Affirmative Action in Centralized College Admission Systems: Evidence from Brazil*

Sebastián Otero[§] (Job Market Paper) Nano Barahona[†] Cauê Dobbin[‡]

November 12, 2021

[Link to the latest version]

Abstract

This paper empirically studies the distributional consequences of affirmative action in the context of a centralized college admission system. We examine the effects of a large-scale program in Brazil that mandated all federal public institutions to reserve half their seats for public high school students, prioritizing those from socioeconomically and racially marginalized groups. After the policy was put in place, the representation of public high school students of color in the most selective federal degrees increased by 73%. We exploit degree admission cutoffs to estimate the effects of increasing affirmative action by one reserved seat on the quality of the degree attended four years later. Our estimates indicate that the gains for benefited students are 1.6 times the costs experienced by displaced students. To study the effects of larger changes in affirmative action, we estimate a joint model of school choice and potential outcomes. We identify the parameters of the model using exogenous variation in test scores—arising from random assignment to graders of varying strictness—that changes the availability of degrees for otherwise identical individuals. We find that the policy creates impacts on college attendance and persistence that imply overall income gains of 1.16% for the average targeted student, and losses of 0.93% for the average non-targeted student. Overall, the policy prompted a negligible increase in predicted income of 0.1% across all students in the population. Taken together, we find that the affirmative action policy had important distributional consequences, which resulted in almost one-to-one transfers from the non-targeted to the targeted group. These results indicate that introducing affirmative action can increase equity without affecting the overall efficiency of the education system.

*We are indebted to Pascaline Dupas, Liran Einav, Matthew Gentzkow, Caroline Hoxby, and Paulo Somaini for invaluable guidance and support on this project. We also thank Claudia Allende, Luis Armona, Nick Bloom, Arun Chandrasekhar, José Ignacio Cuesta, Michael Dinerstein, Marcel Fafchamps, Diego Jiménez, Neale Mahoney, Ursula Mello, Melanie Morten, Christopher Neilson, and seminar participants at Stanford University for valuable comments and suggestions. We are grateful to Anna Popova for her insights, patience, and dedication to making this article better. We gratefully acknowledge financial support from the Stanford King Center on Global Development, the Stanford Center for Computational Social Sciences, the Stanford Institute for Economic Policy Research (SIEPR), the Haley-Shaw Dissertation Fellowship, and the NAEd/Spencer Dissertation Fellowship Program. All remaining errors are our own. [§]Stanford University. Email: `sotero@stanford.edu`. [†]University of California, Berkeley. Email: `nanobk@berkeley.edu`. [‡]Stanford University. Email: `dobbin@stanford.edu`.

1 Introduction

About one quarter of countries across the world use some form of affirmative action (AA) in an effort to increase the numbers of underrepresented groups in higher education (Jenkins and Moses, 2017). AA works by placing preferential weight on applications by students from marginalized groups at the expense of displacing non-targeted students. The prioritization of one group over another makes AA a highly contentious regulation in higher education policy (Arcidiacono and Lovenheim, 2016).[1] The discussion surrounding the value of AA policies centers around a potential trade-off between equity and efficiency.[2] While most policy efforts have been motivated by the equity merits of AA, much of the debate has focused on its impacts on the efficiency of the higher education system. At the heart of the efficiency discussion is the fact that AA pushes targeted students into selective degrees by displacing allegedly more qualified candidates.

There is contention around whether AA results in net positive efficiency benefits because theory fails to provide unambiguous predictions of the effects on winners and losers. On the one hand, AA policies could reduce allocative efficiency by matching high-performing non-targeted students to less-selective colleges if there are complementarities between college selectivity and academic preparedness in terms of returns (Ellison and Pathak, 2021). AA may even harm targeted students by placing them in schools for which they are ill-prepared (also known as the "mismatch theory") (Sander, 2004). On the other hand, AA could increase efficiency if displaced individuals have access to alternative high quality private colleges that are not affordable to targeted students. Quantifying this efficiency trade-off is thus an empirical question, but evidence on the distributional consequences of AA is limited.

This paper endeavors to fill this gap by estimating the effects of AA on the academic and implied labor market outcomes of winners and losers. To do this, we leverage an AA regulation in college education implemented in Brazil between 2013 and 2016. This is large-scale AA program comparable in magnitude to AA programs in India (Sales and Moses, 2014; Bagde et al., 2016). The regulation mandated all federal public universities to increase the number of seats reserved for students from public high schools ("targeted students") to 50% in every degree program.[3] The goal was to increase the representation of these groups closer to the 82% share of the population for which they account. Among public high schools students, the policy heavily targeted those from low-income and marginalized racial groups. After the policy, the share of public high schools students, and public high school students of color, admitted to federal degrees in the top decile of selectivity increased by 47% and 73%, respectively.

[1] AA in the United States has been at the center of this dispute, with the Supreme Court having rejected a number of university admission procedures explicitly favoring disadvantaged groups (e.g., Regents of the University of California vs. Bakke, 1997).

[2] Equity arguments highlight the role of AA in reinforcing the equalizing role of colleges and promoting diversity as a pillar of a sustainable and healthy democracy (Singer, 2011; Alon, 2015; Chetty et al., 2020).

[3] In contrast to the United States, prospective students apply in advance to particular degree programs, defined as a specific major at a specific institution (e.g. Law at the University of Sao Paulo).

We combine several datasets to characterize the population affected by the AA policy, and to determine its impact on several outcomes of interest. We have access to detailed administrative education data, including students' performance on the national university entrance exam, application portfolios to public universities, and admission offers from these institutions. We combine these administrative data with the Brazilian Higher Education Census to track students' progress and academic success in the universe of institutions and degree programs over time. Finally, we use matched employer-employee records of pre-policy cohorts comprising virtually all formal employment in Brazil, to estimate degree-specific labor market returns. This allows us to assess the impact of the AA policy on predicted income. Overall, these datasets provide a rich and comprehensive characterization of all individuals potentially affected by the AA policy.

Brazil is an informative setting for this analysis for two reasons. First, a major challenge in studying AA is that researchers seldom have access to detailed information on how college admission decisions are made. Brazil's AA policy allows us to overcome this difficulty because it is embedded in a centralized system, which produces transparent application data and clear admission rules. Students are admitted to a given degree if they score above the admission cutoff for that degree and type of seat (reserved or open).[4] Second, public federal institutions in Brazil are both highly selective and highly segregated. Federal universities play a similar role to flagship state universities in the United States, as they are typically the most prominent, elite, and selective universities in their specific states. These institutions are both free and higher quality than most of their private counterparts, which makes them very attractive to high-performing students from low- and high-socioeconomic status (SES) backgrounds alike.

We begin by producing regression discontinuity estimates of AA on academic outcomes for marginally benefited and marginally displaced students. We exploit the fact that reserved and open seats have different admission cutoffs, to construct different impact estimates for individuals entering through each type of seat. In contradiction with the predictions of the mismatch hypothesis, we find that gaining admission to a federal degree through a reserved seat improves the quality of the degree where the marginal targeted student is enrolled four years later.[5] We calculate that the gains for marginally benefited students are 1.2 and 1.6 times the costs experienced by marginally displaced students in terms of college enrollment and degree quality, respectively. This result is mostly explained by differential access to higher quality degrees outside of the federal system (i.e. outside options) across AA groups.

Understanding the overall impacts of the AA policy also requires investigating the consequences for individuals scoring further away from the admission cutoffs. With this purpose, we estimate a joint model of school choice and potential outcomes. The first part of our model involves recovering student preferences over degrees. We use the information on preferences contained in the centralized platform, and summarize student preferences by fitting random utility models to the application

[4] Prospective students apply in advance to a particular degree program: a specific major at a specific institution (e.g. Law at the University of Sao Paulo).

[5] We measure degree quality as the average score of students admitted in that degree before the AA policy

behavior of students for each AA group. We allow students' indirect utility to include degree-specific fixed effects, parameters that vary with observed student characteristics, and an unobserved degree-specific taste. Following Fack et al. (2019), we pin down the choice model parameters by leveraging the stability property of the deferred acceptance mechanism. This property implies that students are matched with their favorite degree among those available to them given their test scores, thus revealing their true preference over degree programs.

The second part of our model involves estimating the outcomes that would be realized under the counterfactual degree assignments. Estimating the treatment effects of attending a given degree is complicated due to selection bias; that is, potential outcomes could vary for students with different unobserved tastes for degrees in a way that is not captured by students' observables. To address this identification issue, we control for a rich set of covariates, including test scores on each component of the national exam. In addition, we use a selection correction approach by following the multinomial logit control function estimator of Dubin and McFadden (1984) and Abdulkadiroglu et al. (2020). This approach jointly models the choice of degrees together with the potential outcomes. Our model is flexible, as it can accommodate a variety of unobserved selection schemes, including selection based on student- and degree-specific unobserved matching effects.

To identify our model, we rely on an exogenous score shifter that mimics random shocks to student test scores. The score shifter interacts with degree admission thresholds to alter the degree programs that are available to otherwise identical individuals. The identification assumption is that, while score shifters affect degree availability, they do not enter into the potential outcome equation. The score shifter exploits two exogenous sources of variation in test scores. The first source stems from random assignment to essay graders of varying strictness. The second source arises from plausibly exogenous assignment to multiple choice examination booklets of varying difficulty. We use these to create a leave-one-out measure characterizing the overall exam difficulty faced by a given student. We show that the score shifter is uncorrelated with student characteristics but is highly correlated with test scores. Moving from the 5th to the 95th percentile of the distribution of the score shifter increases test scores by 0.17 standard deviations. In addition, receiving a positive score shifter is strongly associated with a higher probability of attending a federal degree, and with attending a higher quality degree within the federal system.

Because we do not observe income for the impacted cohorts, we assess the effects of the AA policy on predicted income, which we construct as an aggregator of intermediate academic outcomes.[6] We use our model to show that the implied effects on predicted income are equivalent to those resulting from realized earnings under a surrogacy assumption (Athey et al., 2019). This assumption states that realized earnings are independent of degree choice after conditioning on the student covariates and academic outcomes.

Our potential outcome estimates imply that, conditional on test scores, targeted individuals

[6] Our admissions data starts in 2016, and given that degree programs take between 4-6 years to complete, it is too soon to observe labor market outcomes for the individuals in our sample.

reap higher returns from attending federal degrees than do non-targeted individuals. This result is consistent with the fact that non-targeted students have access to better outside options than their targeted counterparts: conditional on test scores, non-targeted students attend private institutions that are higher quality and have higher tuition fees. Interestingly, we also find that test scores are similarly important in determining predicted income across federal degrees and outside options. This suggests that the potential test score mismatch induced by the AA regulation is unlikely to affect realized outcomes.

We then use our model to estimate the distributional consequences of a 50% AA schedule. We simulate the allocation of students with and without the AA policy, and find that most of the change in degree composition occurs within the most selective degrees. The AA policy differentially affects the probability of admission for targeted and non-targeted individuals, by changing the admission thresholds for reserved and open seats. Under the 50% AA policy, the admission thresholds for reserved seats are, on average, 44 points (or 0.6 standard deviations) lower than the admission thresholds for open seats.

We use our potential outcomes estimates to calculate the gains and losses resulting from a 50% AA schedule, given the changes in student admissions it implies. We find that, conditional on test scores, the policy causes targeted individuals to experience large gains in academic outcomes that predict large gains in income, while imposing a smaller cost on non-targeted individuals. Integrating over all affected students, our results indicate that the average targeted student gains 1.16% in predicted income relative to their baseline income without AA. In contrast, the average non-targeted student experiences losses equivalent to 0.93% of their baseline income. Overall, the policy prompted a negligible increase in predicted income of 0.1% across all students in the population. Taken together, we find that the AA policy had important distributional consequences, which resulted in almost one-to-one transfers from the non-targeted to the targeted group. These results indicate that introducing AA can increase equity without affecting the overall efficiency of the education system.

Our paper is related to several strands of literature. First, it contributes to the literature studying the effects of admission into selective colleges on educational and labor market outcomes. Several studies have leveraged admission discontinuities to document significant benefits of attending selective institutions for the marginally admitted student (Hoekstra, 2009; Zimmerman, 2014; Canaan and Mouganie, 2017; Zimmerman, 2019; Sekhri, 2020; Jia and Li, 2021; Bleemer, 2021b). A minority of papers have leveraged actual AA regulation, which induces variation in the admission thresholds for targeted and non-targeted students, to estimate the returns to attending a selective institution for each of these groups separately. Among these, Bagde et al. (2016) study a large-scale AA policy in India, finding positive academic impacts on targeted students. Francis-Tan and Tannuri-Pianto (2018) then examine a race-based AA policy in one federal institution in Brasil, finding positive impacts on targeted students in terms of both academic and labor market outcomes. Relative to this body of work, our contribution is two-fold. First, we combine admission thresholds

with an exogenous score shifter, which allows us to estimate effects for individuals away from the discontinuity. Second, we use these variation to estimate a model of the admissions system that allows us to simulate counterfactuals and compute overall efficiency gains.

A closely connected literature studies heterogeneity in returns to selective degrees by examining the consequences of the academic match between students and colleges (Bound et al., 2010; Dillon and Smith, 2017, 2020). Relative to these papers—because we know how admissions work within our centralized context and we observe how far individuals are from their counterfactual match—we are able to quantify the potential mismatch between students and degrees produced by AA. A key concern is the possibility that AA may create an extreme mismatch between the academic preparedness of targeted students and the difficulty of the degrees they attend, such that it would harm them. Several papers study the "mismatch hypothesis" with varying results, depending on the construction of counterfactuals (Ayres and Brooks, 2004; Sander, 2004; Chambers et al., 2004; Alon and Tienda, 2005; Fischer and Massey, 2007; Rothstein and Yoon, 2008; Arcidiacono et al., 2011). We provide new evidence on the potential mismatch effects by leveraging admission discontinuities which provide a credible counterfactual for targeted students in the absence of AA.[7]

Several papers are conceptually closer to ours in conducting structural policy analysis of college admissions (Arcidiacono, 2005; Howell, 2010; Bodoh-Creed and Hickman, 2019; Chetty et al., 2020; Kapor, 2020; Bleemer, 2021b; Tincani et al., 2021). These papers emphasize different margins of college admissions and simulate how AA could affect student enrollment and outcomes. While these studies have to explicitly model college admissions, we are able to simulate them by following the rules of the centralized mechanism. Relative to this existing work, our contribution is that we study the AA consequences not only on targeted students but also on displaced individuals, allowing us to estimate the distributional and the efficiency impacts of AA. One particularly relevant paper, that—like our own—estimates the effects of AA on both targeted and displaced students is Black et al. (2020), who study the trade-offs between winners and losers of the Top Ten% policy in Texas. They find that while benefited students increased college enrollment, graduation rates and earnings, students displaced by the policy do not see declines in these outcomes. Similar to our paper, this finding highlights the role that the outside option plays for displaced students.

Finally, our work also relates to a large literature on the estimation of joint models of treatment selection and outcomes (Heckman, 1979; Dubin and McFadden, 1984; Bjorklund and Moffitt, 1987), and to a literature connecting such models with IV estimators (Vytlacil, 2002; Kline and Walters, 2016; Brinch et al., 2017; Kline and Walters, 2019; Abdulkadiroglu et al., 2020). Our empirical approach is also related to a recent literature that endeavors to extrapolates regression discontinuity treatment effects away from the discontinuity. Several of these studies are motivated by understanding the impacts of AA on students of varying skill levels (Rokkanen, 2015; Angrist and

[7] One exception is Bleemer (2021a) who uses a differences-in-differences to evaluate the impacts of switching from race-sensitive to race-neutral admission policies at California public universities. He finds no evidence of mismatch in wages. Relative to this paper, we have the advantage that we observe the AA status of every student, and thus we do not need to infer the effects of AA from the overall effects for students from a given demographic group.

Rokkanen, 2015; Dong and Lewbel, 2015; Bertanha and Imbens, 2019). Relative to these papers, we use a selection model and identify the treatment effects for individuals away from the discontinuity by combining admission discontinuities with a continuous instrument that shifts students' test scores.

The remainder of the paper is structured follows. Section 2 introduces the setting and provides extensive details of the AA policy and its implementation. Section 3 discusses the data and provides descriptive facts about Brazilian college admissions. Section 4 presents regression discontinuity evidence. Section 5 presents a model that characterizes the impacts of AA in a centralized admission system. Section 6, introduces an empirical joint model of school choice and potential outcomes, and Section 7 presents the parameters estimates. Section 8 estimates the effects of a 50% AA regulation. Section 9 concludes.

2 Institutional background and regulation

2.1 Higher education in Brazil

The Brazilian post-secondary education system is a mixed system composed of public and private institutions that serves around 8 million students. The system is highly liberalized and market-oriented; approximately 76% of students are enrolled in private, fee-charging colleges.[8] The remaining 24% attend public institutions. These institutions charge virtually no tuition fees and are associated with the federal, state, or municipal government, depending on the source of funding. In 2016, there were a total of 107 federal institutions, which served nearly 1.25 million students, and offered more than 6,000 different degrees. Federal institutions comprise 63 universities and 44 vocational institutions, with the former accounting for the vast majority of federal enrollment. Appendix Table A.1 presents the number and share of students, institutions, and degrees, by sector.

Federal public universities in Brazil are, in most cases, elite and highly selective. These institutions play a similar role to flagship state universities in the United States. They are typically the most prominent universities in their state and are, on average, of substantially higher quality than their private counterparts, as measured by student learning, infrastructure, and the quality of peers and faculty.[9] The median state in Brazil has just one selective federal university. As a result of their high quality and tuition-free policy, public institutions are highly over-subscribed. This contrasts with private universities which, on average, fill only 50% of their offered spots.

Prospective students in Brazil, as is common in many countries, apply directly to specific degree programs: specific majors at specific institutions (e.g. Law at the University of Sao Paulo). In addition to selecting a degree, students must choose the teaching shift in which to study their degree:

[8] As a benchmark, in the United States, 75% of students attend public colleges, while only 25% attend a private institution.

[9] Appendix Figure A.1 shows the distribution of degree quality as measured by a Ministry of Education index ranging from 1 to 5.

morning, evening, night, or full-time. Obtaining a postsecondary degree normally requires 3-6 years for bachelor's degrees, 4-5 years for teaching degrees, and 2-3 years for vocational degrees.[10]

2.2 College admissions: ENEM and SISU

Several million students in Brazil take the national university entrance exam (ENEM) at the end of each academic year, with the aim of gaining access to higher education. The ENEM exam is an extremely competitive standardized exam, that consists of four multiple choice tests—Math, Language, History, and Science—and one written essay. ENEM test scores govern admission to most public institutions through a centralized admission system (SISU), and are the only merit-based criterion used by the Ministry of Education to assign financial aid to students attending private institutions.

Until 2010, college admissions in Brazil were fully decentralized, with each institution administering its own entrance exam. Since then, federal and state institutions can opt to participate in SISU, a centralized digital platform that matches students to degree programs according to their ENEM test scores, and degree preferences. By 2016, over 93% of all incoming students in federal institutions entered through this centralized admission system. In the same year, about 57% of the 4.8 million ENEM takers applied to a degree program using the SISU platform. The remaining ENEM takers attend private institutions, or take ENEM as a means to high school certification. The Brazilian academic year runs from March to December. Students take the ENEM test in December—at the end of a given academic year—and can use their test scores to apply to higher education through the SISU in January and begin their studies in March of the following academic year. Some degrees also open additional seats for a second rounds of admission through SISU in August.[11]

Students participating in SISU can submit and rank program choices among the set of available programs in the system. A program is defined as a degree, institution, and shift tuple. Students are then matched following an iterative deferred acceptance mechanism based on ENEM scores. The centralized system assigns students to degree programs based on their priority scores. Students' priority scores are given by a weighted average of their test scores and degree-specific weights for each of the ENEM components. In contrast to the standard deferred acceptance process, students are sequentially asked to submit rank ordered lists over the course of several "trial" days. Students can rank up to two program choices in each round. At the end of each day, the system produces a cutoff score representing the lowest score necessary to be accepted into a specific program, and students are allowed to change their degree program preferences given the newly available information. The results of the last day are final and determine the admission offers for every degree program.[12] Individuals who are not offered admission in their top preference can

[10] A student can only attend one shift and is not allowed to switch between shifts.

[11] Note that students taking the ENEM in a given year can only use their score to participate in the SISU process in the following year, either on the first or the second round of applications.

[12] This is the same as the mechanism currently used in university admissions by the province of Inner Mongolia in China. See Bó and Hakimov (2019) for formal properties of this mechanism.

opt-in to be waitlisted in that degree program.

2.3 The affirmative action regulation

In August 2012, the Brazilian federal government passed the *Law of Social Quotas* (Lei de Cotas, no. 12.711/2012) requiring all federal institutions to reserve half of their admission spots in each degree program for students coming from public high schools.[13] This group of students represent 82% of all ENEM test takers. Among these students, the policy heavily targeted low-income and marginalized racial groups—who account for most public high school enrollment—under the rationale that members of disadvantaged groups should not be underrepresented relative to their proportion in the population. The regulation sought to reverse racial and income inequality in university access.[14]

Under the regulation, only students coming from public high schools are eligible to compete for the AA vacancies, while the other half of vacancies remains open for broad competition.[15] Reserved seats are then further distributed among students from low-income families, and who are of African or indigenous descent.[16] Targeted students are eligible for either a reserved or an open seat, while non-targeted students can only apply to open seats. Throughout the paper, we refer to "targeted" and "non-targeted" students as the two different AA groups. After submitting an application, students can feasibly be admitted to a given degree program if its admission cutoff is below the student's priority score in the centralized admission system. These cutoffs are specific to each degree program and type of seat (reserved or open).

Figure 1 summarizes the distribution of seats under the policy by presenting the shares of AA students coming from different demographics. For every 100 university spots, 50 are AA spots for public high schools graduates. Of those reserved spots, 25 go to poor students with monthly household income per capita below 1.5 times the minimum wage (about 1,500 BRL, equivalent to 360 USD), and 25 go to non-poor students. Finally, a percentage of the spots in both of the former categories is set aside for black, brown and indigenous students, in accordance with the racial makeup of each of Brazil's 27 states.[17] Overall, this results in four different AA categories

[13] Students attending private high schools with full scholarships are also eligible to compete for AA vacancies. These students represent only 2.5% of all ENEM takers in 2016.

[14] In the 2000s, several federal and state Brazilian universities began to adopt socioeconomic- and race-based AA in admissions to address persistent disparities in college access. See for example, Francis and Tannuri-Pianto (2012a,b) and Estevan et al. (2018) who study AA policies at the University of Brasilia and the University of Campinas, respectively. The *Law of Social Quotas* was an effort to unify these regulations.

[15] Such systems are also known as reservation policies. The design of these reservations has been studied in several real-life applications. These include school choice in Boston (Dur et al., 2018) and Chicago (Dur et al., 2020), college choice in Brazil (Aygun and Bó, 2021), allocation of H-1B visas in the United States (Pathak et al., 2020), and public sector jobs in India (Sonmez and Yenmez, 2021).

[16] In the United States, selective universities doing AA implicitly inflate the application score of students who belong to a given minority. In the rest of the world, many countries use quotas and reserve a share of seats in their college admissions. Some examples include India (cast, gender), France (residence), South Africa (race), Malaysia (ethnicity), Sri Lanka (residence), Nigeria (residence), Romania (ethnicity), China (ethnicity), New Zealand (ethnicity), and Chile (socioeconomic status).

[17] Brazil has a large black and mixed-race population. According to the 2010 Census, the Brazilian population is

plus open seat category to which prospective students can apply.

In the first semester of 2016, the centralized admission system offered 200,877 spots across 4,900 degrees in 101 federal institutions. Table 1 presents summary statistics on the number of spots and applications by admission pool. The admission pools include the four AA categories plus one open pool. Nearly 46% of spots at federal institutions were open for anyone to apply, 48% were reserved for targeted individuals as mandated by the regulation, and the remaining 6% were reserved for institution-specific quotas (e.g. place-based AA). Notably, across all admission pools, selectivity—as measured by the ratio of spots to applications—was approximately 5%, comparable to many Ivy League colleges in the United States.

