeting

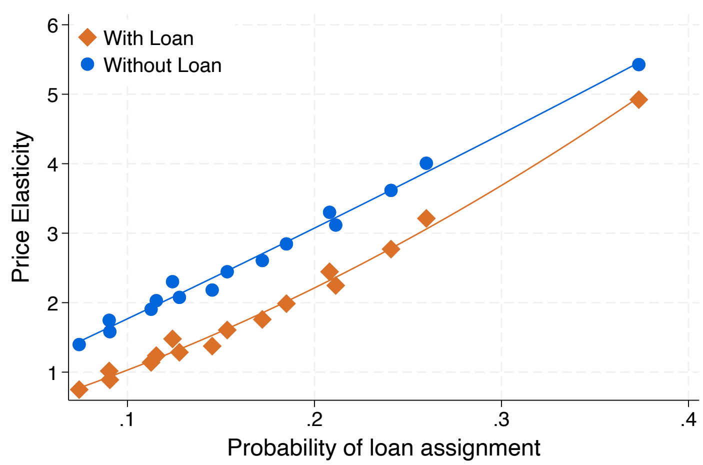

Notes: This figure illustrates the relationship between price sensitivity and loan targeting, as well as the effect of loans on price sensitivity. The x-axis reports the probability of receiving a loan, computed using Equation (14) and the estimated parameters from Section 6.3. Students are grouped based 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.

Model fit and validation. We evaluate the model's fit by comparing its predictions with observed data. Table E.1 shows that the targeted micro-moments align closely with empirical moments. To assess the model's out-of-sample predictive performance, we test its ability to replicate the observed price effects of the 2015 loan reform using pre-reform data. As shown in Figure E.3, the model performs well in predicting post-reform price trends, which reinforces its reliability for counterfactual policy simulations. Further details are provided in Appendix E.4.

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.

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 (15), with firm-specific discount allocation parameters, ρ_f^D, also fixed at their 2014 levels.[24]

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.

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 (14) and assign loans based on the estimated parameters $\hat{\rho}_w^L$ and $\hat{\rho}_f^L$.[25]

[24] 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.

[25] 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

(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. Eligibility 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 access.

(4) *Merit-based:* Loans are allocated based on academic performance h_i. Eligibility 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 eligibility.

In loan allocation schemes (2)-(4), the parameter ρ represents the proportion of students eligible for a loan. As discussed in Section 5.4, we do not model eligibility and take-up separately; therefore, all eligible students take up a loan if they enroll.

For each scheme k and eligibility share 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.[26] Specifically, the budget function $B^k(\rho)$ is given by

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

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 eligibility share 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.[27] 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.

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

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.

[26] 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.

[27] 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.

determined by Equation (19), while market shares are governed by Equation (16). 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 share of students assigned a loan, ρ. Markers along each line denote outcomes under the budget-neutral version of each scheme, with colors indicating the corresponding targeting scheme.

We begin by examining the price effects under the baseline loan allocation, indicated by the black triangle in Figure 5. Relative to the no-loan scenario, this allocation increases equilibrium tuition by 1.92%. In this counterfactual, loans are assigned according to their empirical distribution in the data, where loans are imperfectly targeted to low-income students (see Figure A.6). The net price effect combines an upward force from the direct effect with a downward force from the composition effect, because loan recipients tend to be more price elastic (see Section 6.3, Figure 4).

Figure 5: Impacts of loans on prices

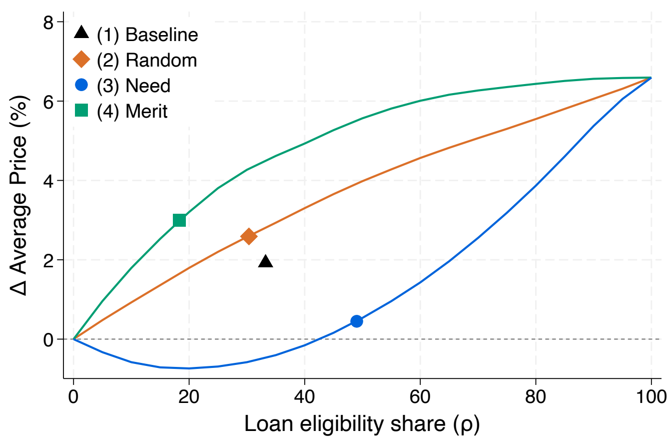

Notes: This figure shows the effect of student loans on tuition under alternative targeting schemes. The y-axis reports the percentage change in average tuition relative to a scenario with no loans. Average tuition is calculated as the enrollment-weighted average of full and discounted prices across degrees. Results are based on the baseline estimated parameters (black triangle) and counterfactual simulations under three alternative targeting schemes: random allocation (orange line), need-based targeting (blue line), and merit-based targeting (green line). Markers (square, diamond, and dot) indicate loan programs with the same total budget as the baseline (budget-neutral). The x-axis represents the loan availability share, ρ. In the baseline scheme, ρ denotes the share of student-degree pairs receiving loans. In the alternative schemes, ρ represents the share of students receiving loans, with each loan covering any degree in which the student enrolls.

Next, we examine the effects of random targeting, in which loans are allocated independent 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 2.58% relative to the no-loan scenario, as represented by the orange diamond in Figure 5. Since only the direct effect operates, tuition increases monotonically with the loan eligibility share ρ, as illustrated by the solid orange line.

We then explore two widely used financial aid allocation strategies: need-based and merit-based loan targeting. Unlike random targeting, these approaches introduce a combination of direct and composition effects, which influence tuition through different channels. Under need-based targeting, loans are directed toward low-income students, which produces a negative composition effect that exerts downward pressure on tuition. As shown by the blue line in Figure 5, this effect dominates when up to 42% of students receive a loan, and results in equilibrium tuition below the no-loan benchmark. However, as the share of eligible students expands and the income of the marginal loan recipient rises, the composition effect weakens. Consequently, beyond the 42% threshold the direct effect dominates, which causes tuition to increase.

By contrast, merit-based targeting allocates loans to students with the highest test scores—a group that, as shown in Panel B of Figure A.6, is disproportionately drawn from higher-income backgrounds. These students tend to be less sensitive to tuition prices, which means the composition 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 loan eligibility share expands.

Overall, we find substantial heterogeneity in tuition responses across targeting schemes, driven by both direct and composition effects. Under a budget-neutral constraint, merit-based targeting raises tuition by 3%, while need-based targeting yields only a 0.45% increase. These outcomes are indicated by the green square and blue circle in Figure 5, respectively. The magnitude of this difference suggests that tuition responses can meaningfully amplify—or dampen—the broader effects of loan programs on college enrollment, which we examine in the next section.

7.4 Equilibrium impacts of loans on enrollment

We now turn to the effects of loan targeting on college enrollment, focusing on two key outcomes: total enrollment and enrollment in high-quality degree programs. We define "top" degree programs as those in the top 10% of the quality distribution, as measured by degree value added.[28] All results are expressed as percentage changes relative to enrollment levels 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. Additional results that explore variation in loan eligibility share are reported in Figure A.7.

[28] Appendix D.2 outlines the methodology used to construct the degree quality measure. Figure A.8 reports robustness results using alternative definitions.

In the demand-only counterfactual, need-based targeting generates the largest increase in total enrollment (32.5%), followed by the baseline scenario (24.6%), random allocation (18.2%), and merit-based targeting (11.5%). 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 these 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 1.9 times more likely to enroll without a loan than those with below-median scores, and enroll in degrees that are 13% more expensive, on average. As a consequence, the merit-based scheme—which targets high-score students—generates a weaker impact on enrollment.

While overall enrollment responses vary substantially across schemes, enrollment in top programs rises by roughly 30% under all scenarios. This convergence 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 impact of need-based targeting on top program enrollment.[29]

In the equilibrium counterfactual, price adjustments amplify the differences in overall enrollment across targeting schemes. Under all schemes, enrollment falls relative to the no-loan scenario; this reflects the fact that, under a budget-neutral constraint, tuition rises in all schemes (as shown in Figure 5). The modest price increase under need-based targeting reduces the enrollment gain only slightly, from 32.5% to 31.1%. In contrast, the sharper tuition increases under merit-based targeting offset much of the initial enrollment gains, with the increase falling from 11.5% to just 2.9%. These patterns highlight how endogenous price responses can critically shape the aggregate impact of loan programs.

The effects on enrollment in high-quality programs vary substantially across targeting schemes once tuition adjustments are taken into account. Under need-based targeting, price effects slightly amplify enrollment in top programs, which increases the impact of loans on top enrollment from 31.7% to 33.0%. This amplification occurs because tuition declines for top programs in this scenario. Figure A.9 illustrates these program-specific price changes. By contrast, under merit-based targeting, tuition increases are especially pronounced for top programs and reduce the enrollment gain from 31.8% to 15.4%.

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 poten-

[29]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.

Figure 6: Equilibrium effects of loans on enrollment

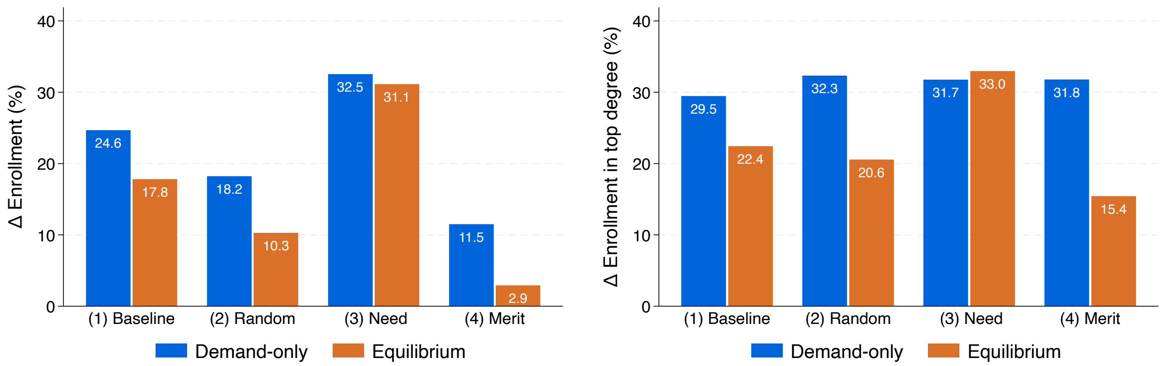

(a) Total enrollment

(b) Enrollment in top degree programs

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 an eligibility share 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 value-added distribution. Orange bars allow for supply-side responses by letting tuition adjust endogenously; blue bars hold prices fixed at their observed levels.

tially 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 article, 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 increase tuition by only 0.4%, 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 3.0%. Moreover, enrollment gains relative to the no-loan scenario range from 31.1% under need-based targeting to just 2.9% 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). 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.

Angrist, Joshua D., Peter D. Hull, Parag A. Pathak, and Christopher R. Walters, "Leveraging Lotteries for School Value-Added: Testing and Estimation*," *The Quarterly Journal of Economics*, May 2017, *132* (2), 871–919.

Angrist, Joshua, Peter Hull, and Christopher Walters, "Methods for Measuring School Effectiveness," in "Handbook of the Economics of Education," Vol. 7, Elsevier, 2023, pp. 1–60.

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.

Bettinger, Eric P., Bridget Terry Long, Philip Oreopoulos, and Lisa Sanbonmatsu, "The Role of Application Assistance and Information in College Decisions: Results from the H&R Block Fafsa Experiment*," *The Quarterly Journal of Economics*, August 2012, *127* (3), 1205–1242.

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.

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

Cordeiro, Fernando and Alvaro Cox, "College Quality and Tuition Subsidies in Equilibrium," *Working Paper*, 2023.