The regulation was implemented in a staggered fashion over four years. Starting in 2013, affected institutions were mandated to reserve a minimum of 12.5% of their vacancies for eligible students, with the minimum share increasing by 12.5 percentage points per year until it reached 50% in 2016. Institutions had the freedom to choose whether to adjust to the full regulation (50%) immediately or to adjust gradually, as long as they were complying with the mandated minimum share of AA spots in a given year.[18]

3 Data

3.1 Data sources

Test scores: The first dataset we employ contains test score data for all students taking the university entrance exam, ENEM. We observe these data for 2009-2015, which corresponds to the 2010-2016 admission period. These data include individual-level test scores on each of the components of the exam, as well as answers to a detailed survey including questions on socioeconomic background and student perceptions.

Centralized admissions process: We complement test score data with records from the centralized admission system, SISU. These data cover applications to federal and state institutions. We focus our attention only on applications to federal institutions, as state institutions were not mandated to comply with the AA regulation.[19] The dataset is at the application level, and only includes records from the final rank-ordered list submitted by each student.[20] For every application, we observe the ranking of a given degree program in the student preference list, the seat type of requested by the applicant, and the student ranking among all applicants from a given degree

47.5% white, 43.4% brown, 7.5% black, 1.1% asian, and 0.42% indigenous (IBGE, 2011).

[18] Appendix Figure A.2 illustrates the staggered implementation of the policy between 2012 and 2016. In every year since the regulation was introduced, we observe increased bunching of quota admissions above the 50% threshold. In addition, in 2013, 2014, and 2015, there is substantial bunching above the minimum thresholds corresponding to those years, until the 50% quota is surpassed in 2016. Between 2013 and 2016, virtually all degree programs complied with the minimum quota mandated for the given period. This illustrates that federal institutions complied with the regulation.

[19] For the structural analysis, we focus only on applications submitted to the state of Minas Gerais.

[20] The final rank-ordered list, refers to the last rank-ordered list submitted by the student throughout the multiple application rounds.

program and seat type combination. We also observe the students who were offered admission, and the admission cutoff for every degree program and seat type combination. We observe these data for the 2016 admission period. Throughout the paper, we define two measures of quality. When we compare degrees within the federal system, we follow Bagde et al. (2016) and use the admission cutoff for open seats as a proxy for degree quality. In contrast when we compare degrees inside and outside the federal system, we use the average test score of incoming students in a given year.

Higher education census: The third dataset we use for our analysis is the Brazilian Higher Education Census. This contains information on every student enrolled at any higher education institution in Brazil, allowing us to observe the educational path of every student in any degree program between 2009 and 2019. In a given year, the unit of observation is at the degree-student level, and includes a variable indicating if the individual graduated, dropped out, or is successfully enrolled at the end of the academic year. These data are of very high quality, as most institutions have their own systems integrated with the census in real time. In addition to the student data, this database also includes administrative information at the level of degrees and institutions, allowing us to characterize both over time.

Matched employer-employee data: Finally, we combine the previous data sources with matched employer-employee annual administrative records (also known as RAIS) from the Ministry of Labor. This is considered to be a high quality census of the Brazilian formal labor market.[21] This dataset includes variables at the firm and worker levels, such as payroll, contracted hours, hiring and firing dates, and occupation. We use these data to produce earning profiles for every degree and student type. We observe these data for 2017, and use them to create a mapping from student demographics, degree attendance and progression, to income. As we describe in more detail below, we use these estimates to create a measure of predicted income for the 2016 cohort.

3.2 Four key facts about test scores and college admissions

We highlight four sets of facts in the data. First, non-targeted students substantially outperform targeted students on the national entrance exam. Second, the composition of selective degrees became less segregated after the AA policy was in place. Third, the AA regulation improved the representation of targeted individuals by effectively lowering their admission thresholds. Fourth, non-targeted individuals have access to better and more expensive private degrees than do targeted individuals with similar test scores. The remainder of this subsection describes each of these facts in detail.

Fact 1: Differences in test performance by targeted status. One of the primary motivations behind the implementation of the AA regulation in Brazil is that ENEM test scores are strongly correlated with socioeconomic status and high school demographics. In 2016, targeted individuals represented nearly 81.5% of the 2.7 million SISU applicants. Figure 2 shows the distribution of the

[21] In 2015, the informality rate across all workers was 25%, but only 10% among individuals in the upper quartile of earnings (Engbom et al., 2021).

ENEM scores by targeted status among all SISU applicants. We observe that the average targeted student scored 518 points, while the average non-targeted student scored 590 points. This 72 point difference represents a 0.98 standard deviations (SD) difference in exam performance across targeted and non-targeted individuals. As a result, targeted students represent only 35% of all ENEM takers in the top 5% of scores (671 points).

Fact 2: Change in the student body composition. The student body composition of federal institutions became more diverse after the implementation of the AA policy, driven by an increase in the representation of targeted students. Panel (a) in Figure 3 presents the average share of targeted students in federal institutions by degree selectivity. In 2012, one year before the AA regulation was introduced, public high school graduates represented only 59% of admissions to all federal universities. In 2016, when the AA policy was fully in place, this share increased to 66%, becoming closer to the 82% share of public high school graduates in the population.[22] Most of the change was driven by an increase in the share of targeted students admitted to degrees above the 50th percentile of selectivity. The share of public high school graduates admitted to the top 10% most selective federal degree programs, increased from 35% in 2012 to 52% in 2016, and the share of public high school graduates of color increased from 17% to 29% during the same period (see Appendix Figure A.3).

Fact 3: Lower admission cutoffs for targeted students. We document that the AA policy introduced lower admission cutoff for targeted students. Table 1, column (6) presents the average admission cutoff for open and reserved seats, and Appendix Figure A.4 plots the distribution of the differences of cutoffs between open and reserved seats across all degrees. The difference between cutoffs for open and reserved seats was of 44 points on average, with wide heterogeneity in the magnitude of the admission cutoffs differences.[23] The differences in admission cutoffs partially mitigates the differences in ENEM test scores between targeted and non-targeted students. Panel (b) in Figure 3 plots the probability of admission in federal institutions over student test scores, for targeted and non-targeted students, before and after the regulation. We find increased differences in the probability of admission for targeted and non-targeted students after the law. These differences are larger for students in the upper part of the score distribution, where the regulation is more effective. In 2016, when the policy was fully implemented, targeted students in the top decile of the score distribution of ENEM takers were 18.7 percentage points more likely to attend a federal institution than non-targeted students, after conditioning on test scores. This difference represents a 187% increase relative to the 6.5 percentage point difference observed in 2012.

[22] These results are in line with Mello (2021), who exploits the progressive rollout of the AA policy to show that it increased the share of targeted students enrolled in federal institutions. Specifically, her study shows that the full adoption of the AA regulation, from 0 to 50% of reserved seats, increased enrollments of targeted students by 9.9 percentage points for the average degree program. It is important to notice that, in 2012, the average degree program was already admitting 22% of their students under reserved spots as a result of other institution-specific AA initiatives.

[23] We construct the average cutoff for reserved seats as the average cutoff across all four different type of reserved seats (weighted by the number of seats). The distribution of ENEM average test scores across all SISU applicants has a mean of 531 points and standard deviation of 73 points.

Fact 4: Differences in outside options by targeted status. We note that in the Brazilian context, the options that students can choose outside the federal system are substantially different for targeted and non-targeted individuals. Students not attending federal degrees can study at another public institution (either state or municipal), enroll at a private institution, or not attend college. Around 56% of all non-targeted 2016 SISU applicants who do not enroll at a federal university, enroll at a private institution. This share is only 30% for targeted students. Even more, conditional on attending a private institution, non-targeted students attend degrees of higher quality and higher tuition fees. It is important to note that targeted students are substantially more likely to receive financial aid, either through discounted prices, student loans, or grants (Dobbin et al., 2021). This last result helps mitigate the differences in access to more selective degrees in the private sector, thus improving the value of the outside option for targeted students.

4 The impact of affirmative action on marginal students

In this section, we use admission discontinuities to estimate the gains of marginally admitted and marginally displaced students in their preferred feasible degrees. In our setting, many students apply to federal degree programs while simultaneously deciding whether or not to attend a private institution, or to not study at all. Thus, some admitted students decline their offers, which opens spots for additional offers further down the admissions waiting list.[24] We compare individuals in similar waiting lists who were just above the admission threshold and offered admission, to those who were just below the threshold and thus not admitted.

We exploit the fact that reserved and open seats have different admission cutoffs, to construct credible instruments from discontinuities for individuals entering through each of these channels. As described in Section 2, there are four different admission tracks for reserved seats, and one for open seats. Thus, for every degree program we observe five different waitlists. We pool all admission cutoffs and normalize the data so that zero represents the position of the last individual admitted in a given waitlist.[25] Because, by definition, the last admitted student in the waitlist always accepts the offer, applicants receiving and not receiving offers are not statistically comparable in terms of their propensity to accept offers. We follow de Chaisemartin and Behaghel (2020), and drop the last admitted student in every waitlist to restore comparability across groups. For simpler presentation,

[24] Individuals in our setting have to report their top two degree program choices in the centralized system. They may either be (i) accepted at their first choice, (ii) accepted at their second choice and waitlisted at their first choice, or (iii) rejected from both their first and second choices but waitlisted at their first choice. Initially admitted students are offered admission through the centralized platform. After the acceptance deadline, the Ministry of Education forwards the list of waitlisted students to participating institutions, with the objective of filling the remaining vacancies. Thus, some students are offered spots immediately, whereas others may be offered seats closer to the beginning of the academic year.

[25] One natural alternative would be to pool the data using the difference between applicant test scores and the admission cutoff (Kirkeboen et al., 2016). Instead, we use the relative position in the waitlist. We do this because test score density varies drastically across degree programs. For example, while Medicine has hundreds of applicants within a 20-point bandwidth around the admission threshold, other degrees have only a handful. Using waitlist position as the index variable determining admissions allows us to have the same number of students across degrees, and below and above the admission threshold.

we also pool individuals across all four type of reserved seats, and present their estimates as a single group.

We provide statistical evidence of the effects of gaining admission into a federal degree on academic outcomes. We present this evidence graphically in Figure 4. To produce this figure, we first construct residualized outcomes, removing waitlist fixed effects. For readability, we rescale the residualized outcomes by the outcome means.[26] Seat types are indexed by $t \in \{0, 1\}$, were 0 denotes open seats, and 1 represents reserved seats. Figure 4 presents regression discontinuity estimates of the effects of AA for a number of academic and labor market outcomes and time periods. We plot binned residualized mean outcomes for individuals in each type of seat. As a result of pooling all targeted groups, reserved seat bins contain 4 individuals per degree program, while open seat bins contain only 1 individual per degree program. The confidence intervals in a given cell represent the standard error of a given outcome mean. We also include estimated linear regression lines on each side of the cutoff. To make groups comparable, we include observations within a small interval above and below the last admitted student.[27]

We now focus our attention on describing the main results in Figure 4. We begin by estimating the first stage relationship between students' admission offers and enrollment in their preferred degree during the first academic year. Panel (a) shows that the first stage coefficient for open seats is 30%. Take-up rates on the left-hand side of the discontinuity are different from zero because some programs offer a second round of applications during the second semester of the academic year. Take-up rates on the right-hand side of the discontinuity are low for two reasons. First, these correspond to the take-up rates of individuals in the waitlist. Every degree program manages their waitlist independently (following the order established by the centralized system). Sometimes it takes degrees a few weeks to clear their waitlists, and in the meantime, many of the waitlisted students enroll in private degrees. In addition, some degrees request students to come in person to sign their admission contracts in fixed dates, which makes it difficult for many individuals to take up on their offers. Second, these take-up rates mask great levels of heterogeneity; while highly selective degrees as Medicine has an average take-up of 70%, other less selective degrees at vocational institutions have take up rates that are substantially lower. As a result of lower take-up and lower reapplication rates in the process during the second semester of the academic year by targeted students, the first stage coefficient for reserved seats is 5 percentage points lower. This difference highlights how frictions in the admission process can differentially affect students from different backgrounds. Panel (b) presents enrollment rates in the preferred degree four years after gaining admission. Interestingly, even though dropout rates are high, at 43 percent, we find no differences in drop out rates by seat type. In Panel (c), we observe enrollment at any federal degree

[26] Formally, we regress the outcome on a constant, α, and a waitlist-specific fixed effect λ_{jq}, where j denotes the preferred degree j and q indicates one of the five admission tracks. The residualized outcome is given by $\tilde{Y}_i = Y_i - (\alpha + \lambda_{jq}) + \bar{Y}$.

[27] Note that since our running variable is discrete in nature, we do not follow the existing literature in calculating the optimal bandwidth, as those methods are developed for assignment variables with density $g(x)$, where $g(.)$ is continuous (Calonico et al., 2014).

four years after admission. This reveals that non-targeted students are substantially more likely to re-apply to a degree in the federal system.

Next, we compute the threshold-crossing effects on the quality of the degree attended. We measure degree quality by calculating the average test score of incoming students in a given degree. Since we have access to the universe of degrees and institutions in Brazil, we are able to compare the quality of the degrees inside and outside of the federal system.[28] In Panel (d), we look at what happens during the first year after admission. This figure provides three insights. First, students in open seats attend degrees of substantially higher quality than do their counterparts in reserved seats. This is a result of both having access to better degrees outside the federal system, and having higher take-up rate of admission to federal degrees. Second, gaining access to a federal degree substantially increases the quality of the degree attended by the student. Third, even though the regression discontinuity estimate is equivalent across open and reserved seats, the implied IV estimate indicates that the impact of assigning a seat in a federal degree on the quality of the degree attended is 20% larger for students in reserved seats.

In Panel (e), we focus on the quality of the degree students attend four years after enrollment. This outcome captures the net effect of gaining admission to a higher quality federal degree and individuals' academic paths after admission (be these inside or outside the federal system, or dropping out). The "mismatch hypothesis" implies that targeted students would experience worse outcomes as a result of being assigned to a higher quality degree. In contrast with these predictions, we find that AA improves the quality of the degree attended by targeted students four years later. Moreover, the IV estimate implies that gaining admission to a federal degree increases degree quality by 64 percent more for targeted students than for non-targeted students.

Finally, in Panel (f), we estimate the effect of gaining admission into a federal degree on predicted income. We do not use realized earnings because it is too early to follow the 2016 cohort in the labor market. Instead, we use individual-level microdata of older cohorts and construct a measure predicted earnings that takes into account the academic trajectories of students together with their demographic characteristics and test scores. This is our primary outcome in the model-based analysis, and its definition is explained and justified in Section 6.4. While the marginally displaced individual does not see a decrease in their expected earnings, the marginally benefited targeted student experiences a 10% income gain relative to the income of targeted students just below the admission threshold.[29] Taken together, these results suggest that targeted students benefit from AA, and that their gains more than compensate for the losses experienced by marginally displaced students.

In Appendix Figure A.5, we assess the validity of our research design by investigating covariate balance around the admission thresholds. We consider several individual characteristics: application score, race, gender, age, household income, and a dummy indicating if the student is a

[28] We assume that everyone not attending college is implicitly attending the same option. The degree quality measure for these students is given by the average test score of everyone who decided not to attend higher education.

[29] This is equivalent to a 270 BRL monthly gain, which is similar to an annualized gain of 1,000 USD.

recent high school graduate. Overall, we find no indication that applicants on different sides of the admission thresholds are different in term of observables.

5 A model of affirmative action in a centralized system

We are interested in estimating the consequences of large changes in the AA policy. For example, in the Brazilian case we would like to estimate the effects of a 50% AA policy relative to a counterfactual scenario without AA. In this section, we propose a framework to characterize such effects, which we use to guide the empirical analysis.

Relative to the analysis on marginal students, this framework allows us to go further in two different dimensions. First, it allows us to characterize the effects for individuals scoring away from the admission discontinuities. Second, we it allows us to capture the equilibrium effects of AA. These effects are important even for marginal increases in the number of reserved seats. To build intuition, consider a case in which we increase AA in the most selective degree by a single reserved seat. In the context of a centralized system, AA will change the equilibrium allocation between students and degrees through two channels: (i) the *direct* effect generated by the marginally benefited and the marginally displaced individual in the selective degree; and (ii) the *indirect* effects on all other degrees resulting from benefited and displaced individuals, respectively making room for or pushing out other students from their preferred degrees.

We consider a centralized system in which students rank their preferred degree programs, and institutions rank applicants using a priority index comprised of standardized test scores. Similar to Dur et al. (2020), our framework embeds an AA regulation within centralized admission systems by reserving admission spots for the targeted group. Because the number of spots is fixed, the AA policy operates by pulling in applicants from the targeted group at the cost of pushing out non-targeted students with higher academic scores.

5.1 Environment

We consider a set of individuals indexed by $i \in \mathcal{I} = \{1, ..., n\}$, applying to a finite set of selective college degree programs in federal institutions through a centralized platform. Let $\mathcal{J} = \{0, 1, ...J\}$ denote the set of degrees indexed by j offered across all federal institutions, where $j = 0$ represents the outside option of either attending a private institution or not attending college.

Students have preferences over degree programs based on a strict ordering, $\succ_i$. Degree programs also have preferences over applicants based on students' priority scores $s_i = \{s_{i1}, ..., s_{iJ}\}$. Scores are very fine so that no tie-breaker is needed. We allow for degree-specific scores, which may arise due to degree-specific exams, or due to degrees assigning different weights to different components of a single entrance exam.[30]

[30] For example, while STEM degrees tend to place higher weight on Math and Natural science components of ENEM, humanities-oriented degrees give more importance to Language and Social Science.

There are different subgroups within the student population, defined by their AA status $t_i \in \mathcal{T}$. For ease of exposition, let $\mathcal{T} = \{0, 1\}$, such that $t_i = 0$ represents non-targeted students and $t_i = 1$ represents targeted students. The status t_i of each student is observable. A student is fully characterized by their type $\theta_i = (\succ_i, s_i, t_i)$. That is, the combination of an applicant's preferences, priority scores across all degree programs, and AA status.[31] We denote the set of all student types by $\Theta = \bigcup_{i \in \mathcal{I}} \theta_i$.

Spots at degree programs are constrained by a strictly positive capacity vector, $q = \{q_0, ..., q_J\}$. An AA regulation is in place such that for every degree program, a share $\omega \in [0, 1]$ of the spots is reserved for applicants from the targeted group. The remaining share is open to all individuals. As such, for any given degree j, ωq_j spots are reserved for targeted students and $(1 - \omega)q_j$ spots are open to all individuals. We assume $q_0 = \infty$ since the outside option of not attending a federal institution is available to all students.

The centralized mechanism applies a student-proposing deferred acceptance algorithm to generate degree assignments. The inputs to the mechanism are student types Θ, school capacities q, and reservation shares ω. When reserved seats are processed, they are assigned to targeted students with the highest priority scores. When open seats are processed, they are assigned to members of either group (targeted or non-targeted) in order of their priority scores. When a student qualifies for both a reserved and an open spot, the admission rules must specify which seats are processed first, i.e. the relative precedence of different admissions channels. In our empirical application, we assume that spots are horizontally reserved. That is, reserved spots are first assigned to targeted students based on priority scores, and next all open spots are assigned to the remaining individuals with the highest priority scores.[32]

Let $\varphi(\Theta, q, \omega) = \mu$ denote the matching produced by mechanism φ for the problem (Θ, q, ω). The matching is a function $\mu : \Theta \to \mathcal{J}$, such that (i) $\mu(\theta_i) = j$ if student i is assigned to degree j, and (ii) no degree is assigned more students than its capacity. Because the mechanism implements a deferred acceptance algorithm, assignments μ between students and degrees are unique and stable (Gale and Shapley, 1962; Abdulkadiroglu, 2005).[33] The stability property of the mechanism implies that each student enrolls in their most preferred degree among those for which they are eligible.

The matching μ has a unique representation in terms of a vector of market clearing admission cutoffs $c_{tj}(\mu)$ for each degree program and AA status combination (Azevedo and Leshno, 2016).[34]

[31] The definition of a student type as a combination of preferences, priorities, and status is similar to that of Abdulkadiroglu et al. (2017).

[32] The literature has established a difference between horizontal and vertical AA depending on which of the spots (reserved or open, respectively) are processed first. A horizontal reservation is a "minimum guarantee" in the sense that it is only binding when there are not enough targeted students who receive a spot on the basis of their test score alone. A vertical reservation works under an "over-and-above" basis. This means that targeted individuals receiving spots on the basis of their priority scores alone do not count towards vertically reserved spots (Sonmez and Yenmez, 2021). Our framework encompasses both types of AA.

[33] A matching μ is stable if there is no student-degree pair (i, j) where: (i) student i prefers degree j to their assignment, and (ii) student i has higher priority than some other student who is assigned to degree j.

[34] Other mechanisms also have a unique representation in terms of admission thresholds, as long as the matches

A cutoff $c_{tj}(\mu)$ is a minimum score s_{ij} required for admission at degree j for students with AA status t. Since targeted students can be admitted through the reserved or the open spots, admission cutoffs for targeted students are always lower than for non-targeted students, i.e. $c_{0j} \geq c_{1j}$ for all $j \in \mathcal{J}$. Because the outside option is always feasible, it has an admission cutoff score of $-\infty$.

As in any strict-priority mechanism, the availability of options will be different for students according to their priority scores. Let $\Omega_i(\mu) = \{j \in \mathcal{J} | s_{ij} \geq c_{t_i j}(\mu)\}$ represent the *feasible choice set* for individual i, defined as those degrees to which they could have gained access based on their score and AA status under a given matching μ. Let the variable $D_i(\mu) = \{j \in \Omega_i | j \succeq k \text{ for all } k \in \Omega_i(\mu)\}$ denote the preferred option in the feasible choice set defined by $\Omega_i(\mu)$. In other words, $D_i(\mu)$ represents the option ranked highest by individual i among the degrees to which they could have been admitted. We refer to this option as the *preferred feasible degree*. From the stability condition we know that $D_i(\mu) = \mu(\theta_i)$.

The realized outcome for student i is given by $Y_i(\mu) = \sum_j \mathbb{1}\{D_i(\mu) = j\} \cdot Y_{ij}$, where $D_i(\mu)$ indicates the degree attended, and Y_{ij} denotes the potential value of some outcome of interest for student i if they attend degree j.

5.2 Gains and losses of affirmative action

One of the advantages of studying AA in the context of a centralized system is that it follows systematic and transparent admission rules. By leveraging the assignment rules from the centralized mechanism, we are able to characterize student allocations in a regime with a different AA schedule ω'. For ease of exposition, we think of ω' as an increase in the share of seats reserved for AA students.

The new schedule, combined with student types and degree capacities, results in a counterfactual matching function $\varphi(\Theta, q, \omega') = \mu'$. In turn, this allocation can be represented by a new vector of of admission thresholds $c_{tj}(\mu')$.[35] In this scenario, the feasible choice sets, preferred feasible degree, and ultimately the realized outcomes, would also change as a result of the changing admission cutoffs. These objects are represented by $\Omega_i(\mu')$, $D_i(\mu')$, and $Y_i(\mu')$, respectively. Combining these elements, we are able to characterize the set of individuals who are effectively benefited and displaced by the AA policy, together with their respective gains and losses.

Our goal is to define how much a student would gain or lose if they were pulled into or pushed out of their preferred feasible degree as a result of an increase in the AA policy. The conditional average treatment effect of increasing ω to ω' for individuals of type θ is:

$$\tau(\theta) = \mathbb{E}[Y_i(\mu') - Y_i(\mu) | \, \theta_i = \theta]. \tag{1}$$

between students and degrees are pairwise stable (Agarwal and Somaini, 2018).