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.

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.

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.

Lee, Vivien and Adam Looney, "Understanding the 90/10 Rule," *Economic Studies at Brookings*, 2019.

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.

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.

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.

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

Rothstein, Jesse, "Teacher Quality in Education Production: Tracking, Decay and Student Achievement," *Quarterly Journal of Economics*, 2010, *125* (1), 175–214.

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

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.

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.

Online Appendix for:

Equilibrium Price Responses to Targeted Student Financial Aid

(Not for publication)

Nano Barahona Cauê Dobbin Sebastián Otero

A Appendix Figures and Tables

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

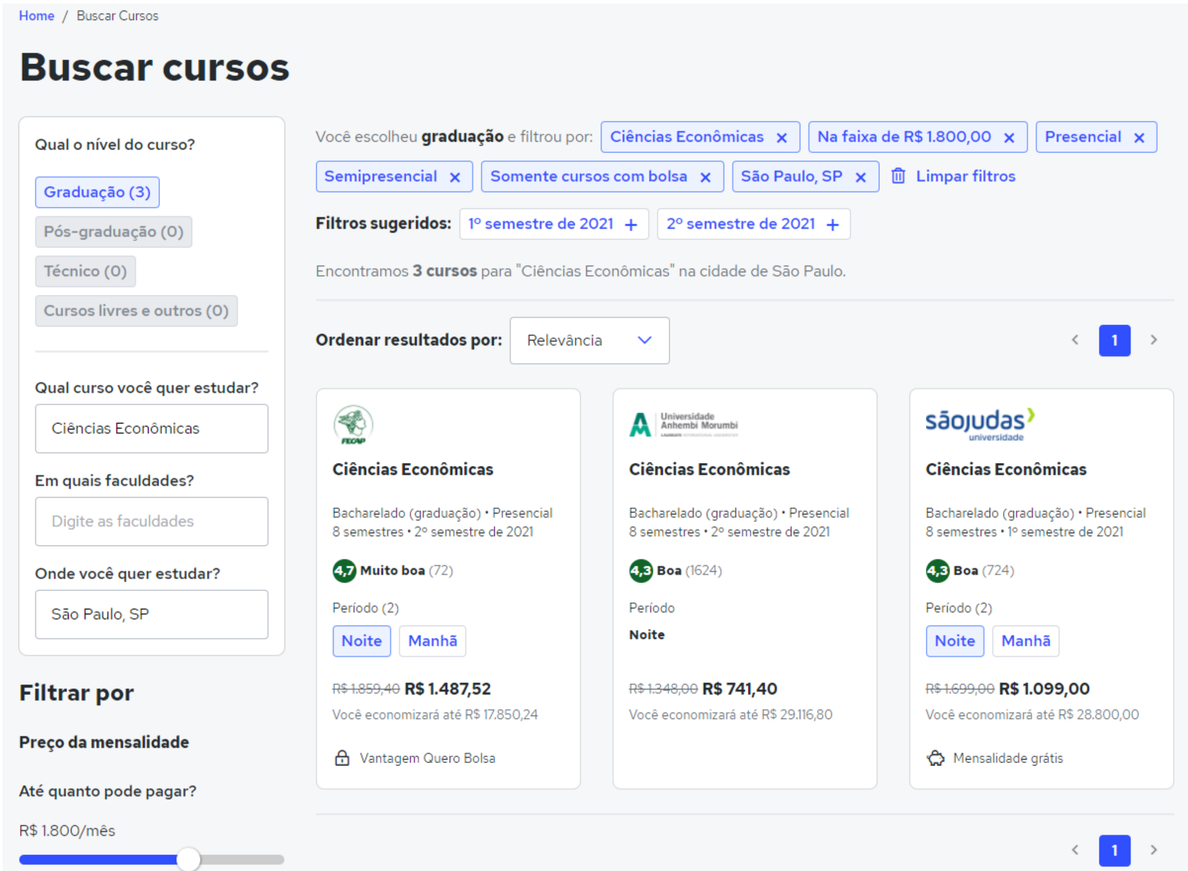

Notes: This figure presents a screenshot from QueroBolsa, Brazil's largest online marketplace for tuition discounts. The platform's interface enables students to filter and select degree programs by location, field of study, and tuition price. Additional details are provided in Section 3.1.

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

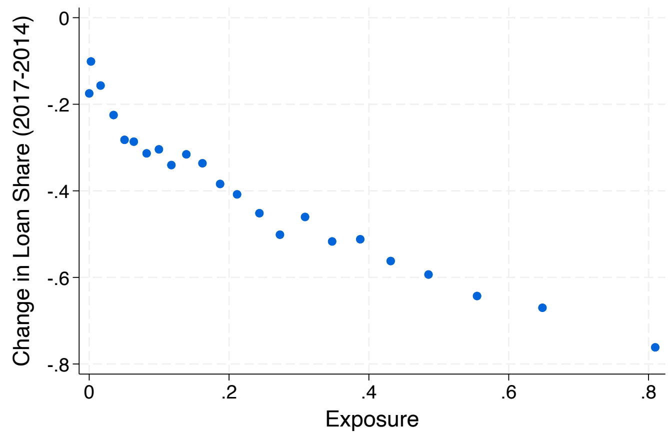

Notes: This figure displays a binned scatter plot illustrating differential exposure to the 2015 FIES reform, based on degree-level data. The x-axis reports the exposure measure defined in Section 4.2, while the y-axis shows the change in the share of incoming students with loans between 2014 (pre-reform)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

Figure A.3: Differential Exposure to the 2015 FIES Reform – Robustness

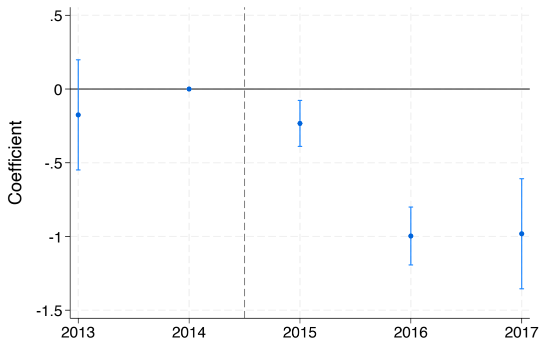

(a) Share with Loans

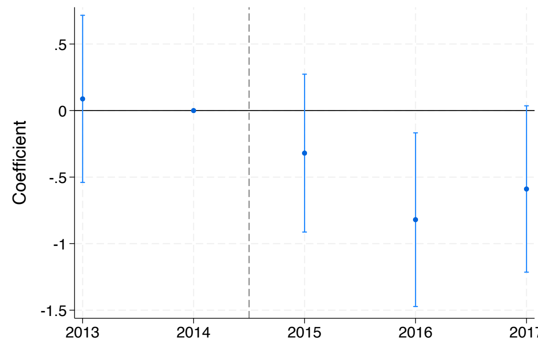

(b) log Enrollment

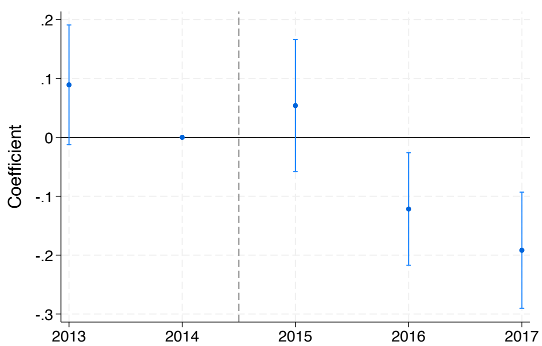

(c) log Tuition

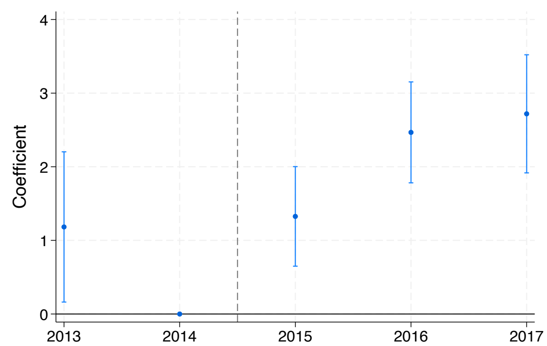

(d) log Enrollment w/o Loan

Notes: This figure presents OLS estimates of β_τ 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 number of incoming students; Panel (c), the log monthly tuition fee; 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.4: Distribution of $\Delta\eta$

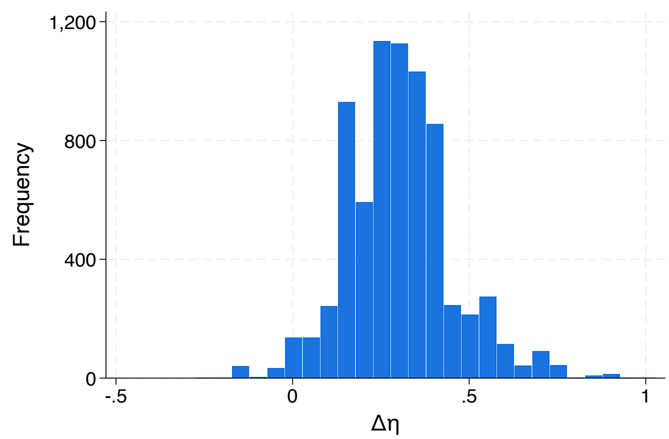

Notes: This figure shows the distribution of $\Delta\eta$, defined as the log difference between the average income of all students at a given college and that of FIES loan recipients at the same college, based on 2012 data. See Section 4.3 for further details.

Figure A.5: Price Elasticities, Marginal Costs, and Markups

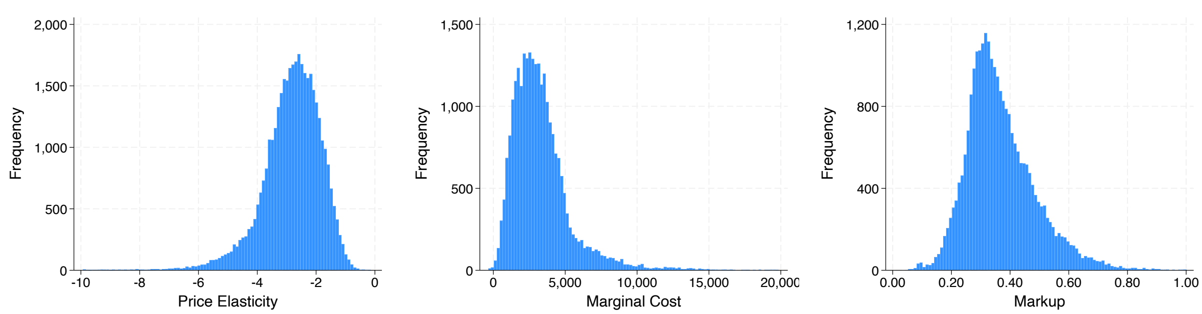

(a) Price elasticities

(b) Marginal Costs

(c) 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.6: Financial Aid Allocation, ENEM Scores, and Income