[35] In the extreme case of $\omega = 0$, applications from targeted and non-targeted students are processed together and spots are allocated solely based on priority scores, and thus every student would face the same admission threshold irrespective of their AA status.

The aggregate effect of increasing AA is given by integrating over all student types. To aggregate these treatment effects, we assume equal weights for all individuals within a given AA group. Thus, the aggregate effects for each AA group t are given by:

$$\Delta_t(\omega', \omega) \quad = \quad \sum_{i \in \mathcal{I}} \tau(\theta_i) \cdot \mathbb{1}\{t_i = t\}. \tag{2}$$

We can use the status-quo and counterfactual admission thresholds, $c_{tj}(\mu)$, and $c_{tj}(\mu')$, to identify the individuals who gain or lose access to any given degree as a result of the policy. For the case of targeted students, individuals with $D_i(\mu') = j$ and priority score $s_{ij} \in [c_{1j}(\mu'), c_{1j}(\mu))$ will benefit and gain admission into degree j when the AA policy is intensified. Conversely, non-targeted individuals with priority scores $s_{ij} \in [c_{0j}(\mu), c_{0j}(\mu'))$, will be displaced from degree j by the regulation.

5.3 A simple example

Under this framework, the aggregate consequences of the AA policy are given by how much pulled-in students gain, and how much pushed-out students lose in terms of a given outcome of interest. For instance, policy makers could be interested in how much income benefited individuals gain, and how much income displaced individuals lose as a result of the policy. These gains and losses are characterized by how admission scores, degree attendance, and the outside options affect the outcome of interest.

In Figure 5, we describe one case in which the income gains for targeted individuals outweigh the losses experienced by displaced individuals. For expositional clarity, assume that $\mathcal{J} = \{0, 1\}$, such that there is only one selective degree available, or an outside option. For simplicity, we also assume that all students have strict preferences for the selective degree ($j = 1$) over the outside option ($j = 0$).

The solid lines show the average potential earnings for students attending the selective institution, conditional on their score and AA type, $\mathbb{E}[Y_{i1} | \, t_i, s_i]$. The dashed lines depict the average potential earnings from attending the outside option, $\mathbb{E}[Y_{i0} | \, t_i, s_i]$. In the absence of AA (i.e. $\omega = 0$), there is a single admission cutoff c', for all students. By focusing on students scoring around c', we can easily observe that implementing AA is efficient for the marginal student. predicted income gains from attending the selective institution are higher for targeted than for non-targeted students. In addition, non-targeted students have a better outside option than targeted individuals. As a result, the returns to attending the selective institution are higher for targeted than for non-targeted students.

Now, assume policy makers implement an AA policy such that targeted and non-targeted students face different cutoffs, represented by $c_{t=1}$ and $c_{t=0}$ respectively. Targeted students with scores $s_i \in [c_{t=1}, c']$ are now offered admission to the selective degree, while non-targeted students with scores $s_i \in [c', c_{t=0}]$ are displaced from it. Area A depicts the aggregate gains for benefited

individuals, while area B depicts the aggregate costs incurred by displaced individuals. In the figure, we observe that the consequences of AA depend crucially on heterogenous returns for benefited and displaced individuals, and also on how this difference in returns varies with admission cutoffs.

6 Econometric model

In our conceptual framework, we characterized the effect of increasing the AA schedule from ω to ω' for individuals of type θ. Using the stability condition, we rewrite Equation (1) as:

$$\tau(\theta) = \mathbb{E}\left[\sum_j \mathbb{1}\{\mu'(\theta_i) = j\} \cdot Y_{ij} - \sum_j \mathbb{1}\{\mu(\theta_i) = j\} \cdot Y_{ij} \;\middle|\; \theta_i = \theta\right], \tag{3}$$

where μ and μ' represent the matching function for each AA schedule, and Y_{ij} denotes the potential value of some outcome of interest for individual i if they attend degree j. Accordingly, estimating $\tau(\theta)$ in the data requires estimating the matching functions for each of the AA schedules, together with the potential outcome of attending a given degree. With that purpose, we start by specifying and parametrizing a joint model of school choice and outcomes. Next, we present the key identifying variation that allow us to estimate the effects of AA away from the admission discontinuities. We end the section introducing and defining our primary outcome variable.

6.1 School choice model

To estimate the matching function, we leverage the rules of the mechanism φ, together with its inputs (Θ, q, ω). In practice, however, researchers never observe the full set of student types Θ, as individuals rarely exhaustively rank the full list of degree preferences $\succ_i$. To recover student preferences, we introduce an empirical school choice model.

We summarize the observed choices in the centralized platform by fitting a random utility model. Specifically, student i's indirect utility for attending degree j is:

$$\begin{aligned} u_{ij} &= V_{ij} + \eta_{ij} \\ &= \delta_j^t + \gamma_j^t \cdot s_{ij} + \kappa^t \cdot d_{ij} + \eta_{ij}, \end{aligned} \tag{4}$$

where V_{ij} captures the part of the utility that varies according to the observed characteristics of students and degrees, and η_{ij} captures unobserved tastes for degrees.[36] We parametrize V_{ij} as a function of degree fixed effects δ_j^t, student test scores s_{ij}, and d_{ij} indicating whether the student lives in a commuting zone that is different to where the degree is offered.[37] We allow flexible

[36] Implicit in Equation (4) is an assumption that preference parameters $(\delta_j^t, \gamma_j^t, \kappa^t)$ can rationalize true utilities for degrees under any realization of the matching function. This assumption implies that preferences for degrees are independent of the observed allocation of students, ruling out preferences for peers. Although this assumption might sound restrictive, most existing empirical approaches abstract away from equilibrium sorting based on preferences for peers (Agarwal and Somaini, 2020).

[37] We use micro regions as proxies for commuting zones, to capture whether individuals need to move to a new

heterogeneity in tastes by estimating preference models separately for each AA type t. We model the unobserved taste η_{ij} as following an extreme value type I distribution, conditional on vectors s_i and $d_i = (d_{i1}, ..., d_{iJ})$. The outside option aggregates all degrees offered by municipal, state, and private universities, as well as not enrolling at all. This outside option is available to everyone and has a deterministic utility V_{i0} normalized to zero.

Let $\tilde{\mu}$ be a realized matching that we observe in the data. The mechanism's stability property implies that students are faced with the option of selecting a degree from an observable and personalized choice set $\Omega_i = \Omega_i(\tilde{\mu})$ (Fack et al., 2019). It also implies that the degree to which the student is assigned is also their preferred feasible option ex-post, that is $D_i = D_i(\tilde{\mu}) = \tilde{\mu}(\theta_i) = \arg\max_{j \in \Omega_i} u_{ij}$. One of the main advantages of final stable allocations is that they allow researchers to exploit revealed preference relationships based only on assignment data, instead of having to rely on the information reported in rank-ordered lists.[38] As such, this approach can be implemented in any centralized system using strict priorities with stable assignments. The stability property, as an ex-post optimality condition, is not necessarily guaranteed when students do not have complete information about admission cutoffs. In Appendix B, we show that in the Brazilian setting, students have little uncertainty over the final admission cutoffs due to the iterative property of the mechanism.

Following Fack et al. (2019), we restrict the choice set Ω_i to all degrees available to student i based on their score and on the admission cutoff specific to their degree and AA status. The logit model implies that the probability that individual i selects degree j is given by:

$$Pr(D_i = j) = \frac{\exp V_{ij}}{\sum_{k \in \Omega_i} \exp V_{ik}}.$$

We estimate $(\delta_j^t, \gamma_j^t, \kappa^t)$ for $t \in \{0, 1\}$ using a maximum likelihood estimator.

6.2 Potential outcomes model

Next, we switch our focus to constructing selection-corrected estimates of a potential outcome model using a control function approach. We follow a similar approach to that of Abdulkadiroglu et al. (2020), who link school choices to potential outcomes to estimate schools' value-added in New York City. We use our estimated parameters to predict the potential outcomes for attending a given counterfactual preferred degree.

residence as a result of geographical distance from a given degree. There are a total of 558 micro regions in Brazil, and they are defined by IBGE as bordering cities with "common social characteristics, geography, and spatial articulation."

[38] An alternative route to estimating degree preferences would be to rely on the rank-ordered lists submitted by students to the centralized system. The deferred acceptance mechanism is strategy-proof, which implies that it is in the best interests of students to rank schools truthfully (Abdulkadiroglu and Sonmez, 2003). As a result, a large body of research has used submitted rank-ordered lists to infer student preferences over degree programs. We take a different approach for two reasons. First, in the SISU system students can only rank up to two options, which destroys the strategy proofness of the mechanism (Haeringer and Klijn, 2009; Calsamiglia et al., 2010). Second, as pointed out by Fack et al. (2019), in strict priority mechanisms where individuals are ranked by test scores, students face limited uncertainty about their admission outcomes and may "skip the impossible" and choose not to apply to degree programs that are out of their reach.

We project the potential outcome Y_{ij} of student i at counterfactual degree j on degree-specific fixed effects and student and degree characteristics:

$$Y_{ij} = \alpha_j^t + X_{ij}'\beta_j^t + \varepsilon_{ij} \tag{5}$$

where α_j^t and β_j^t are population parameters for AA group t, implying $\mathbb{E}[\varepsilon_{ij}] = \mathbb{E}[X_{ij}\varepsilon_{ij}] = 0$. The variable X_{ij} is a vector of observed covariates, including student test scores s_{ij}, and d_{ij} which indicates whether student's i lives in the same commuting zone where the degree is offered. Our goal is to recover the parameters of the potential outcome equations defined above.

The mean outcome Y_i observed in the data for a given matching $\tilde{\mu}$ is given by:

$$\mathbb{E}[Y_{ij}|X_{ij}, D_i = j] \quad = \quad \alpha_j^t + X_{ij}'\beta_j^t + \mathbb{E}[\varepsilon_{ij}|X_{ij}, D_i = j] \tag{6}$$

The OLS estimation of this equation would likely yield biased parameters due to selection into degrees based on unobservable preferences. To recover unbiased estimates we would need to assume that $\mathbb{E}[\varepsilon_{ij}|X_{ij}, D_i = j] = 0$, thus implying that degree choices and potential outcomes are not correlated after accounting for student and degree observed characteristics.

To account for selection on unobservables, we link the outcome equation to the school choice model by conditioning Equation (5) on the vector of unobserved tastes $\eta_i = (\eta_{i0}, \eta_{i1}, ..., \eta_{iJ})$:

$$\mathbb{E}[Y_{ij}|X_{ij}, \eta_i] \quad = \quad \alpha_j^t + X_{ij}'\beta_j^t + \mathbb{E}[\varepsilon_{ij}|X_{ij}, \eta_i] \tag{7}$$

$$= \quad \alpha_j^t + X_{ij}'\beta_j^t + g^t(\eta_i). \tag{8}$$

This model allows expected potential outcomes to vary across students with different preferences for degrees in a way that is not captured by students' observables.[39] To estimate Equation (8) we use the multinomial logit selection model of Dubin and McFadden (1984), which imposes a linear relationship between potential outcomes and the unobserved logit errors. Imposing such a parametric approximation on $g^t(\cdot)$ yields:

$$\mathbb{E}[Y_{ij}|X_{ij}, \eta_i] = \alpha_j^t + X_{ij}'\beta_j^t + \sum_{k=0}^{J} \psi_k^t \cdot (\eta_{ik} - \bar{\eta}) + \rho^t \cdot (\eta_{ij} - \bar{\eta}) \tag{9}$$

where $\bar{\eta} \equiv \mathbb{E}[\eta_{ij}]$ is Euler's constant. As pointed out by Abdulkadiroglu et al. (2020), this parametric relationship allows for a wide range of selection on unobservables in the context of school choice. The parameter ψ_k captures the effect of the preference for degree k that is common across all potential outcomes. For example, students with high preferences for a given type of degree may have higher outcomes in all other degrees in a way that is not fully captured by student observables (e.g., students enrolling in medicine may be also high in motivation and thus do well in any other

[39] From Equation (7) to (8) we impose a separability assumption that implies that the conditional expectation of ε_{ij} as a function of η_i does not depend on X_{ij}. This is a common assumption in applied work that uses instrumental variables to identify selection models (Kline and Walters, 2016; Brinch et al., 2017; Abdulkadiroglu et al., 2020).

degree). We refer to this term as selection on levels. The coefficient ρ represents the match effect of preferring degree j. This unobserved match component allows, for instance, for students to sort into degrees based on potential outcome gains. We refer to this type of selection as selection on gains (Roy, 1951).

By iterated expectations, the mean outcome observed in the data is:

$$\mathbb{E}[Y_i|X_{ij}, \Omega_i, D_i = j] \quad = \quad \alpha_j^t + X_{ij}'\beta_j^t + \sum_{k=0}^{J} \psi_k^t \cdot \lambda_{ik}(\Omega_i) + \rho^t \cdot \lambda_{ij}(\Omega_i), \tag{10}$$

where $\lambda_{ik}(\Omega_i) \equiv \mathbb{E}[\eta_{ik} - \bar{\eta}|X_{ij}, \Omega_i, D_i]$ is the expectation of the unobserved preference for a given degree, conditional on the student's characteristics X_{ij}, their feasible choice set Ω_i, and their preferred feasible degree D_i. These functions have a closed-form solution and can be computed using the logit functional form. These objects serve as control functions to correct for selection on unobservables.

The estimation of the outcome model proceeds in two steps. First, we compute $\hat{\lambda}_{ik}(\cdot)$ using the estimated preference parameters and the logit functional form. Next, we plug $\hat{\lambda}_{ik}(\cdot)$ into Equation (10), and estimate parameters $(\alpha_j^t, \beta_j^t, \psi_j^t, \rho^t)$ using separate OLS regressions for each AA group t.

6.3 Model identification

To identify our model, we rely on an exogenous score shifter z_i that mimics random shocks to student test scores. The score shifter exploits two exogenous sources of variation in test scores. The first source stems from random assignment to essay graders of varying strictness. The second source arises from plausibly exogenous assignment to multiple choice examination booklets of varying difficulty. We describe these exogenous sources of variation in detail in the next subsection. In terms of our model, individuals with positive score shifters will see an exogenous increase in their personalized choice sets $\Omega_i(z_i)$ (i.e. they could be admitted to a larger set of degree programs), while individuals with negative score shifters will experience an exogenous decrease in their personalized choice sets.

We use the score shifter z_i, together with admission discontinuities, to identify the potential outcomes (α_j^t and β_j^t) and selection parameters (ψ_k^t and ρ^t) from Equation (10). Identification of the parameters α_j^t and β_j^t comes from exogenous variation in choice sets $\Omega_i(z_i)$ available to otherwise identical individuals (in terms of both observed characteristics and unobserved preferences). To provide intuition, suppose there are two available degree choices, A and B, and two individuals, 1 and 2, who share an identical preference for degree B. Assume that both individuals have access to degree A but individual 2 also has access to degree B as a result of receiving a positive score shifter $z_i = z^+$, which allowed them to cross the admission threshold. The difference in outcomes for these individuals pins down the treatment effect of attending degree B over degree A.

Identification of the selection parameters ψ_k^t and ρ^t comes from the variation in available choice

sets for students who enroll in the same degree program. To provide intuition, suppose there are two additional individuals, 3 and 4, who are identical on observables but may have different unobserved preferences. Assume both individuals have access to degree A, but individual 4 also has access to degree B as a result of receiving a positive score shifter $z_i = z^+$. If both individuals choose degree A, we can use a revealed preference argument to learn that individual 4 has an unobserved taste for option A which is higher in expectation than that revealed by individual 1. The selection parameters capture whether this expected difference in unobserved preferences is relevant for the potential outcome. In Appendix C, we provide a formal identification proof.

The identification of Equation (10) relies on the assumption that choice sets and score shifters are exogenous to unobserved tastes η_i and potential outcome errors $\varepsilon_i = (\varepsilon_{i0}, \varepsilon_{i1}, ..., \varepsilon_{iJ})$ after conditioning on X_{ij}; that is $(\varepsilon_i, \eta_i | X_{ij}) \perp z_i, \Omega_i(z_i)$. While the independence of z_i is satisfied because of the random assignment of graders and booklets to students, the independence of $\Omega_i(z_i)$ is implied by the model, since personalized choice sets are a function of student observable characteristics and the score shifter $\Omega_i = f^t(X_{ij}, z_i)$. Implicit in this assumption is the fact that students take admission cutoffs as given and cannot manipulate them through their application behavior. Finally, we need an exclusion restriction that ensures that while test scores affect the availability of degrees, only corrected test scores enter into the potential outcome equations.

Identifying variation in the data: Next, we describe the nature of our exogenous score shifter z_i. The ENEM exam consists of a written essay and four multiple choice tests.[40] We leverage features of the test implementation to construct exogenous score shifters for each of these components. These shifters mimic random shocks that impact test scores but are uncorrelated with students' latent ability. The score shifter exploits two exogenous sources of variation. We discuss them in turn.

The first source of variation stems from random assignment of the essay component of the national exam to graders of varying strictness. Each essay is marked by two randomly assigned graders. Each year, over 10,000 graders are involved in the grading process, each of whom is assigned to an average of 1,100 exams.[41] Graders receive the exams via an online platform, which ensures that they are blind to students' identity and characteristics. The only restriction imposed in the assignment process is that graders must be from different states to their assigned students. The essay is marked based on five different competencies, each scored on a scale from 0 to 200. The final grade is the simple average of the total points given by each of the graders. If the score difference across graders is large for any exam, that exam is graded by a third grader, and the final grade is then the average of the two closest scores.[42]

[40] The essay usually has a 20% weight for most degrees. This holds true across different fields of study.

[41] In Appendix Figure A.6 we show the distribution of the number of exams assigned to each grader.

[42] A third grader intervenes in two situations: (i) if the difference between total scores across graders is larger than 100 points, and (ii) if the difference ibetween scores for any of the individual competencies is larger than 80 points. The grading rubric is available from the following link:
http://download.inep.gov.br/educacao_basica/enem/guia_participante/2018/manual_de_redacao_do_enem_2018.pdf

We construct a leave-one-out measure of grader leniency for every student using the first two randomly assigned graders. Specifically, we define the essay score shifter z_i^e as:

$$z_i^e = \frac{1}{2}\left(\underbrace{\frac{1}{|N_{g_{i1}}-1|}\sum_{m\in\left(N_{g_{i1}}\setminus\{i\}\right)} s_m^e}_{\text{strictness of grader 1}} + \underbrace{\frac{1}{|N_{g_{i2}}-1|}\sum_{m\in\left(N_{g_{i2}}\setminus\{i\}\right)} s_m^e}_{\text{strictness of grader 2}}\right),$$

where g_{ik} denotes the kth grader assigned to individual i; s_m^e represents the essay score of student m, where m indexes the set of all other individuals assigned to the same grader as student i; and N_g denotes the set of exams received by grader g. In words, for a given student i, the leave-one-out score shifter is the average score their graders gave to all other students they graded.

Panel (a) in Figure 6 presents the distribution of the score shifter for the SISU sample. It shows a wide dispersion with a range of 165 points. Given the large number of essays marked by each grader, in the absence of any leniency differences, the leave-one-out mean leniency for each grader should be concentrated around 576 points, the mean essay score. The dashed red line shows a local linear regression of the first-stage relationship between our score shifter and the essay score. The relationship is strong, positive, and mostly linear. The dashed blue line serves as a balance test and plots a local linear regression of the average multiple choice score components of the test. Consistent with a satisfactory randomized assignment, this figure shows no correlation between the score shifter z_i^e and other measures of student ability, as captured by the average score in the multiple choice components of the test. We present the first stage coefficients and additional balance tests in Appendix E.

The second source of variation arises from the plausibly random assignment of students to examination booklets on the multiple choice component of the test. To prevent cheating, students are assigned to one of four different examination booklets. Each booklet type has a different color cover, and their contents vary only in the order in which questions are presented. In every examination location, students are organized alphabetically both across and within rooms, where they are assigned to one of the booklets. Accordingly, booklet assignment is plausibly random.[43]

We observe substantive differences in test scores across examination booklets, which we present in Panel (b) of Figure 6. For instance, the average score in the Math score is 6 points higher for red booklets than for yellow booklets. In Appendix E, we show that booklet assignment is uncorrelated with students characteristics. In Barahona et al. (2021), we show that, despite including the same questions, some booklets are more difficult than others because they include the easier questions later in the test. This means that students may run out of time or steam before they reach the easier questions. For each component k in {Math, Language, History, Science} in the multiple

[43]Examination locations receive a tag for each participant, with their name, ID number, examination room, and desk number. Students can only take the exam if the tag matches the ID attached to the student's desk

choice part of the test, we construct the following leave-one-out measure of booklet difficulty:

$$z_i^k = \frac{1}{|N_{b_{ik}} - 1|} \sum_{m \in \left(N_{b_{ik}} \setminus \{i\}\right)} s_m^k,$$

where b_{ik} denotes the booklet assigned to student i, s_m^k is student ms score on component k, and N_{b_k} is the set of students assigned to booklet b for each component k. For a given student i, the leave-one-out score shifter z_i^k captures the average score of every other student assigned to the same examination booklet.

We combine these two sources of variation and create a corrected measure of student test scores by netting out the effect of the score shifters from the observed student test scores. Specifically, for a given component k in {Essay, Math, Language, History, Science} we impose the following parametric restriction:

$$\begin{aligned} \tilde{s}_i^k &= f^k(s_i^k, z_i^k) \\ &= s_i^k + \phi^k(z_i^k - \bar{z}^k), \end{aligned}$$

where s_i^k is the observed score, z_i^k is the score shifter, $\bar{z}^k$ is the average score shifter, and $\tilde{s}_i^k$ represents the corrected test score. An underlying assumption in this parametric model is that the score shifter z_i^k affects all students homogeneously, that is $\phi_i^k = \phi^k$ for all i.[44] We use these to construct corrected test scores that are relevant for the centralized system as $\tilde{s}_{ij} = \sum_k w_j^k \cdot \tilde{s}_i^k$, where w_j is the weight given by degree j to component k.

In Figure 7, we present the reduced-form relationship between our instrument and educational outcomes. Panel (a) shows that a student who received a score shifter in the top 5% of the distribution is 1.55 percentage points (or 12%) more likely to attend a federal institution the year after than one who got a score shifter in the bottom 5% of the distribution. In Panel (b), we show that the score shifter also meaningfully impacts student's predicted income.

6.4 Predicted income as primary outcome

An important challenge is that we do not observe long-term outcomes Y_i for our sample of interest (which we denote by $\mathcal{N}$). To overcome this issue, we write our outcome of interest Y_{ij} as a function of an observed intermediate academic outcome P_{ij} (e.g. income as a function of college completion rates). We allow these variables to vary by degree-specific fixed effects and by student and degree characteristics:

$$P_{ij} = \pi_j^t + X_{ij}\psi_j^t + \xi_{ij}/\zeta^t \tag{11}$$

$$Y_{ij} = \gamma_j^t + X_{ij}\rho_j^t + \zeta^t P_{ij} + \nu_{ij}, \tag{12}$$

[44] This assumption implies the monotonicity assumption which is standard in the so-called "judge designs" that leverage similar sources of exogenous variation.

where the term ξ_{ij}/ζ^t denotes the unobserved component of the intermediate outcome, and ν_{ij} represents the unobserved component of the long-term outcome equation after accounting for the intermediate outcome.[45]

We use Equations (11) and (12) and rewrite the long-term outcome Y_{ij} as:

$$\begin{aligned} Y_{ij} &= \alpha_j^t + X_{ij}'\beta_j^t + \varepsilon_{ij} \\ &\equiv \alpha_j^t + X_{ij}'\beta_j^t + \xi_{ij} + \nu_{ij}, \end{aligned}$$

where $\varepsilon_{ij} \equiv \xi_{ij} + \nu_{ij}$, and α_j^t and β_j^t are the same population parameters indicated in Equation (5). The conditional expectation of Y_i given degree assignment D_i is:

$$\begin{aligned} \mathbb{E}[Y_i|X_{ij}, D_i = j] &= \alpha_j^t + X_{ij}'\beta_j^t + \mathbb{E}[\varepsilon_{ij}|X_{ij}, D_i = j] \\ &\equiv \alpha_j^t + X_{ij}'\beta_j^t + \mathbb{E}[\xi_{ij}|X_{ij}, D_i = j] + \mathbb{E}[\nu_{ij}|X_{ij}, D_i = j]. \end{aligned}$$

Note that in this context, unobserved selection can arise from two sources. The first source is $\mathbb{E}[\xi_{ij}|X_{ij}, D_i = j]$, which implies that selection into degrees may correlate with intermediate outcomes in a way that is not fully captured by student observables. The second source is $\mathbb{E}[\nu_{ij}|X_{ij}, D_i = j]$, which encapsulates the effect of unobserved selection after accounting for the intermediate outcome.

Because we do not observe the long-term outcome, we follow Athey et al. (2019) and transform our intermediate outcome into a "surrogate index;" that is, the predicted value of the long-term outcome given the intermediate outcome and student covariates. Specifically, we define the "surrogate index" for individuals in our sample of interest as:

$$\tilde{Y}_{ij} = \mathbb{E}[Y_i|X_{ij}, P_i, D_i = j,\, i \in \mathcal{N}].$$

To recover $\tilde{Y}_{ij}$, we proceed in two steps. First, we estimate the parameters of Equation (11) using older cohorts (which we denote by $\mathcal{O}$), for whom we observe both intermediate and long-term outcomes. Second, we use these coefficients to predict a measure of the long-term outcome $\tilde{Y}_{ij}$ for our sample of interest $\mathcal{N}$.

To recover unbiased parameters of the long-term outcome equation using the "surrogate index," we establish two assumptions:

$$\begin{aligned} \text{A1 Surrogacy:} \quad & \mathbb{E}[\nu_{ij}|X_{ij}, P_{ij}, D_i = j] = 0 \\ \text{A2 Comparability:} \quad & \mathbb{E}[Y_i|X_{ij}, P_i, D_i = j,\, i \in \mathcal{N}] = \mathbb{E}[Y_i|X_{ij}, P_i, D_i = j,\, i \in \mathcal{O}]. \end{aligned}$$

Assumption A1 is usually referred to as the "surrogacy" assumption and states that the long-term outcome is independent of degree assignment, conditional on the intermediate outcome. In our ap-

[45] We scale ξ_{ij} by ζ^t for expositional clarity.

plication, this assumption implies that unobserved selection can affect the quality and area of the degree from which students graduate but does not affect earnings after accounting for graduation rates. The second assumption ensures that the conditional expectation functions of long-term outcomes are comparable for samples $\mathcal{O}$ and $\mathcal{N}$. That is, the mapping from individual characteristics, degree assignment, and intermediate outcomes to long-term outcomes, is fixed and invariant to the sample. This assumption rules out, for instance, any general equilibrium effects as a result of AA. Given these assumptions, we can rewrite the predicted long-term outcome $\tilde{Y}_{ij}$ for individuals in sample $\mathcal{N}$ as:

$$\tilde{Y}_{ij} = \alpha_j^t + X_{ij}'\beta_j^t + \varepsilon_{ij}. \tag{13}$$

For our sample of interest, consisting of the universe of SISU applicants in the 2016 academic year, we observe college enrollment and detailed academic progress over four years, from 2016 to 2019. For older cohorts, we observe both academic progress and income.[46] We use individual-level microdata for cohorts entering college between 2010 and 2012, and their observed labor market income in 2017, to estimate a model linking academic progress, student demographics, and test scores to income. We then use this model to estimate predicted incomes for students in our sample based on their academic trajectories. We use this predicted income measure as our main outcome of interest. We discuss the econometric implementation of this procedure in Appendix D.

Using predicted income as a primary outcome measure has two important advantages relative to using other academic measures, such as dropout rates. First, it allows us to compare individuals with different trajectories that would otherwise be impossible to weigh up against each other.[47] Second, under the surrogacy assumption discussed above, we can interpret the predicted income results as the long-term effects of the AA policy.

7 Parameter estimates

In this section, we present the parameter estimates from our model. We first discuss the preference parameters that arise from estimating the school choice model. To estimate the model, we focus on applications to federal degrees through the SISU platform in the first semester of 2016. We restrict the sample to every applicant submitting an application to any institution in the state of Minas Gerais. We do this exclusively due to computational power constraints, and future versions

[46] After being admitted to a given degree, individuals face a number of choices that determine their academic trajectories. They can choose whether or not to enroll in the degree to which they are admitted, and subsequently whether to continue or to drop out. Conditional on dropping out, they could re-apply to other federal degrees, enroll in degrees outside the federal system, or remain out of college.

[47] Take, for example, the case of three individuals who enroll together in Nursing. Assume that, after four years, the first individual remains enrolled in the same degree, while the second individual drops out and switches to Medicine, and the third individual drops out of college. In this case, dropout rates would not be a good metric to compare the current academic status of the three students. High switching rates are common in centralized systems. See Larroucau and Rios (2021), for instance, for evidence on switching rates into lower and higher quality degrees in Chile.

of this paper will scale the analysis to the rest of the country. This is the largest state in terms of number of applications, and represents 17% of all applicants in the system. In total, we have 11 institutions offering 502 degrees (including the outside option). Our sample size consists of 329,621 students. We define a student as targeted if we observe them submitting any application through the reserved seats admission pools. In contrast, non-targeted students are defined as those who submit an application only through the open seats pool.

In Panel (a) of Table 2, we present the parameters estimates from the school choice model. We find negative coefficients for the distance parameter, indicating that students have strong preferences for studying close to where they live. The distance coefficient is -2.73 for targeted individuals, and slightly smaller for non-targeted students, at -2.36.

Parameters related to degree fixed effects (δ_j^t) or those associated with student test scores (γ_j^t) are very high-dimensional, as we have one different estimated coefficient per degree and AA status tuple. To visualize these parameters, we estimate the average valuation for a given degree $\bar{V}_j^t$. To make estimates comparable across targeted and non-targeted groups, we fix the test score variable to equal the admission cutoff score for open seats for that degree, c_{0j}, as well as fixing the distance variable equal to 0 (i.e., the student lives in the commuting zone where the degree is offered). Specifically, we define $\bar{V}_j^t$ as:

$$\bar{V}_j^t = \hat{\delta}_j^t + \hat{\gamma}_j^t c_{0j}$$

The value of $\bar{V}_j^t$ indicates the valuation of a given degree relative to the outside option. The larger $\bar{V}_j^t$, the more likely it is that the student would choose degree j over the outside option if only these two degrees were offered to them.

In Panel (a) of Figure 8, we show how these valuations vary by degree selectivity and AA status. First, we note that, relative to the outside option, degrees have a low valuation. This is not surprising as the outside option pools a collection of different alternatives (e.g., working, attending a private institution, waiting another year to go to college) and its value is the maximum value of all these alternatives. Second, we find that targeted students put a lower value on the outside option relative to attending a federal degree when compared to their non-targeted counterparts. If non-targeted students can attend good private degrees outside of the federal system, they will be more likely to choose the outside option rather than to enroll in a given federal degree. Third, there is a very strong correlation between degree selectivity and average degree valuation. Fourth, there are a handful of degrees at the top of the selectivity distribution that are remarkably more valued than other degrees. These degrees mostly consist of Medicine and Engineering-related programs.

We use these parameters to assess the in-sample model fit by comparing the degree admission thresholds of degrees predicted by the model to those observed in the data. We discuss the construction of the predicted admission thresholds in Appendix F. In Appendix Figure F.1, we show the scatterplot of simulated and observed admission thresholds. The correlation coefficient between

simulated and observed admission thresholds is virtually 1 for both open and reserved seats.

Next, we switch our attention to the outcome model. Parameters α_j^t and β_j^t of the potential outcome equation govern the structural relationship between predicted income, degree attendance, and student characteristics. Our model estimation produces a vector of these estimates for each degree program and AA tuple. These estimates are unbiased but noisy measures of the underlying outcome parameters. To improve our estimates, we follow Abdulkadiroglu et al. (2020) and use an empirical Bayes shrinkage estimator to reduce the sampling variance. This yields empirical Bayes posterior means of degree program-specific parameters, which we then use in our subsequent analyses (see Appendix G for a discussion on the implementation of this procedure).

To visualize these parameters, we define the value-added of attending a given degree as the gains from attending that degree relative to those from attending the outside option. To calculate the value-added, we use the potential outcome parameters α_j^t and β_j^t, and fix the test score variable to equal the admission cutoff score for open spots c_{0j}, as well as fixing the location dummy equal to 0:

$$\text{VA}_j = (\hat{\alpha}_j^t - \hat{\alpha}_0^t) + (\hat{\beta}_j^t - \hat{\beta}_0^t) \cdot c_{0j}.$$

In Panel (b) of Figure 8, we plot the value-added of each degree against its selectivity level. We find that more selective degrees offer higher value-added. In contrast to the degree valuations from Panel (a), we find that the average federal degree offers higher value than the outside option (i.e., positive value added). Moreover, targeted students obtain higher value-added than their non-targeted counterparts. This is likely driven by non-targeted students having outside options with high returns, such as attending similar programs in private universities.

8 The impact of affirmative action on winners and losers

In this section, we use the parameter estimates from our model to compute student assignments, as well as their potential outcomes under different AA schedules. We then compare the overall predicted income gains and losses for targeted and non-targeted individuals, and for the system as a whole.

8.1 Estimating counterfactual assignments and outcomes

The goal is to simulate student assignments for any given AA schedule ω by leveraging the rules of the centralized mechanism. We start by recovering the inputs of the matching function $\varphi(\Theta, q, \omega) = \mu$, namely student types, Θ, and degree capacities, q.

The first step is to compute the set of student types $\Theta = \bigcup_{i \in \mathcal{I}} \theta_i$, where $\theta_i = (\succ_i, s_i, t_i)$. We recover preferences, $\hat{\succ}_i$, by evaluating indirect utilities from the school choice model introduced in Section 6.1 using preference parameter estimates from Section 7. To recover the remaining inputs

to Θ, we assume no behavioral responses to the regulation. This allows us to recover inputs directly from the data. Specifically, we assume that the composition of applicants, $\mathcal{I}$, priority scores, s_i, and AA status, t_i, are fixed and invariant to changes in the AA schedule.[48] The main concern related to ruling out behavioral responses is whether the policy effects are well predicted by ignoring them.

The second step is to recover degree capacities, $q = \{q_0, ..., q_J\}$. This variable, as opposed to the inputs of Θ, is a policy choice that is decided together with the AA schedule. Since we are interested in learning about the consequences of AA, we keep this variable fixed and recover it from the observed data. After recovering $\hat{\Theta}$ and q, we simulate the matching function as $\varphi(\hat{\Theta}, q, \omega) = \hat{\mu}$.

The next step, after simulating student assignments using the matching function, is to compute the realized outcomes associated with these assignments. We define our object of interest for individual of type θ_i as $\mathbb{E}[\tilde{Y}_i(\hat{\mu}(\theta_i)) \mid \theta_i]$, where $\hat{\mu}$ denotes the matching function estimated above. Using Equations (3), (8), and (9), we parametrize this object as:

$$\begin{aligned}
\mathbb{E}[\tilde{Y}_i(\hat{\mu}(\theta_i)) \mid \theta_i] &= \sum_j \mathbb{1}\{\hat{\mu}(\theta_i) = j\} \cdot \mathbb{E}[\tilde{Y}_{ij} \mid \theta_i] \\
&= \sum_j \mathbb{1}\{\hat{\mu}(\theta_i) = j\} \cdot \left(\alpha_j^t + X_{ij}'\beta_j^t + \mathbb{E}[\varepsilon_{ij} \mid \theta_i]\right) \\
&= \sum_j \mathbb{1}\{\hat{\mu}(\theta_i) = j\} \cdot \left(\alpha_j^t + X_{ij}'\beta_j^t + \sum_{k=0}^{J} \psi_k^t(\eta_{ik} - \bar{\eta}) + \rho^t(\eta_{ij} - \bar{\eta})\right) \qquad (14)
\end{aligned}$$

This parametrization assumes that potential outcomes are invariant to the AA regulation. There are two implications that we deem important to discuss. The first is that this assumption rules out, for instance, peer effects in the production function of degrees, as well as changes in the value of Y_{ij} coming from stigmatization, or from a reduction of the signaling value of degrees as a result of the AA regulation. The second implication of this assumption is that the value of the outside option is fixed. This assumption would be violated if non-targeted students displaced from public institutions were to crowd out other students from private institutions. This effect could potentially create a crowding-out cascade extending throughout the whole higher education system. However, most private institutions, including those comparable to federal universities, are far from being capacity constrained. Indeed, in our setting, only 1.9% of degree programs in the top 10% of selectivity are capacity constrained, thus crowding-out concerns are of second order.[49]

We use Equation (14) together with the selection, and potential outcome parameter estimates, to simulate student allocations for a given schedule ω. We describe the simulation procedure in H .

[48] The assumption of no behavioral responses to priority scores is consistent with findings by Francis and Tannuri-Pianto (2012b) and (Estevan et al., 2018) in public institutions in Brazil. They study AA regulations that increased the representation of disadvantaged groups at the University of Brasilia and the University of Campinas, respectively, and find no evidence of behavioral reactions regarding examination preparation effort.

[49] We calculate these statistics based on the degree capacities reported by private degrees, and the number of students actually enrolled at the degree. We measure selectivity as the average test score of students admitted in 2016.

8.2 Counterfactual admission thresholds

Next, we perform counterfactual exercises where we vary the share of reserved seats ω from 0 to 100%. For each AA schedule ω, we compute students' allocations and estimate their expected potential outcome in terms of predicted income as described in Section 6.4. We proceed by computing the admission thresholds associated with counterfactual allocations.

For expositional purposes, in Figure 9 we show the distribution of admission cutoff scores across degrees for reserved and open spots under two different counterfactual scenarios. The first one is a counterfactual similar to the current policy that reserves 50% of the seats to targeted students ($\omega = 0.5$). The second one corresponds to a laissez-faire situation in which there is no AA at all ($\omega = 0$). Under this scenario, all students compete over the same spots and thus face the same admission thresholds.

We find that under the AA counterfactual, where $\omega = 0.5$, cutoffs for open spots students are substantially higher than those reserved for targeted ones. This implies that targeted students can get admitted into selective degrees with much lower scores. In the absence of AA, when $\omega = 0$, the distribution of cutoff scores (dashed black line) becomes uniform across AA types and closer to the distribution of cutoff scores for open seats in the presence of AA. Appendix Figures A.8 and A.9 present quantile-quantile plots that show how admission thresholds change for open and reserved spots under different AA schedules ranging from 0 to 100%. The overall results suggest that, by removing the AA program, admission into selective degrees becomes much harder for targeted students but not substantially easier for non-targeted ones.

8.3 Counterfactual outcomes

In Figure 10 we show the expected predicted income for targeted and non-targeted students under each of the counterfactuals. The red dashed line denote a counterfactual scenario without AA, while the blue solid line indicates a counterfactual with 50% reserved seats. The grey line in the background (and measured by the right-hand side axis) shows the distribution of students over ENEM scores. In Panel (a) of Figure 10 we show that the AA program induces large gains on targeted individuals, especially on those students with high scores that now can access more selective degrees with higher value-added. As expected, targeted individuals at the very top of the score distribution are not affected by the AA regulation as all degrees are within reach even in the absence of the regulation.

These gains for targeted students, however, come at the cost of displacing non-targeted students from those very selective degrees. In Panel (b) of Figure 10 we show that, even though non-targeted students are worse off under the AA policy, their losses are small relative to the gains of targeted individuals. This is mostly explained by the fact that non-targeted individuals have lower chances of ending up in the outside option and that their outside option has a relatively higher value-added.

Next, we estimate the aggregate effects of AA on predicted income. Let $\Delta_t(\omega) = \Delta_t(\omega, 0)$ denote

the aggregate gains for AA group t of moving from no AA to an ω AA schedule, as indicated by Equation (2). The overall aggregate gains over targeted and non-targeted individuals are defined as:

$$\Delta(\omega, \lambda) \quad = \quad \Delta_0(\omega) + \lambda \Delta_1(\omega) \tag{15}$$

where λ denote the welfare weights capturing society's concerns for fairness with respect to the targeted group.

In Figure 11, we present the group-specific and overall gains in terms of predicted income. We normalize these gains in terms of the aggregated predicted income for AA group t when $\omega = 0$. The blue and red lines denote the gains and losses for targeted and non-targeted individuals, respectively. The grey line represents the normalized overall aggregate change using equal welfare weights across groups, that is $\lambda = 1$. We find that, at the current affirmative action schedule of $\omega = 0.5$, the average targeted individual sees a 1.16% increase in their predicted income. In comparison, the average non-targeted student affected by the policy faces a 0.93% reduction in their predicted income. As a result, the predicted income of the average student across AA groups increases by 0.1%. These results are consistent with those presented in Figure 10. Although the gains for targeted students are greater than the losses experienced by non-targeted students, these gains are aggregated over fewer targeted individuals. Thus, the average student in the population experiences a 0.1% increase in their predicted income.

9 Discussion and Conclusion

In this paper, we study the distributional consequences of affirmative action policies in centralized admission systems. In addition to providing a controlled setting, analyzing AA in centralized systems is important because of a rising trend in the number of countries adopting such systems. Today, over 45 countries organize their high education admissions using centralized systems (Neilson, 2020). Several of these systems rely heavily on standardized test scores to assign students to degrees. Because test scores are strongly correlated with socioeconomic status, reliance on such centralized systems will result in highly segregated student bodies. In this context, AA is a relevant policy lever for increasing diversity and the representation of minorities.

To understand the consequences of AA, we develop and estimate a model that links preferences over degrees with the potential outcomes of attending each degree. We find that the AA policy increases the enrollment of targeted students in more selective degrees. In terms of outcomes, we focus on the impact of the policy on academic progress and their implied gains and losses on income. In contrast with the mismatch hypothesis, we find that targeted students benefit from the AA policy in terms of academic persistence and implied income gains. Moreover, our results suggest that the benefits for targeted students are larger than the costs imposed on non-targeted students. However, because the number of affected targeted students is smaller than the number of

affected non-targeted students, the net effects on predicted income are close to zero. These results suggest that introducing AA can increase overall equity without affecting the overall efficiency of the education system.

One key element that we introduced in this paper is that targeted and non-targeted students may face different outside options. This implies that non-targeted students who are displaced from public institutions as a result of the AA policy, may not be significantly harmed because they have good outside options available to them in the private sector. In our context, it is important to note that 15% of targeted students in the top quartile of ENEM performance have access to student loans that allow them to study in the private sector. This number is relatively large compared to the 4.3% of non-targeted students in the top quartile of ENEM performance who have access to loans. In the absence of student loans, targeted students would have different outside options, which would affect our efficiency estimates. We conclude that AA cannot be studied in isolation. The results presented in this may include interactions with the effects of other programs potentially targeting the same population. Exploring this interaction is material for future research.

Finally, it is worth mentioning that this paper focuses on the first-order efficiency trade-off between targeted and non-targeted individuals. However, AA can also affect other education margins not considered in this paper. For example, universities can affect student outcomes through peer effects, and influence intergroup attitudes by increasing diversity. In addition, a more diverse student body composition may change the production function of degrees by affecting students' academic outcomes, social behavior, and preferences. Understanding how these different margins interact with the direct distributional effects of AA is key to understanding its overall role in shaping a higher education sector that fosters social mobility and promotes a healthy democracy.

References

Abdulkadiroglu, A. (2005). College admissions with affirmative action. *International Journal of Game Theory 109*(4), 535–549.

Abdulkadiroglu, A., J. Angrist, Y. Narita, and P. Pathak (2017). Research design meets market design: Using centralized assignment for impact evaluation. *Econometrica 85*(5), 1373–1432.

Abdulkadiroglu, A., P. Pathak, J. Schellenberg, and C. R. Walters (2020). Do parents value school effectiveness? Forthcoming, American Economic Review.

Abdulkadiroglu, A. and T. Sonmez (2003). School choice: A mechanism design approach. *The American Economic Review 93*, 729 – 747.

Agarwal, N. and P. Somaini (2018). Demand analysis using strategic reports: An application to a school choice mechanism. *Econometrica 86*(2), 391–444.

Agarwal, N. and P. Somaini (2020). Revealed preference analysis of school choice models. *Annual Review of Economics 12*(1), 471–501.

Alon, S. (2015). *Race, Class, and Affirmative Action.* Russell Sage Foundation.

Alon, S. and M. Tienda (2005). Assessing the "mismatch" hypothesis: Differences in college graduation rates by institutional selectivity. *Sociology of Education 78*(4), 294–315.

Angrist, J. and M. Rokkanen (2015). Wanna get away? regression discontinuity estimation of exam school effects away from the cutoff,. *Journal of the American Statistical Association 110*(512), 1331–1344.

Arcidiacono, P. (2005). Affirmative action in higher education: How do admission and financial aid rules affect future earnings? *Econometrica 73*(5), 1477–1524.

Arcidiacono, P., E. M. Aucejo, H. Fang, and K. I. Spenner (2011). Does affirmative action lead to mismatch? a new test and evidence. *Quantitative Economics 2*(3), 303–333.

Arcidiacono, P. and M. Lovenheim (2016). Affirmative action and the quality-fit trade-off. *Journal of Economic Literature 54*(1), 3–51.

Athey, S., R. Chetty, G. Imbens, and K. Hyunseung (2019). The Surrogate Index: Combining Short-Term Proxies to Estimate Long-Term Treatment Effects More Rapidly and Precisely. Working paper no. 26463, NBER.

Aygun, O. and I. Bó (2021, August). College admission with multidimensional privileges: The brazilian affirmative action case. *American Economic Journal: Microeconomics 13*(3), 1–28.

Ayres, I. and R. Brooks (2004). Does affirmative action reduce the number of black lawyers. *Stan. L. Rev. 57*, 1807.

Azevedo, E. M. and J. D. Leshno (2016). A supply and demand framework for two-sided matching markets. *Journal of Political Economy 124*(5), 1235–1268.

Bagde, S., D. Epple, and L. Taylor (2016). Does affirmative action work? caste, gender, college quality, and academic success in india. *American Economic Review 106*(6), 1495–1521.

Barahona, N., C. Dobbin, K. Kanopka, and S. Otero (2021). Blues books and the position effect. Unpublished manuscript, Stanford University.

Bertanha, M. and G. Imbens (2019). External validity in fuzzy regression discontinuity designs. *Journal of Business & Economic Statistics 0*(0), 1–39.

Bjorklund, A. A. and R. Moffitt (1987). The estimation of wage gains and welfare gains in self-selection models. *The Review of Economics and Statistics 69*(1), 42–49.

Black, S. E., J. T. Denning, and J. Rothstein (2020). Winners and Losers? The Effect of Gaining and Losing Access to Selective Colleges on Education and Labor Market Outcomes. Working Paper 26821, NBER.

Bleemer, Z. (2021a, 09). Distributional Effects of Race-Blind Affirmative Action. *The Quarterly Journal of Economics 0*(0), 0–0. qjab027.

Bleemer, Z. (2021b). Top Percent Policies and the Return to Postsecondary Selectivity. Working paper, UC Berkeley.

Bó, I. and R. Hakimov (2019, 07). Iterative Versus Standard Deferred Acceptance: Experimental Evidence. *The Economic Journal 130*(626), 356–392.

Bodoh-Creed, A. and B. Hickman (2019). Pre-college human capital investment and affirmative action: A structural policy analysis of u.s. college admissions. Working paper, Haas School of Business, University of California, Berkeley.

Bound, J., M. F. Lovenheim, and S. Turner (2010, July). Why have college completion rates declined? an analysis of changing student preparation and collegiate resources. *American Economic Journal: Applied Economics 2*(3), 129–57.