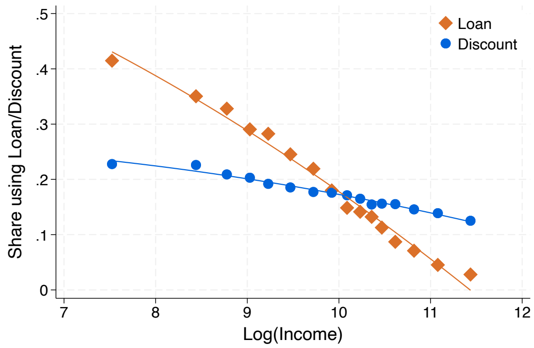

(a) Loans and Discounts

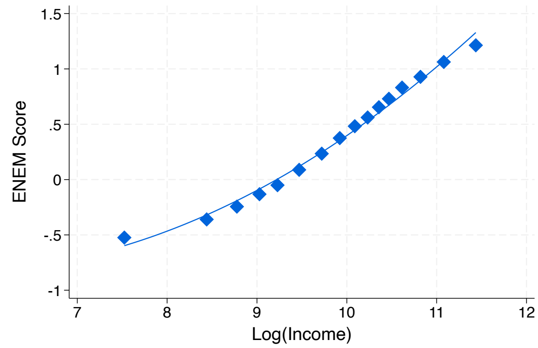

(b) ENEM Scores

Notes: This figure presents the relationship between financial aid allocation, test scores, and household income. In both panels, the x-axis reports log household income (in annual dollars). Panel (a) plots the share of enrolled students receiving loans (orange squares) or tuition discounts (blue dots). Panel (b) shows average ENEM scores. Both panels display 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.7: Equilibrium effects of loans on enrollment by program size

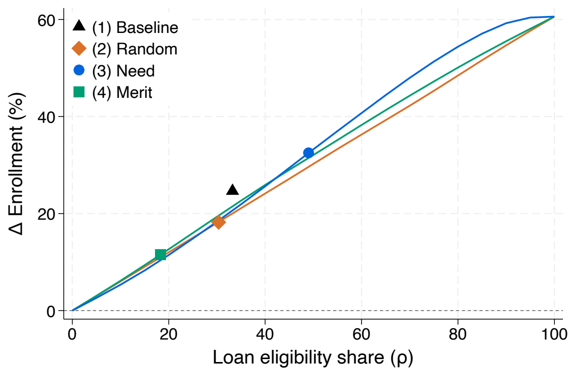

(a) Total enrollment: Demand-only

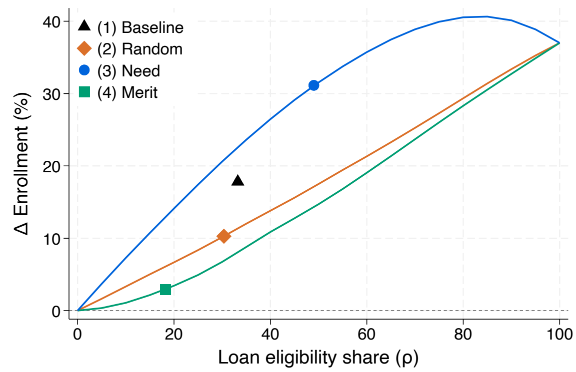

(b) Total enrollment: Equilibrium

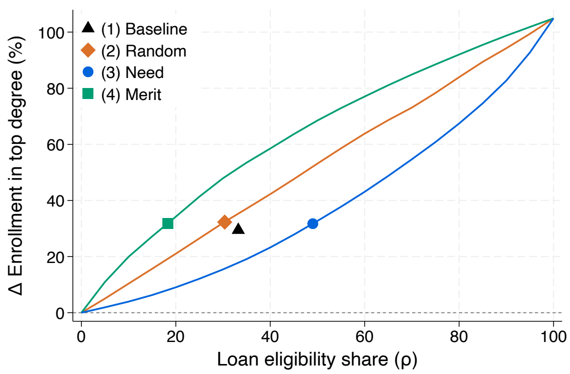

(c) Enrollment in top degree: Demand-only

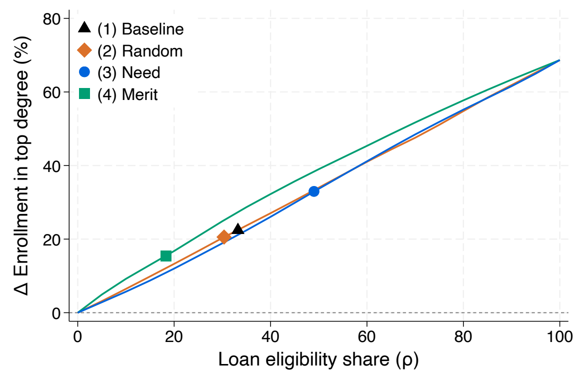

(d) Enrollment in top degree: Equilibrium

Notes: This figure shows the effect of student loans on enrollment under alternative targeting schemes. The y-axis reports the percentage change in enrollment relative to a scenario with no loans. Panels (a) and (b) display total enrollment effects, while Panels (c) and (d) focus on enrollment in programs in the top 10% of the value-added distribution. Panels (b) and (d) incorporate supply-side responses by allowing tuition to adjust endogenously, whereas Panels (a) and (c) hold prices fixed at their observed levels. Results are based on the baseline estimated parameters (black triangle) and counterfactual simulations under three alternative targeting schemes: random allocation (orange line), need-based targeting (blue line), and merit-based targeting (green line). Markers (square, diamond, and dot) indicate loan programs with the same total budget as the baseline (*budget-neutral*). The x-axis represents the loan availability share, ρ. In the baseline scheme, ρ denotes the share of student-degree pairs receiving loans. In the alternative schemes, ρ represents the share of students receiving loans, with each loan covering any degree in which the student enrolls. See Section 7 for additional details on the simulation procedure.