Brinch, C. N., M. Mogstad, and M. Wiswall (2017). Beyond late with a discrete instrument. *Journal of Political Economy 125*(4), 985–1039.

Calonico, S., M. D. Cattaneo, and R. Titiunik (2014). Robust nonparametric confidence intervals for regression-discontinuity designs. *Econometrica 82*(6), 2295–2326.

Calsamiglia, C., G. Haeringer, and F. Klijn (2010, September). Constrained school choice: An experimental study. *American Economic Review 100*(4), 1860–74.

Canaan, S. and P. Mouganie (2017). Returns to education quality for low-skilled students: Evidence from a discontinuity. *Journal of Labor Economics 91*(4), 717–724.

Chambers, D. L., T. T. Clydesdale, W. C. Kidder, and R. O. Lempert (2004). The real impact of eliminating affirmative action in american law schools: An empirical critique of richard sander's study. *Stan. L. Rev. 57*, 1855.

Chetty, R., J. N. Friedman, E. Saez, N. Turner, and D. Yagan (2020, 02). Income Segregation and Intergenerational Mobility Across Colleges in the United States*. *The Quarterly Journal of Economics 135*(3), 1567–1633.

de Chaisemartin, C. and L. Behaghel (2020). Estimating the effect of treatments allocated by randomized waiting lists. *Econometrica 88*(4), 1453–1477.

Dillon, E. W. and J. A. Smith (2017). Determinants of the Match between Student Ability and College Quality. *Journal of Labor Economics 35*(1), 45–66.

Dillon, E. W. and J. A. Smith (2020). The Consequences of Academic Match between Students and Colleges. *Journal of Human Resources 55*(3), 767–808.

Dobbin, C., N. Barahona, and S. Otero (2021). The equilibrium effects of subsidized student loan. Working paper, Stanford University.

Dong, Y. and A. Lewbel (2015). Identifying the effects of changing the policy threshold in regression discontinuity models. *Review of Economics and Statistics 97*(5), 1081–1092.

Dubin, J. A. and D. L. McFadden (1984). An econometric analysis of residential electric appliance holdings and consumption. *Econometrica 52*(2), 345–362.

Dur, U., S. D. Kominers, P. A. Pathak, and T. Sonmez (2018). Reserve design: Unintended consequences and the demise of boston's walk zones. *Journal of Political Economy 126*(6), 2457–2479.

Dur, U., P. A. Pathak, and T. Sonmez (2020). Explicit vs. statistical targeting in affirmative action: Theory and evidence from chicago's exam schools. *Journal of Economic Theory 187*, 104996.

Ellison, G. and P. A. Pathak (2021, March). The efficiency of race-neutral alternatives to race-based affirmative action: Evidence from chicago's exam schools. *American Economic Review 111*(3), 943–75.

Engbom, N., G. Gonzaga, C. Moser, and R. Olivieri (2021). Earnings inequality and dynamics in the presence of informality: The case of brazil. Working paper, New York University.

Estevan, F., T. Gall, and L.-P. Morin (2018, 05). Redistribution Without Distortion: Evidence from an Affirmative Action Programme at a Large Brazilian University. *The Economic Journal 129*(619), 1182–1220.

Fack, G., J. Grenet, and Y. He (2019, April). Beyond truth-telling: Preference estimation with centralized school choice and college admissions. *American Economic Review 109*(4), 1486–1529.

Fischer, M. J. and D. S. Massey (2007). The effects of affirmative action in higher education. *Social Science Research 36*(2), 531 – 549.

Francis, A. and M. Tannuri-Pianto (2012a). The redistributive equity of affirmative action: Exploring the role of race, socioeconomic status, and gender in college admissions. *Economics of Education Review 31*(1), 45 – 55.

Francis, A. and M. Tannuri-Pianto (2012b). Using brazil's racial continuum to examine the short-term effects of affirmative action in higher education. *The Journal of Human Resources 47*(3), 754 – 784.

Francis-Tan, A. and M. Tannuri-Pianto (2018). Black movement: Using discontinuities in admissions to study the effects of college quality and affirmative action. *Journal of Development Economics 135*, 97–116.

Gale, D. E. and L. S. Shapley (1962). College admissions and the stability of marriage. *American Mathematical Monthly 69*(1), 9–15.

Haeringer, G. and F. Klijn (2009). Constrained school choice. *Journal of Economic Theory 144*(5), 1921–1947.

Hastings, J. S., C. A. Neilson, and S. D. Zimmerman (2013, July). Are Some Degrees Worth More than Others? Evidence from college admission cutoffs in Chile. Working Paper 19241, National Bureau of Economic Research.

Heckman, J. J. (1979). Sample selection bias as a specification error. *Econometrica 47*(1), 153–161.

Hoekstra, M. (2009). The effect of attending the flagship state university on earnings: A discontinuity-based approach. *The Review of Economics and Statistics 91*(4), 717–724.

Howell, J. S. (2010). Assessing the impact of eliminating affirmative action in higher education. *Journal of Labor Economics 28*(1), 113–166.

Jenkins, L. D. and M. S. Moses (2017). Global: Affirmative Action Initiatives Around the World. In *Understanding Global Higher Education. Global Perspectives on Higher Education*. SensePublishers, Rotterdam.

Jia, R. and H. Li (2021). Just above the exam cutoff score: Elite college admission and wages in china. *Journal of Public Economics 196*, 104371.

Kapor, A. (2020). Distributional effects of race-blind affirmative action. Working paper, Princeton University.

Kirkeboen, L. J., E. Leuven, and M. Mogstad (2016). Field of study, earnings, and self-selection. *The Quarterly Journal of Economics 131*(3), 1057.

Kline, P. and C. R. Walters (2016, 07). Evaluating Public Programs with Close Substitutes: The Case of Head Start. *The Quarterly Journal of Economics 131*(4), 1795–1848.

Kline, P. and C. R. Walters (2019). On heckits, late, and numerical equivalence. *Econometrica 87*(2), 677–696.

Larroucau, T. and I. Rios (2021). Dynamic college admissions and the determinants of students' college retention. Technical report, Working Paper.

Mello, U. (2021). Affirmative action, centralized admissions and access of low-income students to higher education. Forthcoming, American Economic Journal: Economic Policy.

Neilson, C. (2020). The rise of centralized choice and assignment mechanisms in education markets around the world. Working paper, Princeton University.

Pathak, P. A., A. Rees-Jones, and T. Sonmez (2020). Reversing reserves. Working Paper 26767, NBER.

Rokkanen, M. (2015). Exam schools, ability, and the effects of affirmative action: Latent factor extrapolation in the regression discontinuity design. Working paper, Columbia University.

Rothstein, J. and A. H. Yoon (2008). Affirmative action in law school admissions: What do racial preferences do? Technical report, National Bureau of Economic Research.

Roy, A. D. (1951, 06). Some Thoughts on the Distribution of Earnings. *Oxford Economic Papers 3*(2), 135–146.

Sales, S. and M. S. Moses (2014). Brazil: Enhancing opportunity and justice through new affirmative action policies for black and mixed-race students. In *Affirmative Action Matters Creating opportunities for students around the world*. Routledge.

Sander, R. H. (2004). A systemic analysis of affirmative action in american law schools. *Stanford Law Review 57*(2), 367–483.

Sekhri, S. (2020, July). Prestige matters: Wage premium and value addition in elite colleges. *American Economic Journal: Applied Economics 12*(3), 207–25.

Singer, P. (2011). *Practical Ethics*. Cambridge University Press; 3rd edition.

Sonmez, T. and M. B. Yenmez (2021). Affirmative action in india via vertical, horizontal , and overlapping reservations. Working paper, Boston College.

Tincani, M. M., F. Kosse, and E. Miglino (2021). Subjective beliefs and inclusion policies: Evidence from college admissions. Working paper, UCL.

Vytlacil, E. (2002). Independence, monotonicity, and latent index models: An equivalence result. *Econometrica 70*(1), 331–341.

Zimmerman, S. D. (2014). The returns to college admission for academically marginal students. *Journal of Labor Economics 32*(4), 711 – 754.

Zimmerman, S. D. (2019). Elite colleges and upward mobility to top jobs and top incomes. *American Economic Review 109*(1), 1–47.

Figure 1: Affirmative Action Regulation

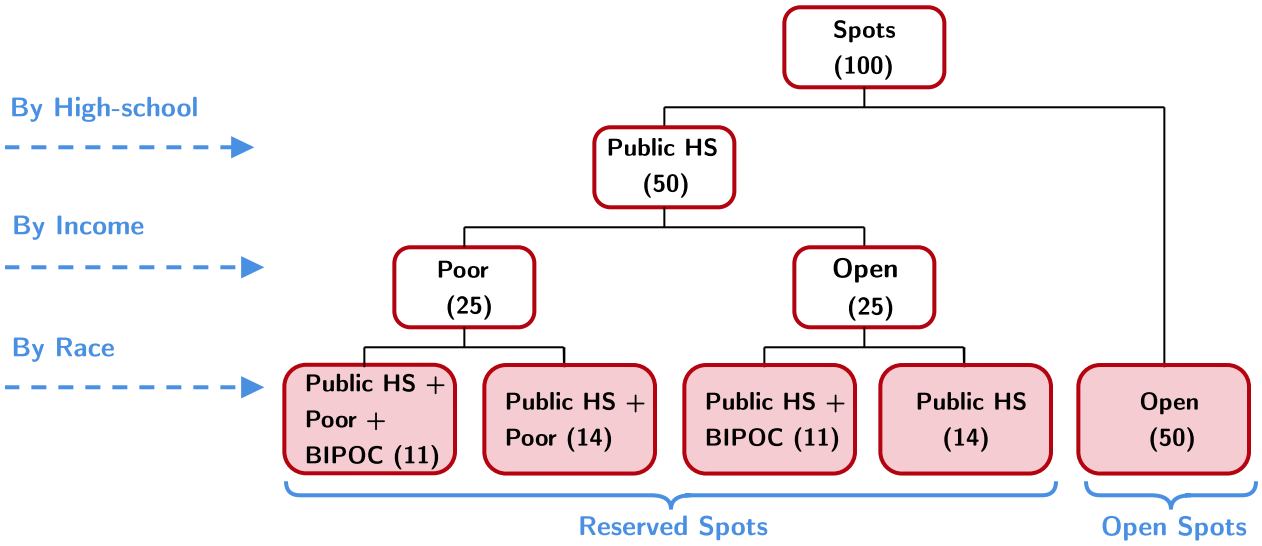

Notes: This figure describes the affirmative action policy. Under the regulation, for every 100 university spots, 50 are affirmative action spots for students who attended public high schools. Those affirmative actions spots are then divided equally by income, with 25 going to poor students and 25 to non-poor students. Finally, each group of 25 spots is distributed to reflect the proportion of non-white individuals in the population of a given state. This example uses a proportion of 54%, which is the combined share of black, brown, and indigenous people in the state of Minas Gerais, as reported by the Brazilian National Bureau of Statistics (IBGE) in 2012.

Figure 2: ENEM score distributions by targeted status

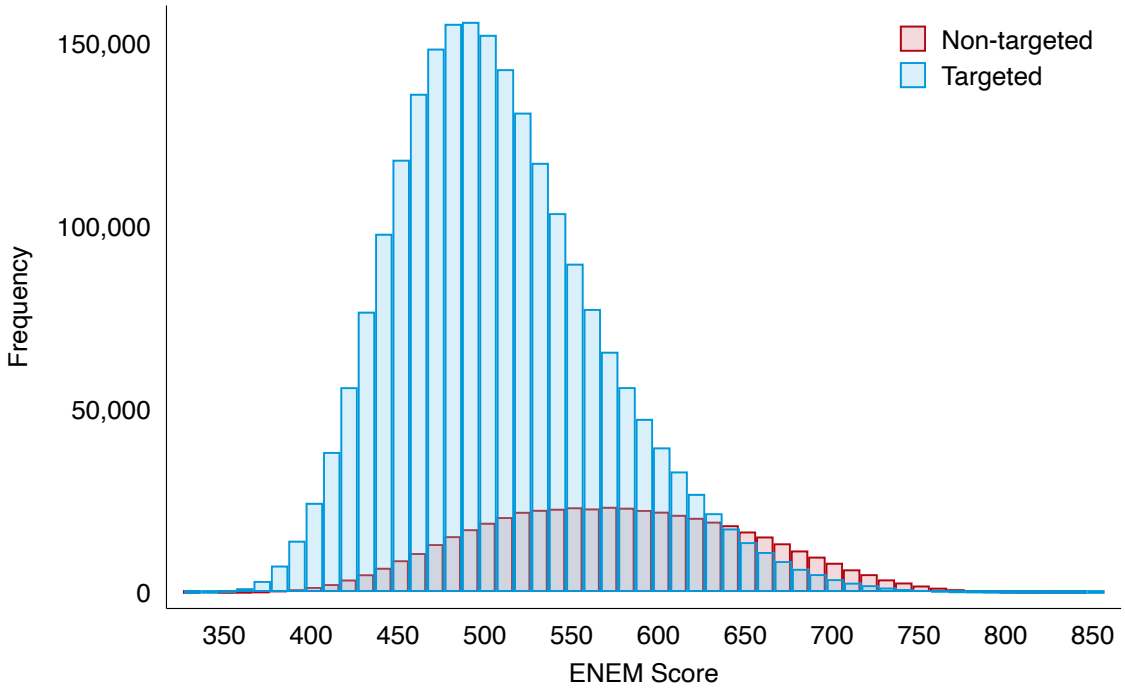

Notes: This figure shows kernel density plots of the average ENEM score distribution of targeted and non-targeted individuals. The sample is the universe of ENEM takers who had positive test scores in the 2015 ENEM test. The average test score include math, language, natural science, and social science. The average score is 504 and the standard deviation is 66 points. Targeted students are defined as those who are eligible for any of the affirmative action vacancies.

Figure 3: Student body composition and attendance to federal institutions

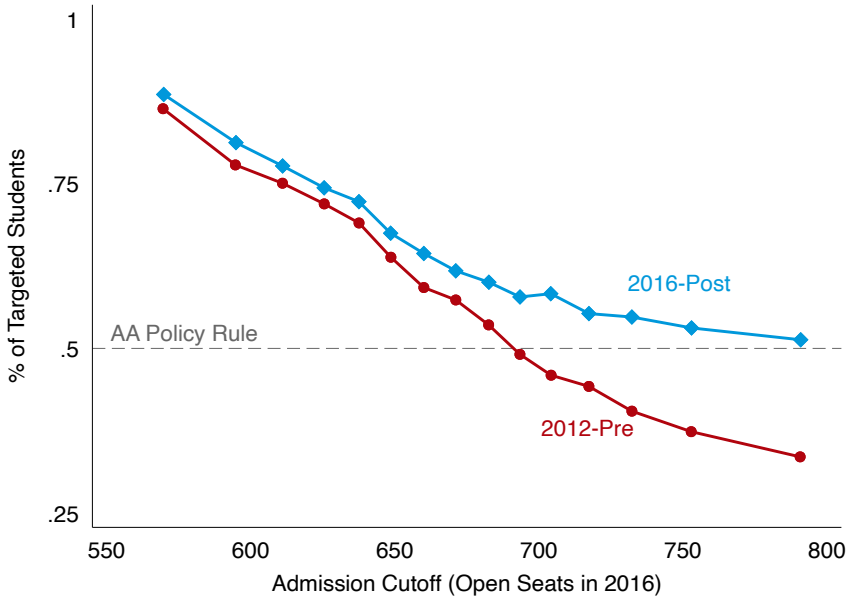

(a) Degree composition

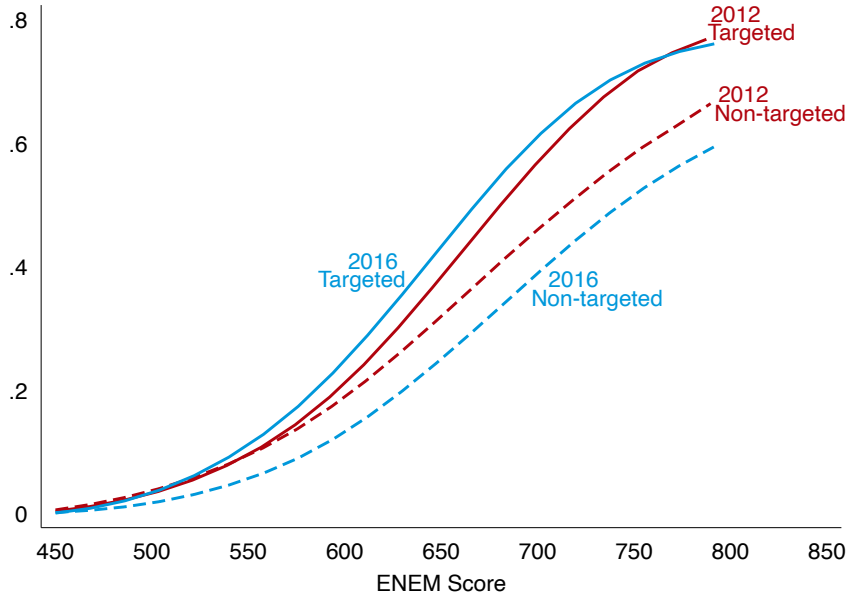

(b) Admission into federal institution

Notes: Panel (a) Figure presents the average share of incoming targeted students in federal institutions by degree selectivity as defined as the admission cutoff of open spots in 2016 SISU. We consider individuals starting a degree program in the first semester of each year. An observation is a degree and shift tuple. We weight each observation by the size of the incoming cohort. We keep degree programs that exist both in 2012 and 2016 and that participated in 2016 SISU. Our sample covers 94% of the total enrollment in federal institutions in 2012 and 2016. The dashed line depicts the 50% AA policy rule mandated by the regulation. Panel (b) plots the admission probability into federal institutions by student test scores. An observation is a student taking ENEM in a who just graduated from high school. The solid lines depict targeted students, and the dashed lines represent non-targeted students. The red lines refer to ENEM valid for academic year 2012, and the blue lines refer to ENEM valid for academic year 2016.

Figure 4: Regression discontinuity estimates

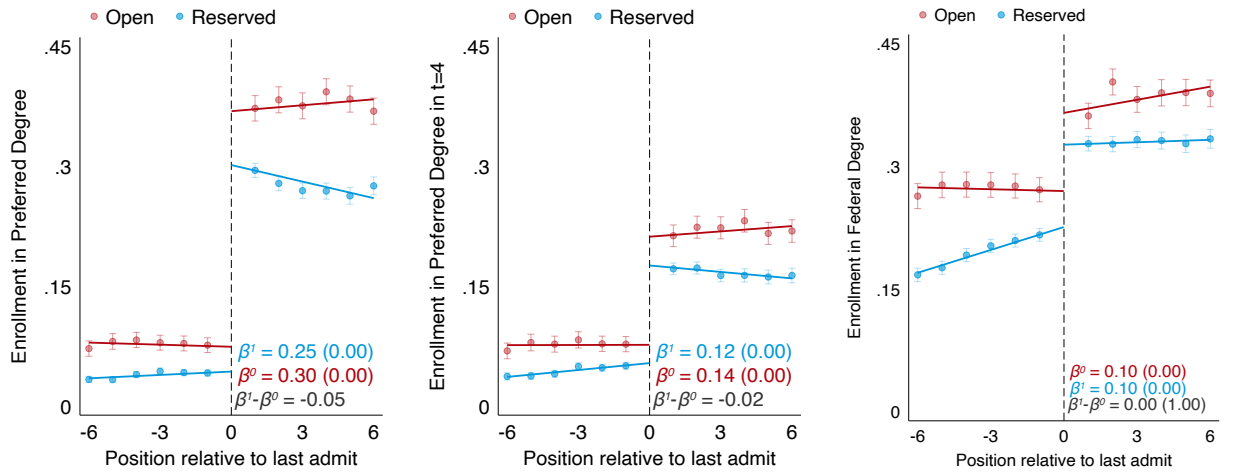

(a) Enrolled in preferred degree, $t = 1$ (b) Enrolled in preferred degree, $t = 4$ (c) Enrolled in federal degree, $t = 4$

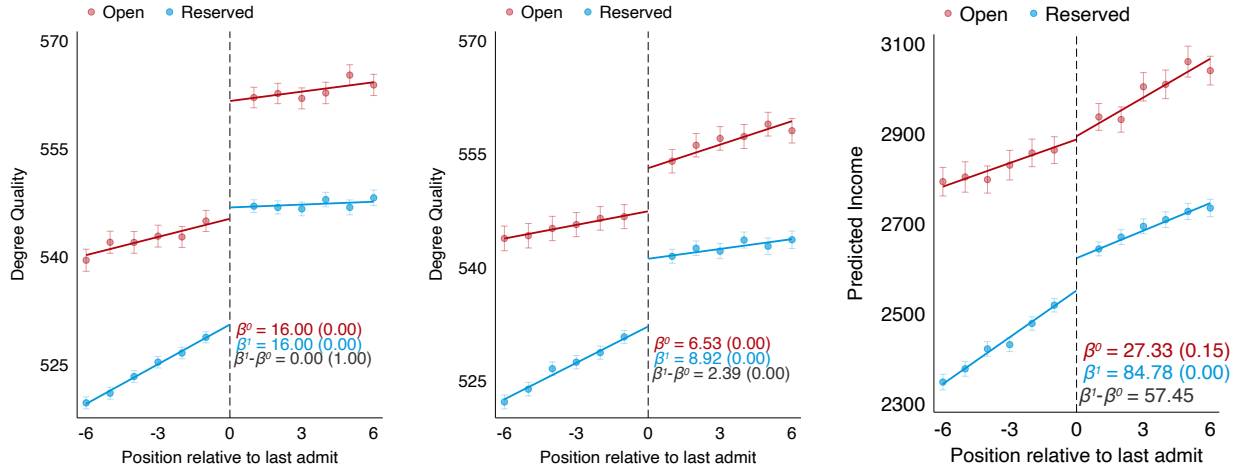

(d) Degree quality in $t = 1$ (e) Degree quality in $t = 4$ (f) Predicted income

Notes: This figure shows the regression discontinuity estimates for a number of academic and labor market outcomes, and time periods. We pool all degree programs and center them around the last admitted individual. For simpler presentation, we also pool individuals across all four type of reserved seats, and present their estimates as a single group. The red and blue lines show the outcomes for students applying for an open seat and a reserved seat respectively. The confidence intervals represent the standard error of the outcome mean. The coefficients at the bottom of the figure indicate the treatment effect for each of the groups, as well as the difference between these coefficients. P-values are reported in parentheses. Standard errors of the estimates are clustered at the waitlist level.

Figure 5: Conceptual Framework

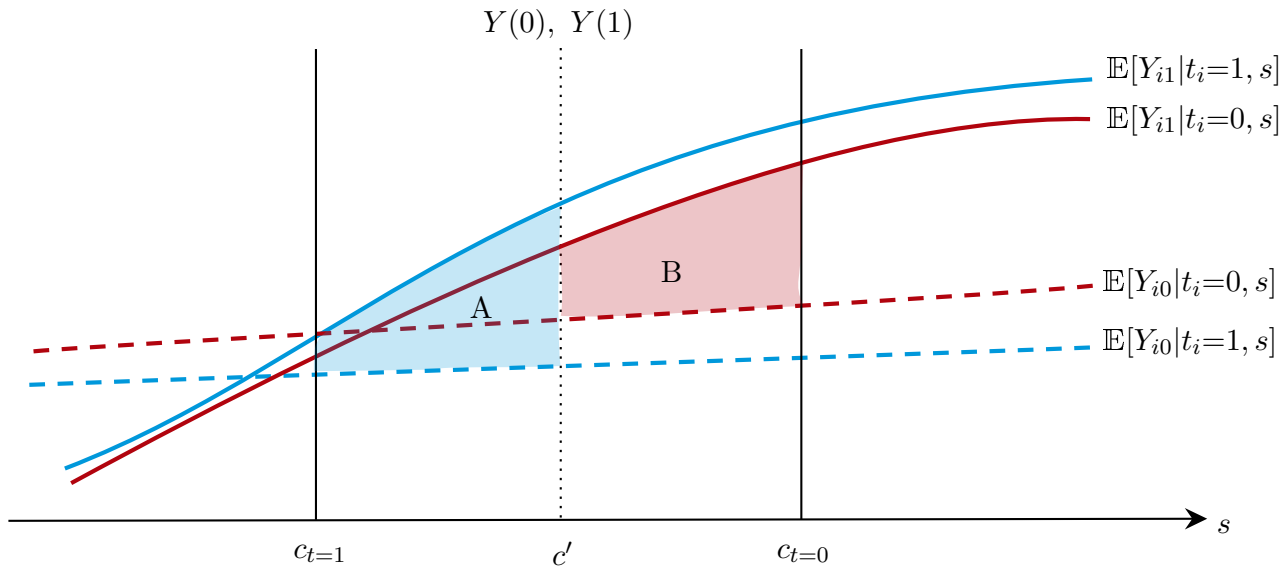

Notes: This figure presents the distributional impacts of an affirmative action policy in a centralized mechanisms. Each of the lines denote the mean potential outcome for targeted (blue line) and non-targeted (yellow line) individuals. While the solid line present the expected outcome of attending the selective degree, the dashed line presents the expected outcome of attending the outside option.