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

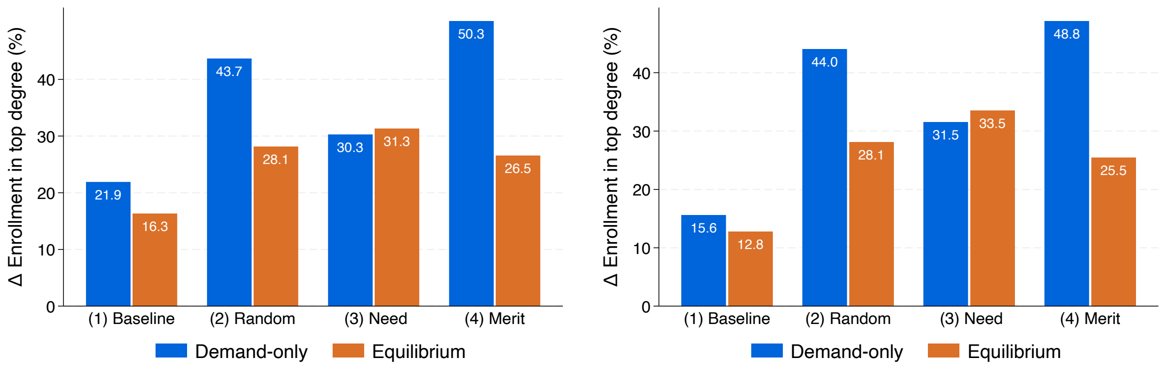

(a) Average Income

(b) Average Score

Notes: This figure presents the effects of student loans on enrollment under alternative targeting schemes. Outcomes are expressed as percentage changes relative to a scenario with no loans. For each scheme, simulations are conducted at an eligibility share ρ that holds the total loan program budget constant at the baseline level (*budget-neutral*). Both panels focus on enrollment in degrees in the top 10% of the quality distribution. Panel (a) ranks degrees by the average income of former students, while Panel (b) ranks them by the average ENEM score of incoming students. Orange bars allow for supply-side responses by letting tuition adjust endogenously; blue bars hold prices fixed at their observed levels. See Section 7 for additional details on the simulation procedure.

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

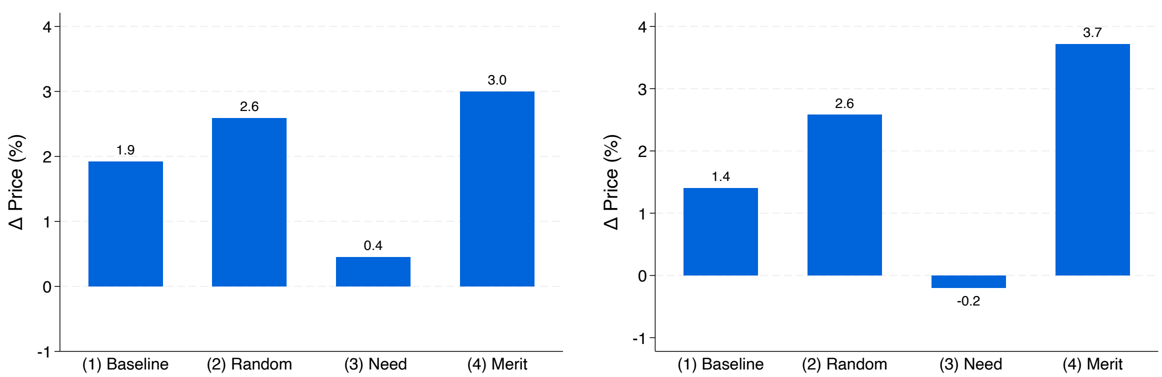

(a) All degrees

(b) Top degrees

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 an eligibility share 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 value-added distribution. See Section 7 for additional details on the simulation procedure.

Table A.1: Direct and Composition Prices 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	22,569	19,911	22,569	19,911	19,911
Degree FE	Yes	Yes	Yes	Yes	Yes
Region-Year FE	Yes	Yes	Yes	Yes	Yes
Major-Region-Year FE	No	Yes	No	Yes	Yes
Δ Inc-Year FE	No	No	No	No	Yes

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	3.532
	(0.008)
β^w	-0.604
	(0.002)
α^0_L	0.171
	(0.002)
α^0_{wL}	-0.583
	(0.003)
α^1	-5.339
	(0.002)
α^1_w	-0.354
	(0.001)
α^1_L	-0.191
	(0.002)
α^1_{wL}	-0.135
	(0.001)
Panel B: Financial Aid Targeting	
ρ^L_w	-0.9564
	(0.0004)
ρ^D_w	-0.0521
	(0.0004)

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.

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 Coverage rate

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 coverage 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 and national priorities. The allocation process had two main stages: degree-level seat allocation and student-level assignment.

Seat Allocation to Institutions. Each semester, the Ministry of Education (MEC) set a national cap on the number of FIES contracts available. Private colleges submitted proposals specifying the number of seats and tuition amounts they sought to have financed. The MEC evaluated these proposals based on degree quality (as measured by Sinaes scores), geographic location (with preference given to micro-regions with lower Human Development Index scores), and field of study. Starting in 2016, the allocation formula incorporated 70% of regional ENEM-based demand and 30% of historical FIES demand, weighted by municipal HDI factors ranging from 1.3 to 0.7. Within each micro-region, 70% of seats were reserved for priority fields—health, engineering, and teacher education—while the remaining 30% were allocated to other disciplines. Final allocations also prioritized programs with higher degree quality.

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), *13/2015* (for 2016 onward).

D Data: Details

D.1 Tuition Fees

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.

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 98.5% of total enrollment. All prices are deflated to 2015 price levels using the Brazilian consumer price index (IPCA).

D.2 Estimating Value Added

We track all university entrance exam (ENEM) participants and assign them to degrees based on their initial college enrollment, or to an outside option if they do not enroll. We then follow these individuals in the labor market 7 years after the exam, recording their highest salary from labor market administrative records (RAIS).

Our value added estimation relies on a standard selections-on-observables model (Rothstein, 2010; Angrist et al., 2017), which compares the labor market earnings of students across various degrees while controlling for test scores and an extensive range of student characteristics. All values are normalized relative to the outside option of not attending college. We allow the outside option to vary by region to account for local labor market conditions.