Figure 6: Score shifters

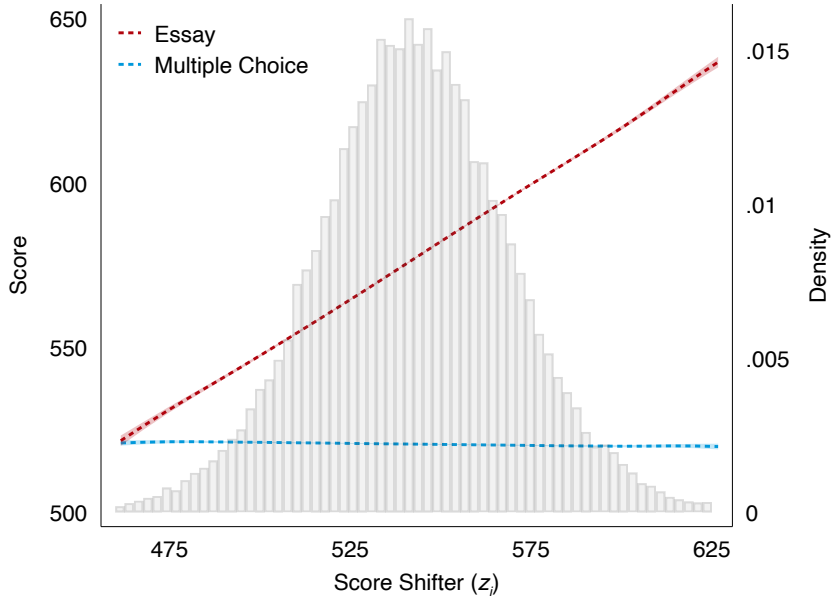

(a) Grader strictness

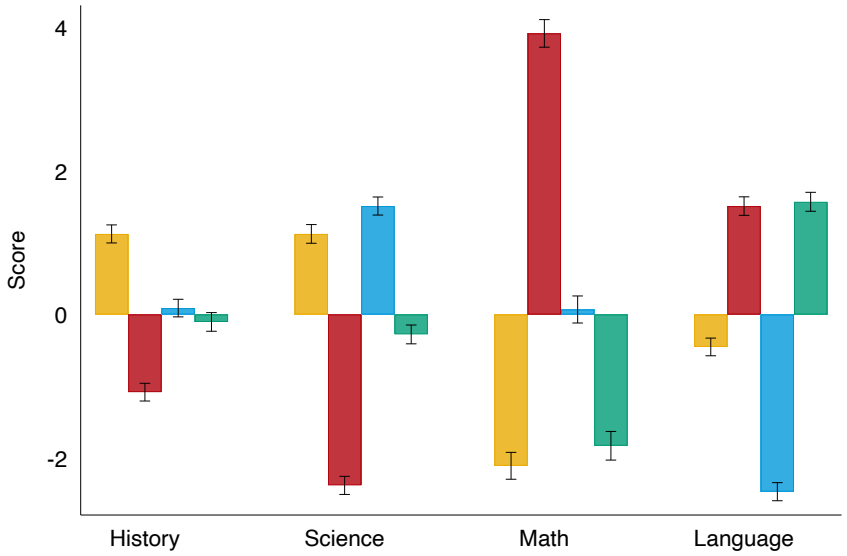

(b) Booklet difficulty

Notes: Panel (a) shows the relationship between the essay score shifter and test scores. The red line shows a strong first stage correlation between essay scores and the essay score shifter. The blue line serves as a balance test, and shows no correlation between the essay score shifter and the average test score in the multiple choice components of the test. The dashed line presents a local linear regression with a second order polynomial. The solid lines represent a 95% confidence bands. The histogram in the background reports the distribution of the shifter. Panel (b) reports the difference in average test scores across examination booklets for each of the different components of the ENEM national exam.

Figure 7: Score shifter, reduced form

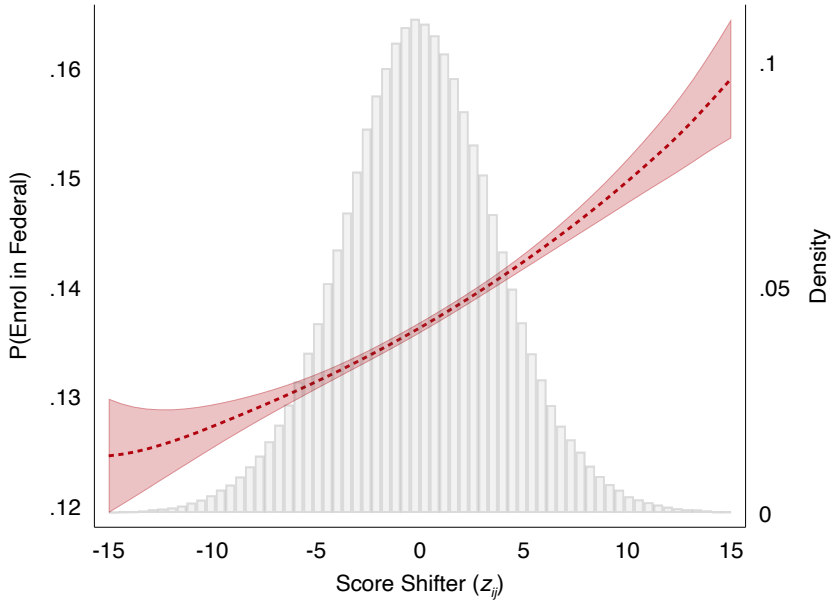

(a) Enrollment in Federal Institution in 2016

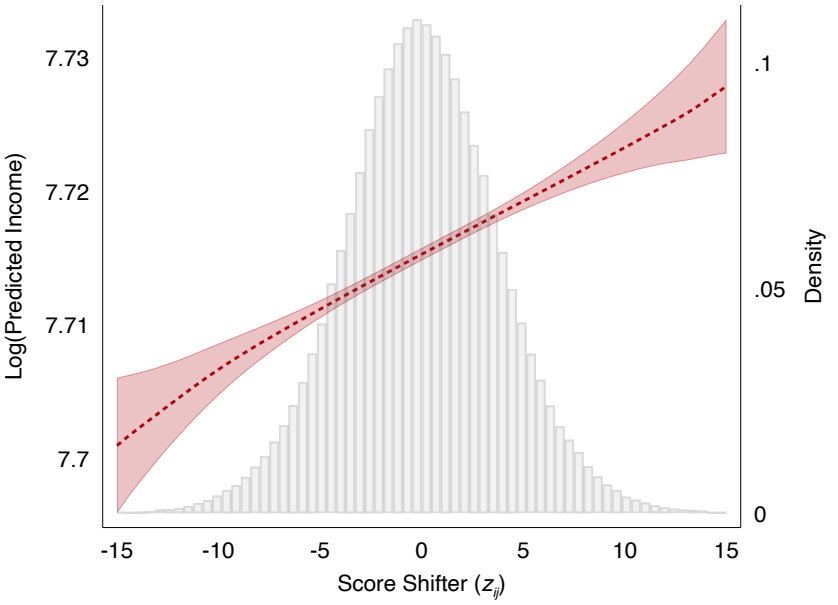

(b) Log(Predicted Income)

Notes: The histogram in the background reports the distribution of the shifter z_i. Panel (a) shows the reduced-form between the score shifter leniency and the probability of attending a federal institution during the year after taking the exam. Panel (b) reports the relationship between the score shifter and log of predicted income based on college attainment between 2016 and 2019 (see Section 6.4). The dashed line presents a local linear regression with a second order polynomial. The solid lines represent a 95% confidence bands.

Figure 8: Degree valuations and Value Added

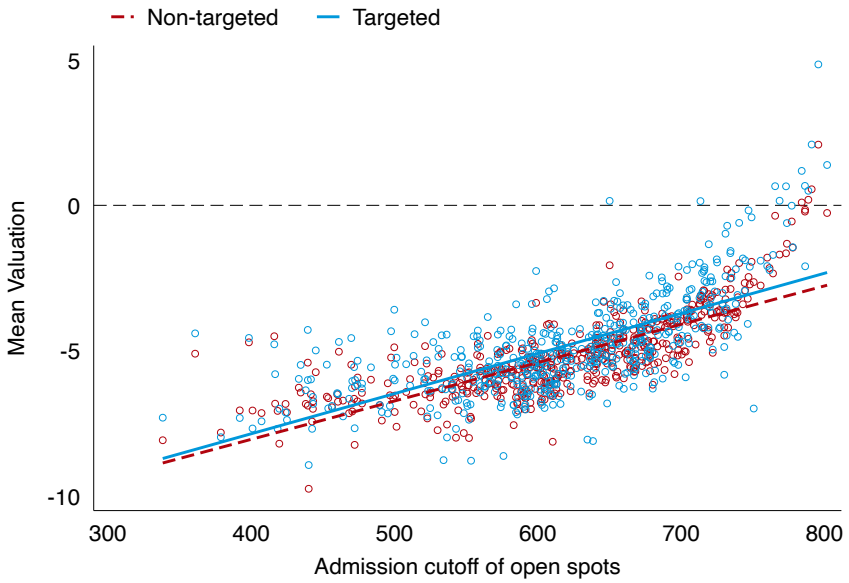

(a) Mean valuation and degree selectivity

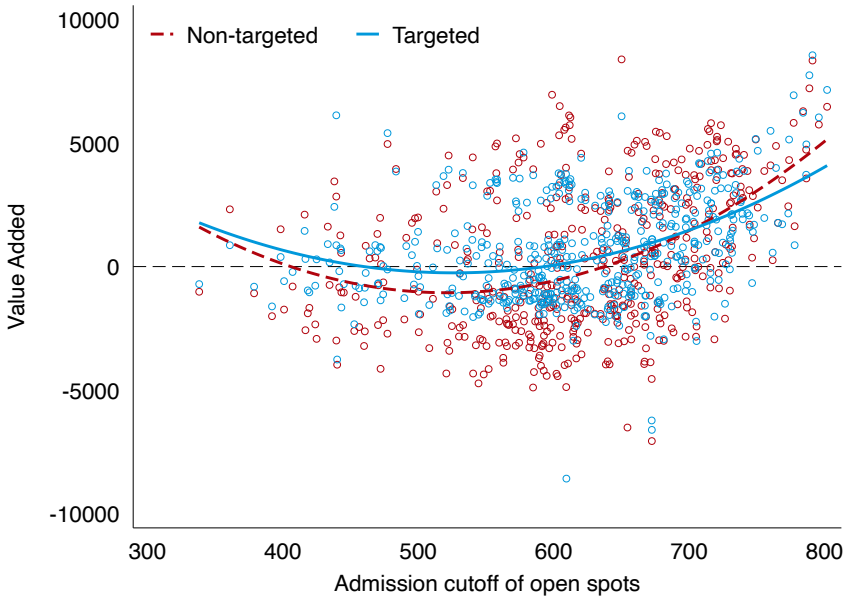

(b) Value Added and degree selectivity

Notes: This figure present the mean valuation and value added of degrees based on the parameter estimates. Panel (a) shows the relationship between degree selectivity and students' mean valuation for them. Panel (b) shows the relationship between degree selectivity and degree's value added. In both panels the x-axis denotes degree selectivity as measured by the admission cutoff of open seats c_{0j}.

Figure 9: Distribution of cutoff scores

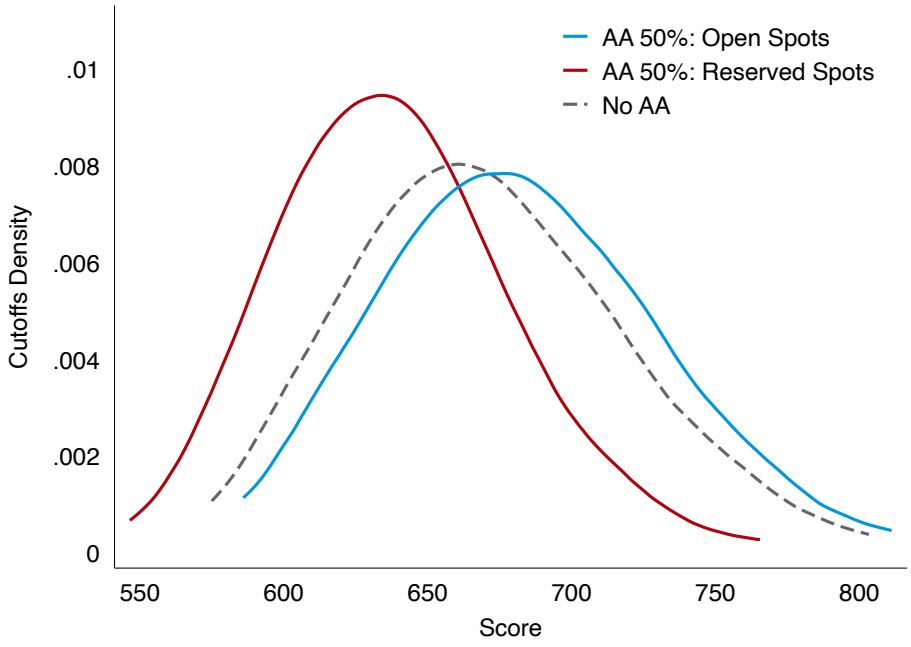

Notes: This figure shows the distribution of equilibrium cutoffs for open and reserved spots under different affirmative action schedules. The blue line shows the distribution of admission cutoffs for open spots when $\omega = 0.5$. The red line shows the distribution of admission cutoffs for reserved spots when $\omega = 0.5$. The dashed black line shows the distribution of admission cutoffs faced by all students when $\omega = 0$.

Figure 10: Students' income with and without affirmative action

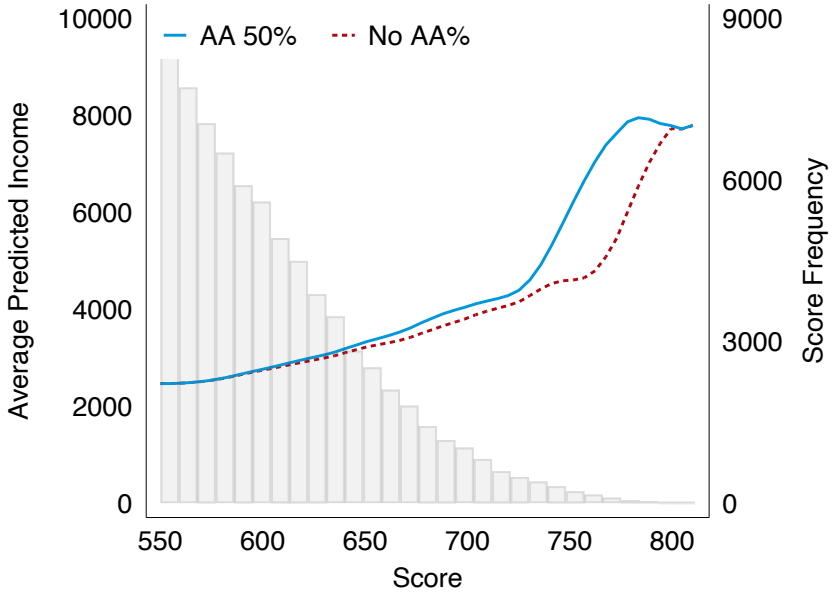

(a) Targeted students

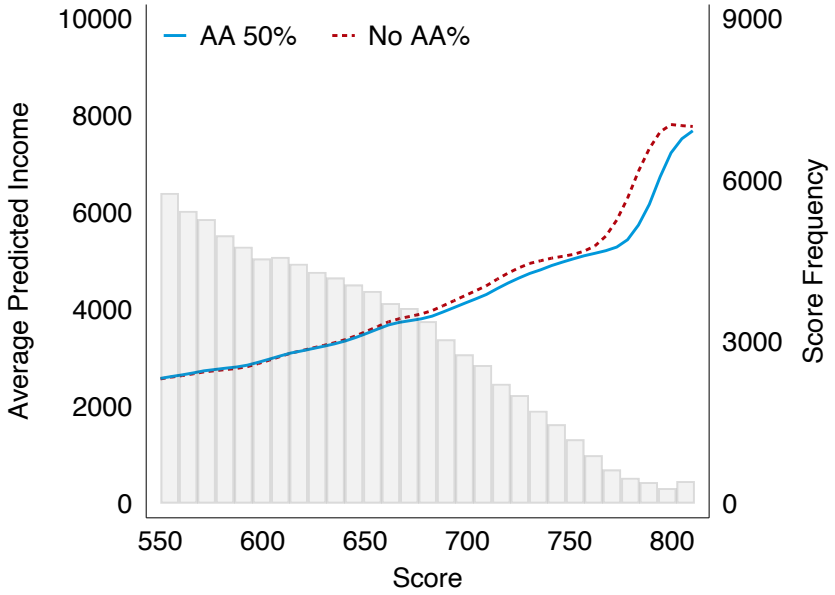

(b) Non-targeted students

Notes: This figure shows the expected outcome for targeted and non-targeted students with and without affirmative action across the score distribution. The blue line represents the expected outcome when $\omega = 0$, and the red line represents the outcome when $\omega = 0.5$.

Figure 11: Predicted income under different affirmative action schedules

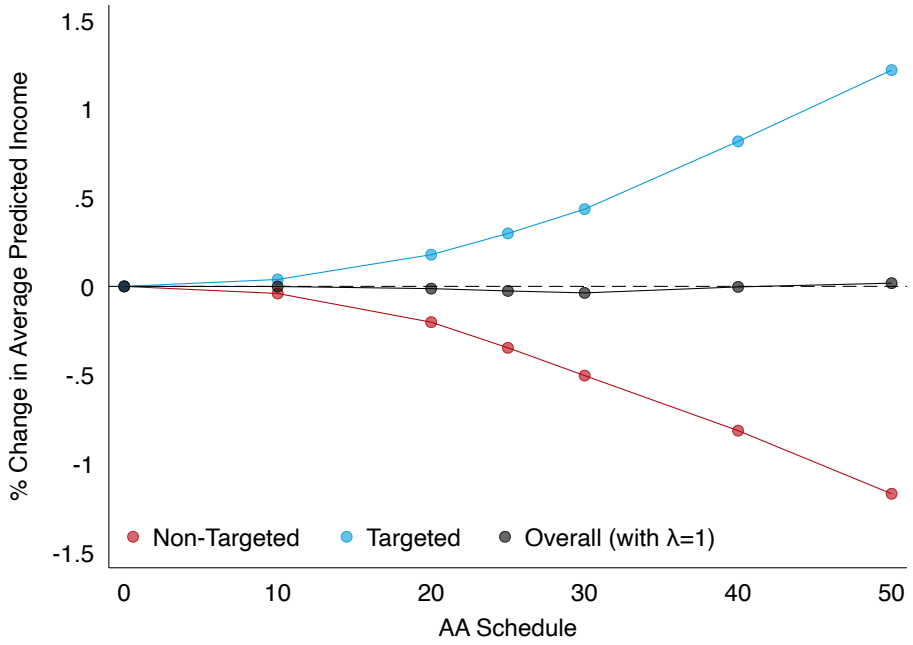

Notes: This figure shows the overall gains and losses of affirmative action in terms of predicted income. We normalize the outcome with respect to the aggregate predicted income in the absence of AA. We perform a separate normalization for each of the groups: targeted, non-targeted and overall. The overall label denotes the outcome when individuals across both groups are given equal weights.

Table 1: SISU Applicants

	(1)	(2)	(3)	(4)	(5)	(6)
	Spots		Applications		# Spots/	
Admission Pool	#	%	#	%	# Apps (%)	Cutoff
Open seats	92,409	46.0	1,978,077	44.0	4.7	664.1
Reserved seats, by Law						
Public HS	18,850	9.4	412,892	9.2	4.6	638.6
Public HS & Income	19,768	9.8	482,761	10.7	4.1	620.3
Public HS & Race	28,475	14.2	522,994	11.6	5.4	614.4
Public HS & Race & Income	29,597	14.7	964,483	21.5	3.1	606.3
Reserved seats, Other	11,778	5.9	135,491	3.0	8.7	-
Total	200,877	100	4,496,698	100		

Notes: Own elaboration based on SISU microdata from 2016. This table shows the breakdown of the number of applicants and spots for each of the different admission pools. The rightmost column shows the admission cutoff of degrees weighted by the number of spots they offer. The "Other" admission channel includes other affirmative action initiatives that are not mandated by the federal regulation. We omit the cutoff of the "Other" admission channel as it is available in some selected degrees, and thus not comparable to the admission cutoff of the other admission pools.

Table 2: Preference Parameters

	Targeted				Non-targeted			
	Mean	SD	P10	P90	Mean	SD	P10	P90
Panel (a): School Choice Model								
Degree FE (δ_j)	11.85	11.34	-0.06	25.97	29.83	17.55	12.34	51.06
Ability (γ_j)	-0.02	0.02	-0.05	-0.01	-0.05	0.02	-0.07	-0.02
Location (κ)	- 2.73	-	-	-	-2.36	-	-	-
	(0.35)	-	-	-	(0.15)	-	-	-
Average Valuation (V_j)	-4.47	1.38	-5.90	-2.97	-3.43	1.15	-4.56	-2.07
Panel (b): Potential Outcomes Model								
Degree FE (α_j)	-278	13,330	-16,974	11,362	1,177	39,553	-49,216	31,888
Ability (β_j^a)	0.96	20.58	-17.73	26.59	-3.97	56.83	-50.27	70.21
Selection in gains (ρ)	72.28	-	-	-	572	-	-	-
	(10.78)	-	-	-	(25.37)	-	-	-
Value Added (VA_j)	695	1,687	-687	3,011	-665	2,230	-2,736	3,161
Share of students	54.8%				44.7%			
Number of students	329,621							
Number of degrees	502							

Notes: This table summarizes the parameters estimates. AA denote the targeted students, while NA refers to non-targeted students. Panel A presents coefficients from the school choice model. Panel B display parameters from the potential outcomes model.

Appendix for:

Affirmative Action in Centralized College Admission Systems: Evidence from Brazil

Sebastián Otero Nano Barahona Cauê Dobbin

November 12, 2021

A Additional tables and figures

Table A.1: Administrative sector and type categories (as of 2016)

	Enrollment (# in 1,000)	Enrollment (share %)	Institutions	Degrees
Panel A: by sector, all students				
Federal	1,249	15.7	107	6,361
State	623	7.8	123	3,541
Municipal	47	0.6	45	264
Private	6,058	75.9	2,110	22,827
Panel B: by institution type, only federal enrollment				
University	1,084	86.7	63	5,005
Vocational	165	13.3	44	1,356

Notes: Calculations based on the Brazilian Higher Education micro data of 2016. Panel A in this table shows the breakdown of the number and share of students enrolled by sector. Panel B shows the enrollment number and share by institutions type among students enrolled in federal institutions. The total number of students is in 1,000s. Sector refers to whether the higher education institution is public (Federal, State, or Municipal administered) or private. Type refers to whether the institution is a university, or a vocational institution.