Specifically, we estimate the following model for each region r:

$$\log(Y_i) = \sum_j \text{VA}_j \cdot D_{ij} + X_i'\beta_r + \varepsilon_i, \tag{D.1}$$

where Y_i represents the wage income of student i; $D_{ij} \in \{0, 1\}$ indicates whether student i is enrolled in degree j; X_i includes students' characteristics, such as gender, age, ENEM score, and a constant; and ε_i captures other determinants of income that are uncorrelated with school enrollment. The term VA_j denotes the value added of each degree. We normalize the value added of the outside option to zero in each region and estimate the model via OLS.

We follow Angrist et al. (2023) and use empirical Bayes shrinkage methods to improve the precision of estimates. The estimates $\widehat{\text{VA}}_j$ from estimating Equation (D.1) via OLS are unbiased but noisy measures of the underlying degree-specific value added. We investigate the distribution of VA_j using the following hierarchical model:

$$\begin{aligned} \widehat{\text{VA}}_j | \text{VA}_j &\sim \mathcal{N}(\text{VA}_j, s_j^2), \\ \text{VA}_j &\sim \mathcal{N}(\mu, \sigma^2), \end{aligned}$$

where s_j^2 is the sampling variance of the estimator $\widehat{\text{VA}}_j$, while μ and σ^2 are the hyperparameters that govern the distribution of the latent parameters across degrees. Method of moments estimates of these hyperparameters are given by

$$\begin{aligned} \hat{\mu} &= \frac{1}{J} \sum_j \widehat{\text{VA}}_j, \\ \hat{\sigma}^2 &= \frac{1}{J} \sum_j \left[\left(\hat{\text{VA}}_j - \hat{\mu} \right)^2 - s_j^2 \right], \end{aligned}$$

where J is the number of programs offered.

The final step in the empirical Bayes estimation is to construct posteriors for the value added of each program. Given the model above, the posterior means are given by

$$\text{VA}_j^* = \left(\frac{\widehat{\sigma}^2}{\widehat{\sigma}^2 + s_j^2}\right)\widehat{\text{VA}}_j + \left(\frac{s_j^2}{\widehat{\sigma}^2 + s_j^2}\right)\hat{\mu}. \tag{D.2}$$

By shrinking the noisy estimate of VA_j toward the prior mean, the posterior mean reduces variance, with more shrinkage for programs with noisier estimates. We use the posterior means as our measure of degree quality.

To estimate the model, we use the universe of ENEM 2010 takers from the 2011 enrollment cycle and use student labor market earnings in 2017, 7 years after their enrollment decisions. Because we use only one cohort, our program value added measure remains constant over time.

E Model Estimation: Details

E.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.

E.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.[30] The sample otherwise follows the definition in Section 6.2 and is consistent with the rest of the paper.

E.1.2 Cutoff-based microments 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

[30] 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.

moments separately for each quartile of household income. Formally, we include the following micro-moment in the estimation:

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

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

E.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 (E.1) and test for manipulation. Specifically, we estimate the following local linear regression:

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

where $\text{Above}_{ij} \equiv \mathbb{1}\{h_i \geq \overline{h}_j\}$ is an indicator for whether student i's ENEM score h_i is above the loan eligibility cutoff $\overline{h}_j$ for degree j. The outcome y_{ij} is either q_{ij} or $p_{ij}q_{ij}$, as defined in Section E.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_{ij} 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 (E.2).

Results are shown in Figure E.1. Panel (a) presents the distribution of the running variable, the relative score $h_i - \overline{h}_j$. 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 E.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 (E.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.[31] Panel (c) plots the relationship between relative scores and tuition expenditures.

[31] 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.

The results show that tuition expenditures are also higher just above the cutoff compared with just below.

Figure E.1: Loan Eligibility: Regression Discontinuity Estimates

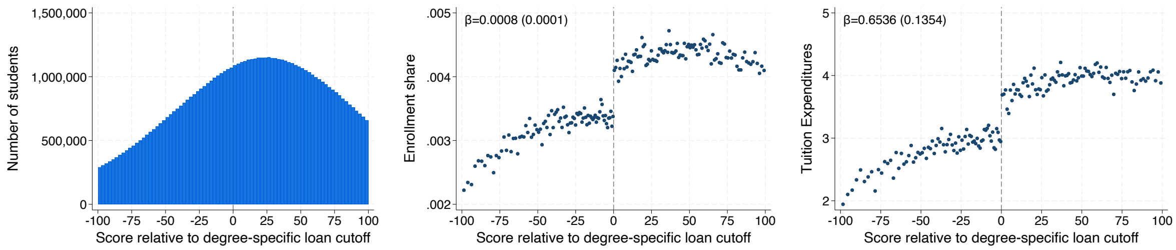

(a) Distribution of scores (b) Enrollment (c) Tuition expenditures

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 (E.2) are reported in both panels. Standard errors, clustered at degree level, are in parentheses.

To assess whether students just above and just below the cutoff differ in important ways, we estimate Equation (E.2) using a series of predetermined student characteristics as outcomes. Figure E.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.

Figure E.2: Loan eligibility: Regression discountinuity balance

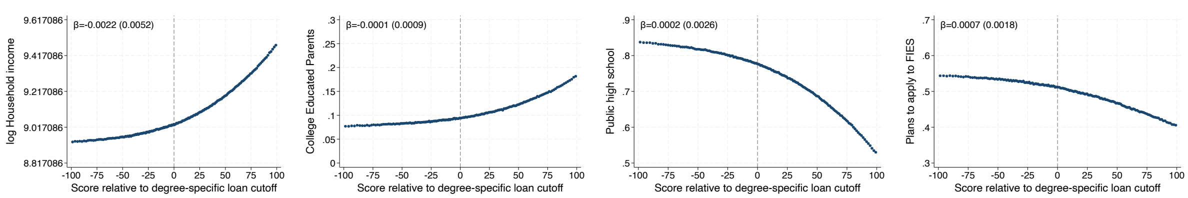

(a) Income (b) Parents' Education (c) Public High School (d) FIES Intention

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 (E.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.