Figure A.1: Quality distribution

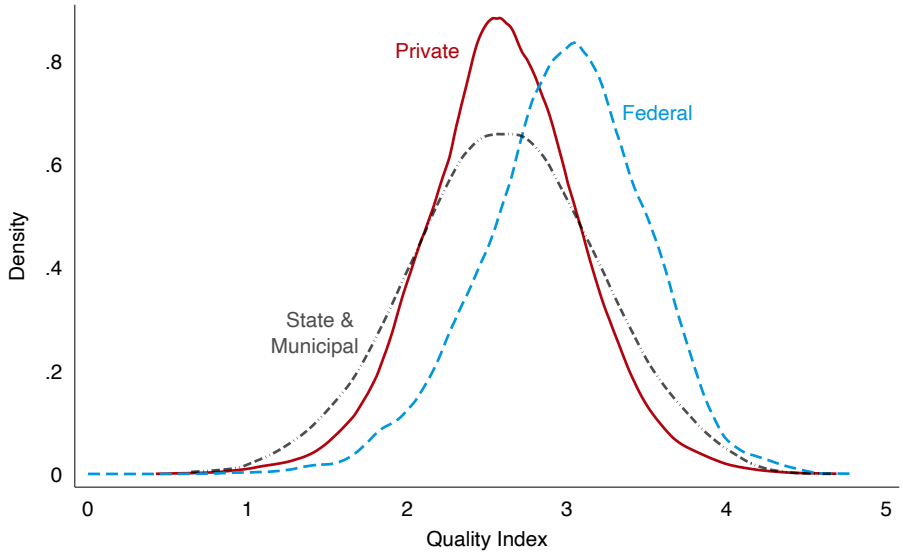

Notes: This figure shows the distribution of quality as measured by an index ranging from 1 to 5 prepared by the Ministry of Education of Brazil to evaluate degrees between years 2014 and 2016. An observation is a degree, and each observation is weighted by the number of students enrolled in the degree. Degrees from state and municipal institutions are pooled together.

Figure A.2: Staggered Implementation of the policy

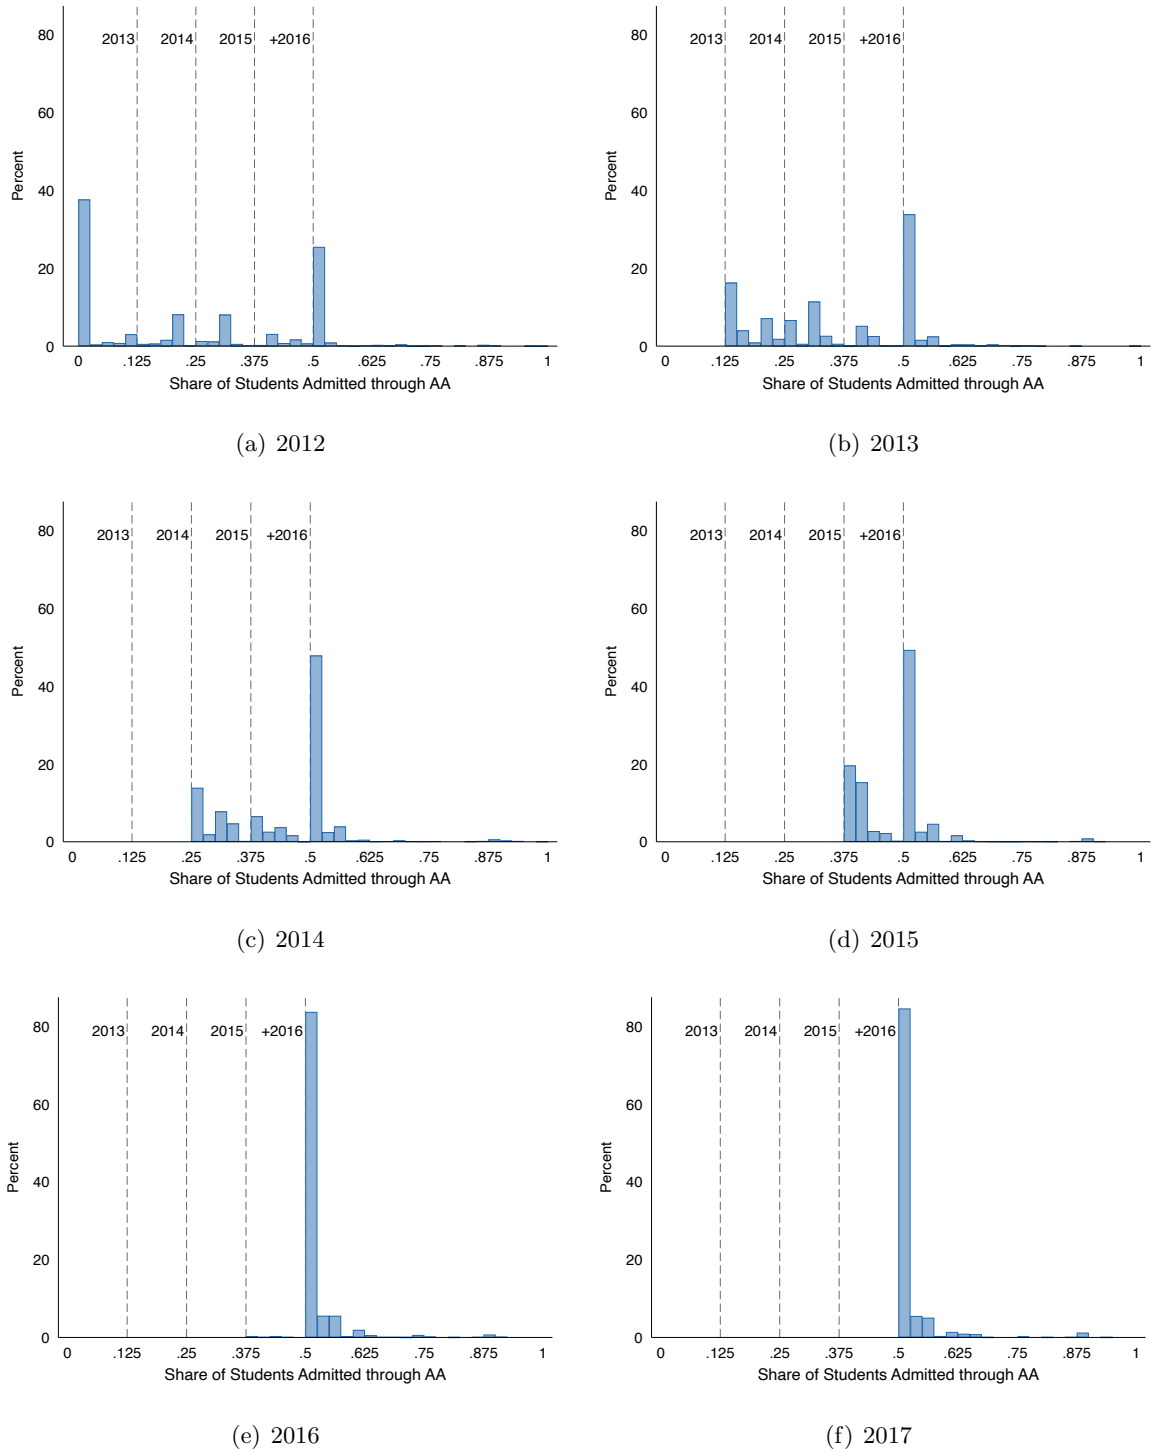

(a) 2012

(b) 2013

(c) 2014

(d) 2015

(e) 2016

(f) 2017

Notes: These figures describe the implementation of the affirmative action regulation. An observation is a degree program, and the y-axis measures the share of student that got admitted through the affirmative action admission track in each of the years. The dashed lines represents 12.5%, 25%, 30% and 50%. This figure only uses data from institutions participating from the SISU process.

Figure A.3: Share of public high school students of color in federal institutions

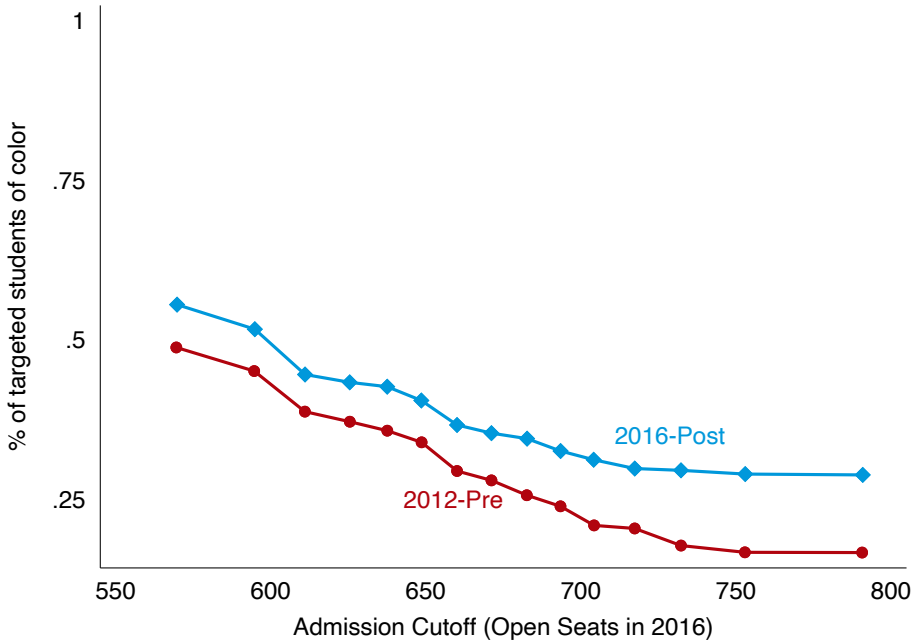

Notes: This Figure presents the average share of incoming targeted students of color in federal institutions by degree selectivity as defined as the admission cutoff of open spots in 2016 SISU. We consider individuals starting a degree program in the first semester of each year. An observation is a degree and shift tuple. We weight each observation by the size of the incoming cohort. We keep degree programs that exist both in 2012 and 2016 and that participated in 2016 SISU. Our sample covers 94% of the total enrollment in federal institutions in 2012 and 2016

Figure A.4: Difference in admission cutoff between open and reserved seats

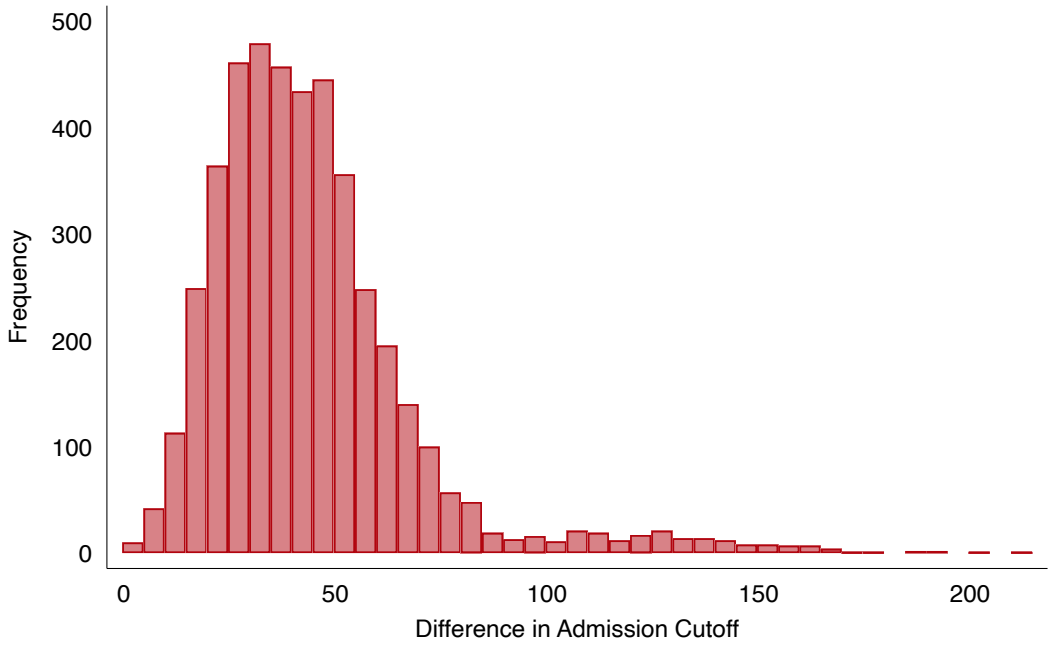

Notes: This Figure presents the distribution of the difference in admission cutoffs for open and reserved seats across all federal degrees offered in the centralized system in 2016. We construct the average cutoff for reserved seats as the average cutoff across all four different type of reserved seats weighted by the number of seats. The distribution of ENEM average test scores across all SISU applicants has a mean of 531 points and standard deviation of 73 points.

Figure A.5: Regression discontinuity estimates, balance tests

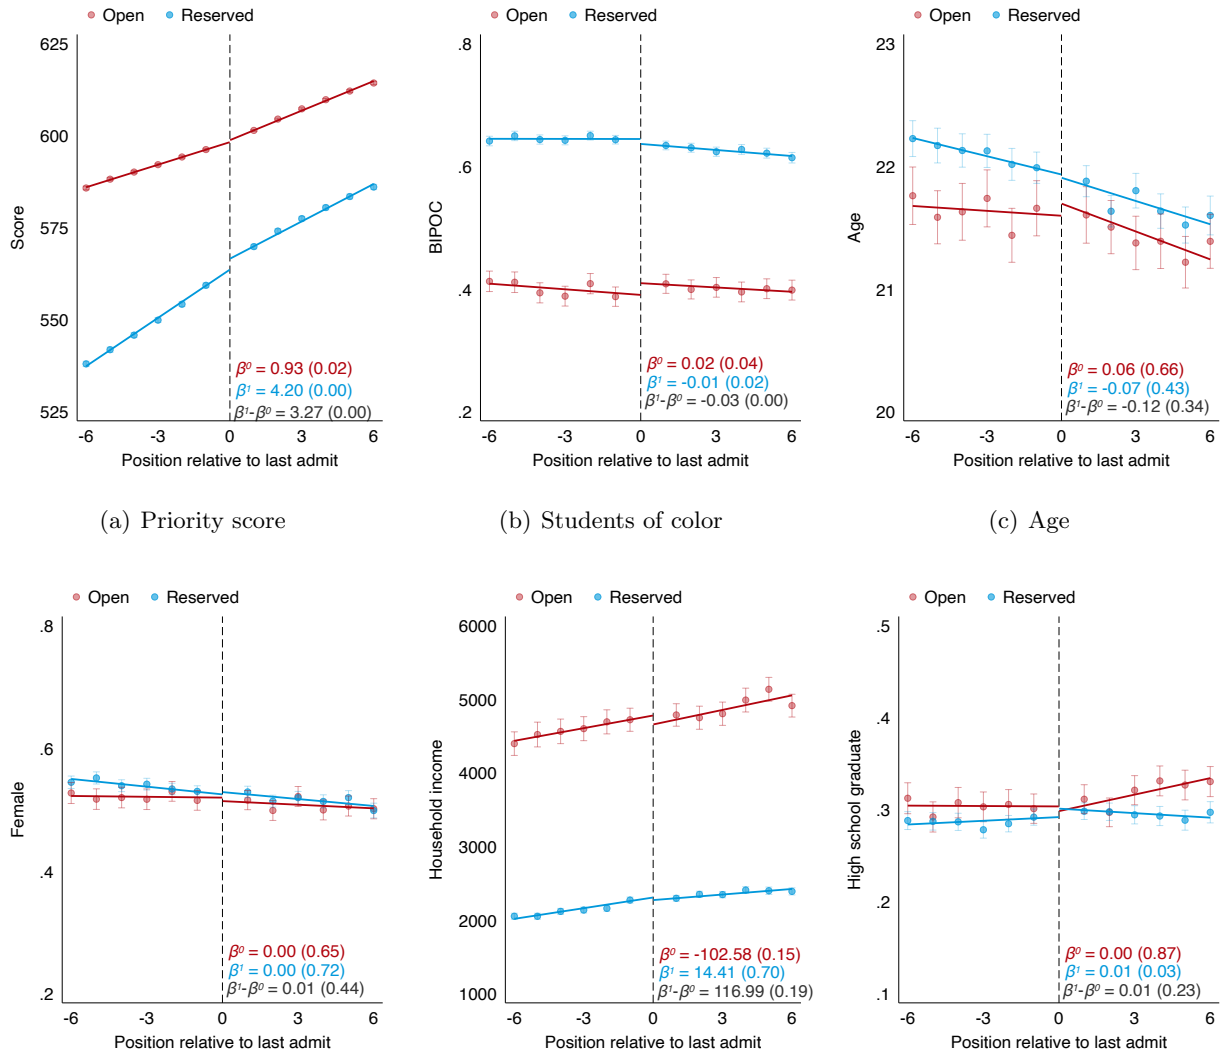

(a) Priority score (b) Students of color (c) Age

(d) Gender (e) Household income (f) High school graduate

Notes: This figure shows the regression discontinuity estimates for a number of academic and labor market outcomes, and time periods. We pool all degree programs and center them around the last admitted individual. For simpler presentation, we also pool individuals across all four type of reserved seats, and present their estimates as a single group. The red and blue lines show the outcomes for students applying for an open seat and a reserved seat respectively. The confidence intervals represent the standard error of the outcome mean. The coefficients at the bottom of the figure indicate the treatment effect for each of the groups, as well as the difference between these coefficients. P-values are reported in parentheses. Standard errors of the estimates are clustered at the waitlist level.

Figure A.6: Number of exams per grader

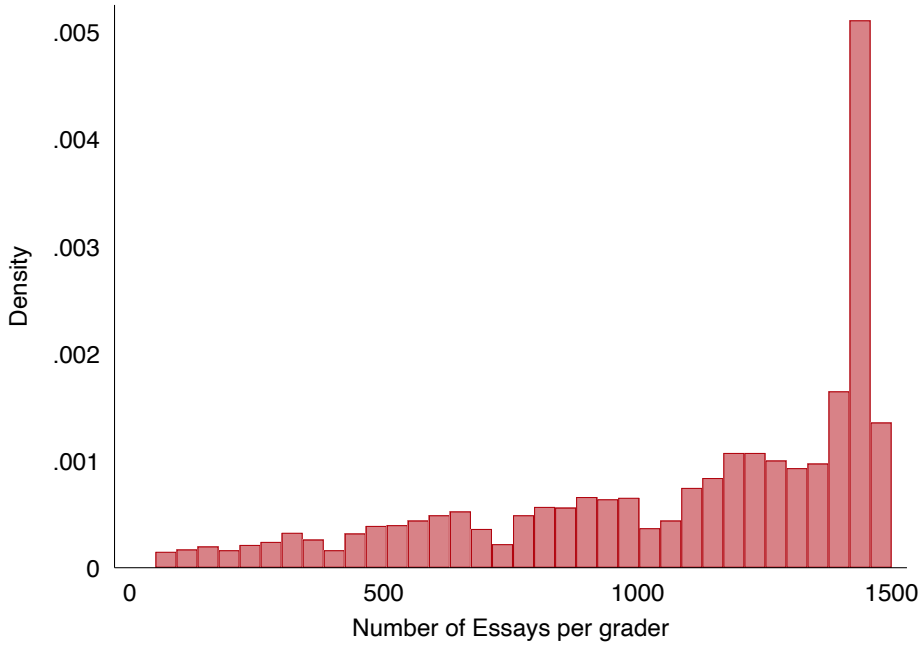

Notes: This figure shows the distribution of the number of essays assigned to each grader. We use only graders assigned to more than 50 essays. There are a total of 10,356 graders, and the average grader receives 1,067 exams to grade.

Figure A.7: Quantile-quantile admission cutoffs, Open vs Reserved

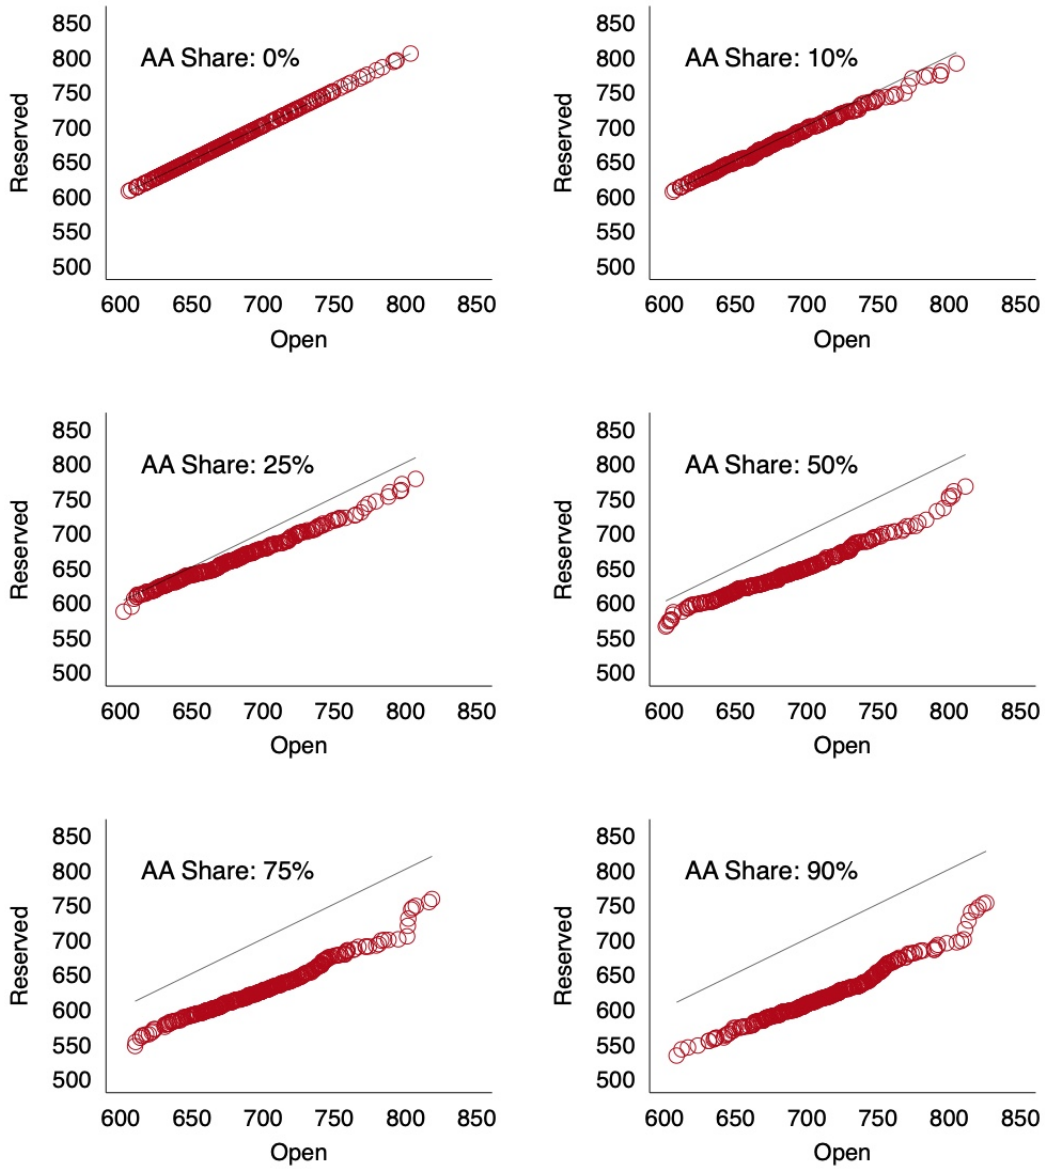

Notes: This figure shows the quantile-quantile plot of admission cutoffs of open and reserved spots, under different affirmative action schedules. The solid line is a 45 degree line. The x-axis and the y-axis represent the admission cutoff of open and reserved spots, respectively. The red circles denote the quantile associated to each of the admission cutoff value.