E.2 Other micro-moments

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

$$\begin{aligned}
\text{Average score of enrolled students:} \quad & \frac{\mathbb{E}\left[h_i q_{ij}\right]}{\mathbb{E}\left[q_{ij}\right]} \\
\text{Average income of enrolled students:} \quad & \frac{\mathbb{E}\left[w_i q_{ij}\right]}{\mathbb{E}\left[q_{ij}\right]} \\
\text{Average income of students in high-price degrees:} \quad & \frac{\mathbb{E}\left[w_i q_{ij} \mid p^F_{jt} > \text{median}(p^F_{jt})\right]}{\mathbb{E}\left[q_{ij} \mid p^F_{jt} > \text{median}(p^F_{jt})\right]} \\
\text{Average income of enrolled students with loans:} \quad & \frac{\mathbb{E}\left[w_i q_{ij} \mid L_{ij} = 1\right]}{\mathbb{E}\left[q_{ij} \mid L_{ij} = 1\right]} \\
\text{Average income of enrolled students with discounts:} \quad & \frac{\mathbb{E}\left[w_i q_{ij} \mid D_{ij} = 1\right]}{\mathbb{E}\left[q_{ij} \mid D_{ij} = 1\right]}
\end{aligned}$$

E.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:

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

where Q is the GMM objective function, defined in Equation (21).

E.4 Model Fit

We now assess the fit of the model. Table E.1 presents the targeted micro-moments. Since the model is overidentified, it does not replicate the data moments exactly, but the deviations are small.

Table E.1: Model Fit: Targeted Micro-Moments

	Data	Model
$\mathbb{E}$[score|enrolled]	0.532	0.541
$\mathbb{E}$[income|enrolled]	9.130	9.174
$\mathbb{E}$[income|enrolled, high price degree]	9.366	9.347
$\mathbb{E}$[income|enrolled w/ loan]	8.79	8.797
$\mathbb{E}$[income|enrolled with discount]	9.006	8.976
Enrollment discountinuity, inc 1	0.505	0.518
Enrollment discountinuity, inc 2	0.496	0.438
Enrollment discountinuity, inc 3	0.514	0.450
Enrollment discountinuity, inc 4	0.484	0.461
Tuition expenses discontinuity, inc 1	8.766	8.779
Tuition expenses discontinuity, inc 2	8.752	8.712
Tuition expenses discontinuity, inc 3	8.758	8.706
Tuition expenses discontinuity, inc 4	8.687	8.692

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 E.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.

Our primary goal with the model is to predict equilibrium prices under counterfactual loan policies. To assess the model's predictive accuracy, we evaluate how well it replicates the observed price effects of the 2015 loan reform. Specifically, we simulate what prices would have been post-2015 if all other factors had remained constant, except for the allocation of loans. In this simulation, we hold mean utilities, δ_{jt}, costs, c_{jt} and κ_{jt}, and discount allocations, D_{ij}, at their 2014 levels, while allowing loan allocation, L_{ij}, to follow its 2016 and 2017 levels.[32]

Figure E.3 presents the results. The model predicts a 4.3% decline in tuition in 2016 and a 5.2% decline in 2017, compared to declines of 5.8% and 7.7% observed in the data, respectively. This implies that, according to the model, between 68% and 74% of the observed post-reform tuition decline can be attributed to changes in loan availability. The remainder is driven by other shifts in market conditions, captured in the model by changes in mean utility on the demand side and costs on the supply side. The model's ability to explain a large share of the observed price decline underscores the central role of student loans in shaping tuition levels, as well as

[32]For degree programs that only existed post-reform, we predict pre-reform mean utility and costs using the program's field of study, region, and the age of its institution.

the model's capacity to predict equilibrium outcomes under alternative loan policy scenarios.

Figure E.3: Model Fit: Price Effects of the 2015 FIES Reform

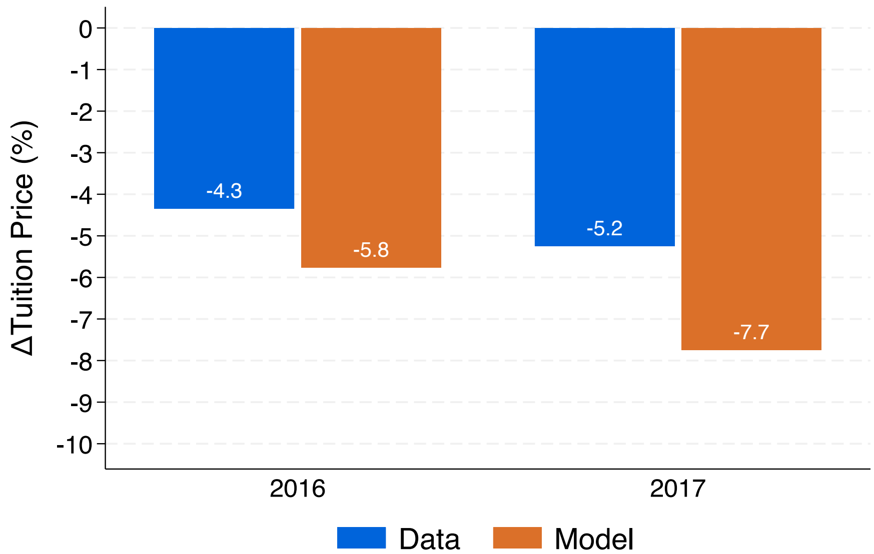

Notes: This figure evaluates the predictive accuracy of the model introduced in Section 5 by assessing its ability to replicate the observed price effects of the 2015 FIES reform. The blue bars represent the average tuition prices observed in each year, while the orange bars display simulated post-reform prices under a counterfactual scenario in which all factors remain constant at their 2014 levels, except for loan allocation. Tuition prices are computed as the enrollment-weighted average of full and discounted prices across degrees. Further details on this simulation are provided in Appendix E.4.