Figure A.8: Quantile-quantile plots, No AA vs Open

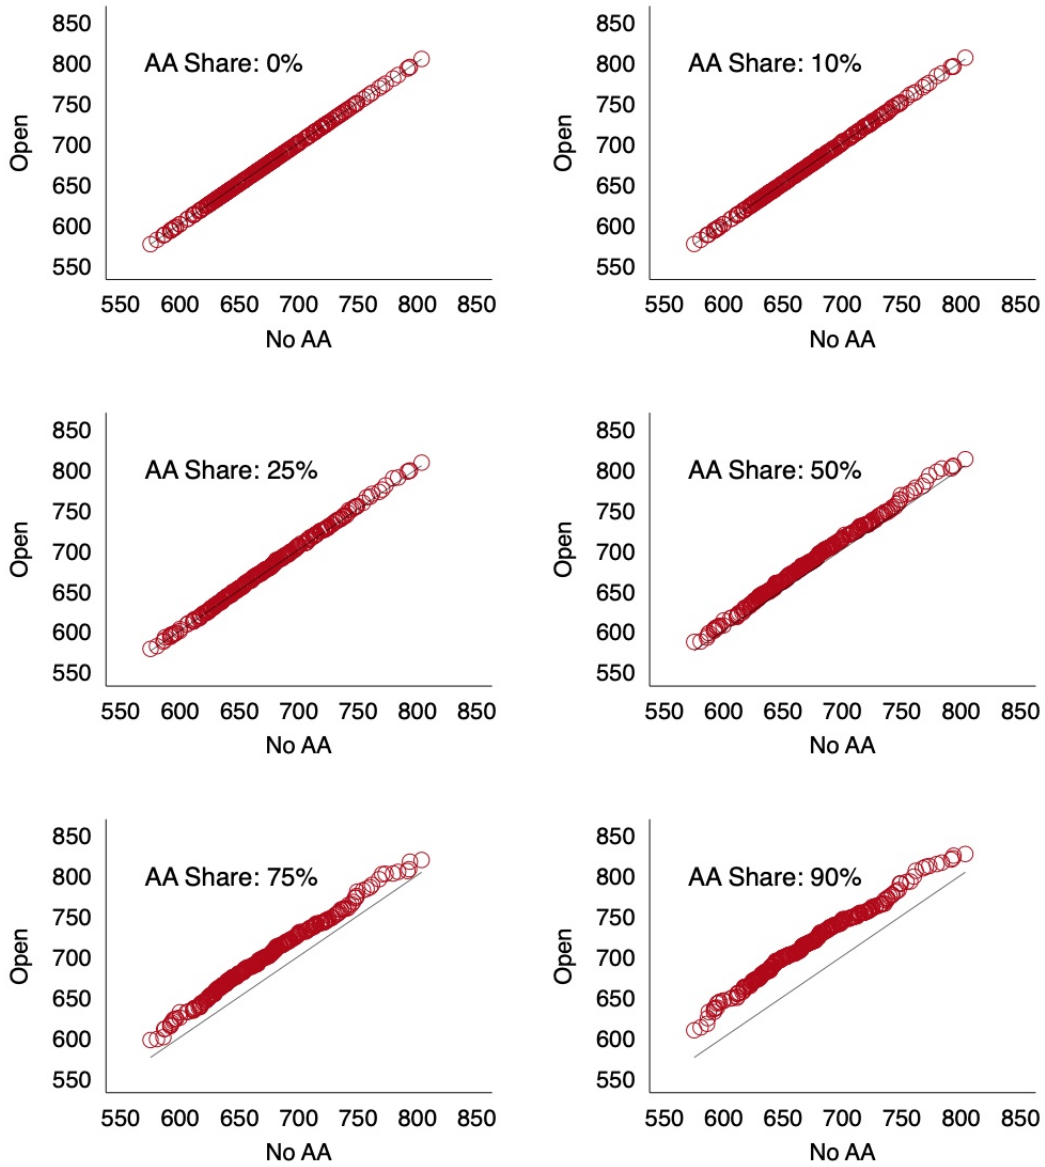

Notes: This figure shows the quantile-quantile plot of admission cutoffs of open spots. The solid line is a 45 degree line. The x-axis represent the admission cutoff of spots in the absence of the AA policy (i.e. $\omega = 0$). The y-axis represent the admission cutoff of open spots under different affirmative action schedules. The red circles denote the quantile associated to each of the admission cutoff values.

Figure A.9: Quantile-quantile plots, No AA vs Reserved

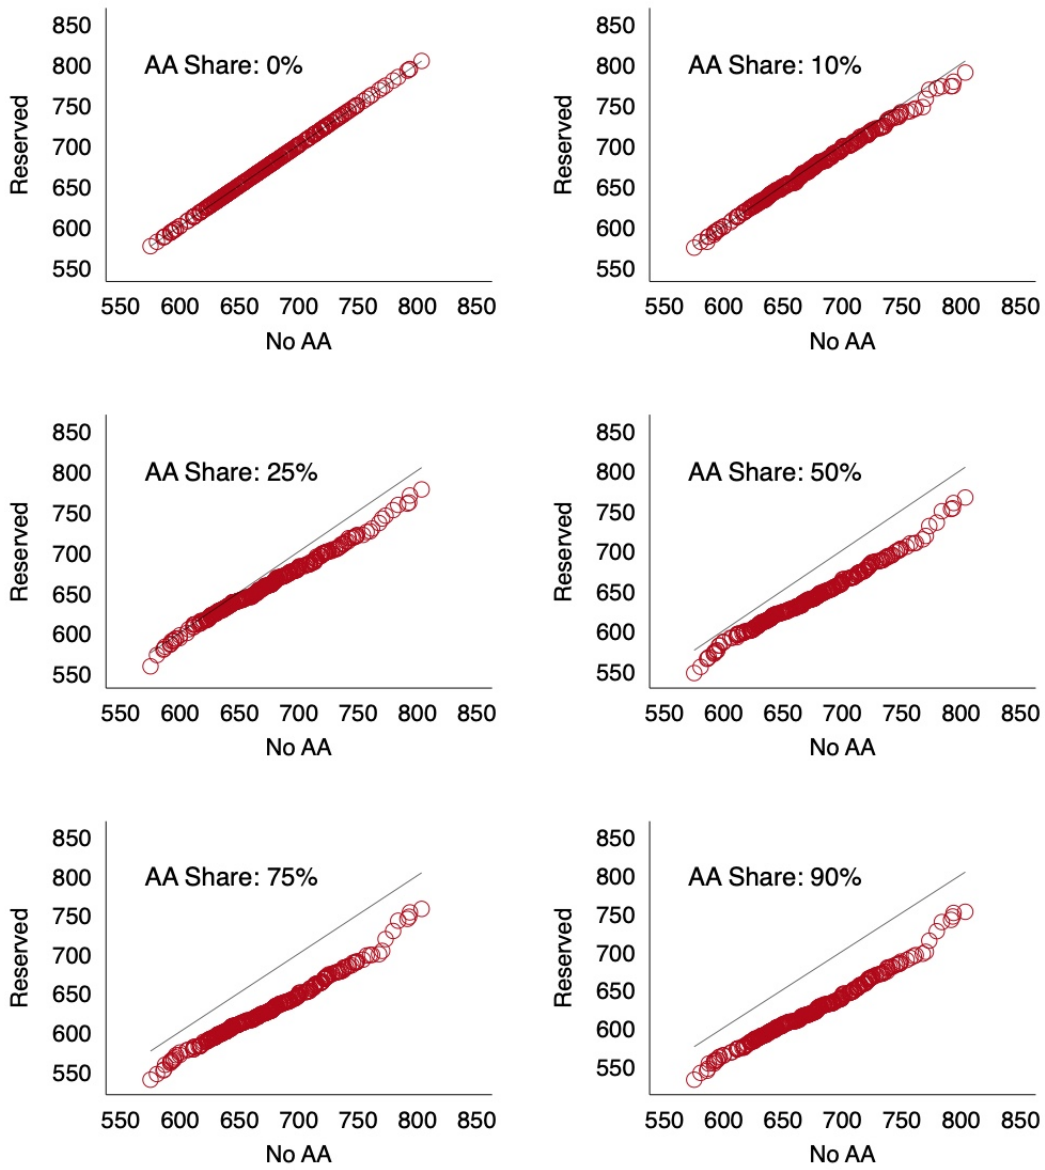

Notes: This figure shows the quantile-quantile plot of admission cutoffs of reserved spots. The solid line is a 45 degree line. The x-axis represent the admission cutoff of spots in the absence of the AA policy (i.e. $\omega = 0$). The y-axis represent the admission cutoff of reserved spots under different affirmative action schedules. The red circles denote the quantile associated to each of the admission cutoff values.

B Iterative Admission Cutoffs

The centralized admission system uses an iterative deferred acceptance mechanism. Under this system, students are sequentially asked to submit rank ordered lists over the course of several "trial" day. At the end of each day, the system produces a cutoff grade representing the lowest grade necessary to be accepted at a specific program. In this Appendix we explore the change in admission cutoffs over the course of the application period.

Unfortunately, there are no administrative records of the admission cutoffs reported by the system throughout the application period. The system only saves the final admission cutoffs. To circumvent this issue, we scrapped online data of the degrees offered by the Federal University of Minas Gerais (UFMG) during the 2021 admission period. In total we observe 450 admission cutoffs (90 degrees with 5 admission tracks each) over the 9 days of the application process. In Figure B.1 we show how the admission cutoffs evolve over time. Panel (a) displays the absolute difference between the final admission cutoffs (day T in the figure) and the one reported in the first day of the admission period. We observe large differences in admission cutoffs, especially for less selective degrees. Panel (b) displays the absolute difference between the final admission cutoffs (day T in the figure) and the one reported in last day of the admission period. The pattern in the data shows that at the last day of the application period, most degrees have converged or are very close to converging to the final admission cutoff.

Figure B.1: Absolute Change in admission cutoffs

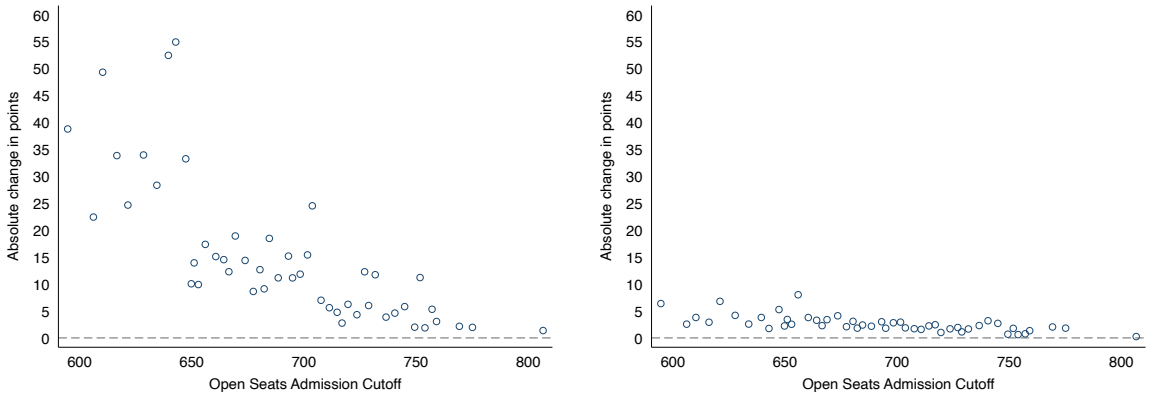

(a) Day 1 - Day T

(b) Day T-1 - Day T

Notes: This figure shows the absolute change in admission cutoffs over degree selectivity, as defined by the degree's final admission cutoffs of the open seats. Panel (a) display the difference between the absolute difference between the final admission cutoffs and those reported in the first day of the system. Panel (b) displays the difference between the absolute difference between the final admission cutoffs in those reported in the penultimate day of the system.

Overall, these data suggest that the ex-ante and ex-post eligibility into degrees are very similar.

C Identification Proof

C.1 Setting

We consider a set of individuals indexed by $i \in \mathcal{I}$, applying to a finite set of selective college degree programs through a centralized platform. Let $\mathcal{J} = \{0, 1, 2, 3, ..., J\}$ denote the set of degrees indexed by j offered in the system, where $j = 0$ represents the outside option with unlimited capacity. Students are characterized by their preferences $\succ_i$, priority scores s_i, and an an exogenous score shifter z_i. For simplicity, assume all degrees give the same weight to all components of the test such that priority scores do not depend on j. In addition, we normalize the score shifter such that it has a mean of zero, $\bar{z} = 0$. Let corrected test scores, $\tilde{s}_i$, denote the test scores students would have gotten if $z_i = 0$. We assume that $\tilde{s}_i$ is additive separable, thus $\tilde{s}_i = s_i - z_i$.

The centralized mechanism uses a deferred acceptance mechanism, such that the matching of students to degrees has a unique representation in terms of a vector of market clearing admission cutoffs c_j (Azevedo and Leshno, 2016). Assume that degrees are ordered by their selectivity, such that their admission cutoffs are given by $c_1 < c_2 < c_3 < ... < c_J$, and $c_0 = -\infty$. Let $\Omega_i = \{j \in \mathcal{J} \mid s_i \geq c_j\}$ represent the *feasible choice set* for individual i, defined as those degrees to which they could have gained access based on their score. Students choose their preferred degree D_i among all degrees available in their choice sets.

Let Y_{ij} denote the potential outcome of student i for attending degree j. We assume that potential outcomes and the score shifter z_i are independent, thus $\mathbb{E}[Y_{ij}|z_i] = \mathbb{E}[Y_{ij}]$. The latter implies that $\mathbb{E}[Y_{ij}|s_i] = \mathbb{E}[Y_{ij}|\tilde{s}_i]$. We write the potential outcomes as:

$$\begin{aligned} Y_{ij} &= \mathbb{E}[Y_{ij}|s_i] + \varepsilon_{ij} \\ &= Y_j(s_i) + \varepsilon_{ij}, \end{aligned} \tag{1}$$

where ε_{ij} is an unobserved term, such that $\mathbb{E}[\varepsilon_{ij}|s_i] = 0$. In the data, we observe student scores s_i, scores-shifters z_i, choice-sets Ω_i, degree choices D_i, and outcomes Y_i. The non-parametric form of the mean outcome observed in the data is:

$$\mathbb{E}[Y_i|D_i = j, \Omega_i, s_i] = Y_j(s_i) + \mathbb{E}[\varepsilon_{ij}|D_i = j,\, \Omega_i,\, s_i]. \tag{2}$$

Our goal is to recover $Y_j(s_i)$.

To simplify the algebra, we define $\Delta_j^{m-n}(s_i)$ as the difference in mean outcomes for individuals with score s_i, attending j with choice set Ω^m relative to those with choice set Ω^n.

$$\Delta_j^{m-n}(\tilde{s}_i) = \mathbb{E}[Y_i|D_i = j, \Omega_i = \Omega^m, s_i] - \mathbb{E}[Y_i|D_i = j, \Omega_i = \Omega^n, s_i] \tag{3}$$

$$= \mathbb{E}[\varepsilon_{ij}|D_i = j,\, \Omega_i = \Omega^m,\, s_i] - \mathbb{E}[\varepsilon_{ij}|D_i = j,\, \Omega_i = \Omega^n,\, s_i], \tag{4}$$

The equality between Equation (3) and (4) follows from Equation (2). Additionally, we use the

exclusion restriction between Y_{ij} and z_i, to show that:

$$\mathbb{E}[Y_i|D_i = j, \Omega_i = \Omega^m, \tilde{s}_i] = \mathbb{E}[Y_i|D_i = j, \Omega_i = \Omega^m, s_i],$$

which implies that $\Delta_j^{m-n}(\tilde{s}_i) = \Delta_j^{m-n}(s_i)$. Finally, note that $\Delta_j^{m-n}(\cdot)$ is a difference of conditional expectation functions that are observed in the data.

C.2 Random Choice Sets

We start by exploring random allocation of choice sets as our benchmark case. We explore identification for different sizes in the availability of degrees.

Case 1: $J = 1$. When there are only 2 degree programs, we can identify the LATE between the inside and outside options. Following the standard treatment effect literature, we define the following objects:

$$\begin{aligned}
\mathbb{E}[Y_{i1}|C, s_i] &= \mathbb{E}[Y_i|D_i = 1, \Omega_i = (1,0), s_i] \\
\mathbb{E}[Y_{i0}|NT, s_i] &= \mathbb{E}[Y_i|D_i = 0, \Omega_i = (1,0), s_i] \\
\mathbb{E}[Y_{i0}|s_i] &= \mathbb{E}[Y_i|D_i = 0, \Omega_i = (0), s_i]
\end{aligned}$$

where C denotes the compliers (i.e., students with $\Omega_i = (1,0)$ that decide to attend degree 1 when given the option to do so), and NT denote the never takers (i.e. students that always choose the outside option). Note that we observe all of these terms in the data. Additionally, we also observe the share of compliers given by $\pi(s_i) = Pr(D_i = 1|\Omega_i = (1,0), s_i)$.

Since Ω_i is randomly assigned, by the law of total probability we have that:

$$\mathbb{E}[Y_{i0}|s_i] = (1 - \pi(s_i)) \cdot \mathbb{E}[Y_{i0}|NT, s_i] + \pi(s_i) \cdot \mathbb{E}[Y_{i0}|C, s_i]$$

which allows us to identify $\mathbb{E}[Y_{i0}|C, s_i]$. Hence, for any given s_i we can identify the LATE given by:

$$\tau^{\text{LATE}}(s_i) = \mathbb{E}[Y_{i1} - Y_{i0}|C, s_i].$$

Case 2: $J \geq 3$. When choice sets are random and $J \geq 3$, we can non-parametrically identify the average treatment effect (ATE) of attending each degree for any given test score s_i. For the case of $J = 3$, an individual with score s_i could observe the following choice sets (the outside option is always included):

$$\Omega_i(s_i) \in \{(0,1,2,3), (0,1,2), (0,1,3), (0,2,3), (0,1), (0,2), (0,3), (0)\}$$

We prove identification of $\mathbb{E}[Y_{ij}|s_i]$ for $j = 1$. From the definition of $\Delta_j^{m-n}(s_i)$ we know that:

$$\begin{aligned}
\Delta_1^{(0,1,2,3)-(0,1,2)}(s_i) &= \mathbb{E}[\varepsilon_{i1}|D_i = 1,\, \Omega_i = (0,1,2,3),\, s_i] - \mathbb{E}[\varepsilon_{i1}|D_i = 1, \Omega_i = (0,1,2),\, s_i] \\
\Delta_1^{(0,1,2,3)-(0,1)}(s_i) &= \mathbb{E}[\varepsilon_{i1}|D_i = 1,\, \Omega_i = (0,1,2,3),\, s_i] - \mathbb{E}[\varepsilon_{i1}|D_i = 1, \Omega_i = (0,1),\, s_i] \\
\Delta_1^{(0,1,2)-(0,1)}(s_i) &= \mathbb{E}[\varepsilon_{i1}|D_i = 1,\, \Omega_i = (0,1,2),\, s_i] - \mathbb{E}[\varepsilon_{i1}|D_i = 1, \Omega_i = (0,1),\, s_i]
\end{aligned}$$

which leads to a system of 3 equations and 3 unknown terms. This means that the unobserved component of equation (2) is identified. Since we know $\mathbb{E}[\varepsilon_{i1}|D_i = 1, \Omega_i, s_i]$, and we observe $\mathbb{E}[Y_i|D_i = 1, \Omega_i, s_i]$, we can recover $Y_1(s_i)$. A similar argument follows for the rest of degrees, which allows us to identify $Y_j(s_i)$ for all $j \in \mathcal{J}$. The same proof applies for $J \geq 3$.

Thus, the ATE for individuals of score s_i for attending degree j over degree k is given by:

$$\tau_{jk}^{\text{ATE}}(s_i) = Y_j(s_i) - Y_k(s_i).$$

C.3 Deferred acceptance with z_i with large support

Next, we consider a case in which choice sets are assigned following a deferred acceptance algorithm. Assume the support of z_i is large enough in the sense that it has the same range as the range of s_i. We consider the following cases:

Case 1: $J = 1$. When there are only 2 degree programs, we can identify the LATE for individuals with any given corrected test score $\tilde{s}_i$. Identification works similarly to when degrees are randomly assigned.

Case 2: $J \geq 3$. Under a deferred acceptance mechanism with $J \geq 3$, we can non-parametrically identify the ATE of attending any degree but the two top most selective degree programs. For those two degree programs we can only identify the LATE. For the case of $J = 3$, an individual with corrected test score $\tilde{s}_i$ could observe the following choice sets (the outside option is always included):

$$\Omega_i(\tilde{s}_i) \in \{(0,1,2,3), (0,1,2), (0,1), (0)\}$$

We use $\Delta_j^{m-n}(\tilde{s}_i)$ defined above, to set the following system of equations for $j = 0$:

$$\begin{aligned}
\Delta_0^{(0,1,2,3)-(0,1,2)}(\tilde{s}_i) &= \mathbb{E}[\varepsilon_{i0}|D_i = 0, \Omega_i = (0,1,2,3), s_i] - \mathbb{E}[\varepsilon_{i0}|D_i = 0, \Omega_i = (0,1,2), s_i] \\
\Delta_0^{(0,1,2,3)-(0,1)}(\tilde{s}_i) &= \mathbb{E}[\varepsilon_{i0}|D_i = 0, \Omega_i = (0,1,2,3), s_i] - \mathbb{E}[\varepsilon_{i0}|D_i = 0, \Omega_i = (0,1), s_i] \\
\Delta_0^{(0,1,2)-(0,1)}(\tilde{s}_i) &= \mathbb{E}[\varepsilon_{i0}|D_i = 0, \Omega_i = (0,1,2), s_i] - \mathbb{E}[\varepsilon_{i0}|D_i = 0, \Omega_i = (0,1), s_i]
\end{aligned}$$

which leads to a system of 3 equations and 3 unknown terms. This means that the unobserved component of equation (2) is identified. Since we know $\mathbb{E}[\varepsilon_{i0}|D_i = 0, \Omega_i, s_i]$, and we observe $\mathbb{E}[Y_i|D_i = 0, \Omega_i, s_i]$, we can recover $Y_0(s_i)$.

For j=1, we set the following system of equations:

$$\begin{aligned}
\Delta_1^{(0,1,2,3)-(0,1,2)}(\tilde{s}_i) &= \mathbb{E}[\varepsilon_{i1}|D_i = 1, \Omega_i = (0,1,2,3), s_i] - \mathbb{E}[\varepsilon_{i1}|D_i = 1, \Omega_i = (0,1,2), s_i] \\
\Delta_1^{(0,1,2,3)-(0,1)}(\tilde{s}_i) &= \mathbb{E}[\varepsilon_{i1}|D_i = 1, \Omega_i = (0,1,2,3), s_i] - \mathbb{E}[\varepsilon_{i1}|D_i = 1, \Omega_i = (0,1), s_i] \\
\Delta_1^{(0,1,2)-(0,1)}(\tilde{s}_i) &= \mathbb{E}[\varepsilon_{i1}|D_i = 1, \Omega_i = (0,1,2), s_i] - \mathbb{E}[\varepsilon_{i1}|D_i = 1, \Omega_i = (0,1), s_i]
\end{aligned}$$

which leads to a system of 3 equations and 3 unknown terms. Following a similar argument as above, this allows us to recover $Y_1(s_i)$. Using a similar argument as with $J = 1$, we can identify the LATE of attending $j = 2$ and $j = 3$.

D Predicted income

In this Appendix we discuss the econometric procedure that we use to predict student's income based on their academic trajectories. Ideally, we would like to use labor market outcomes of students as the main outcome of interest for the analysis. Unfortunately, our sample consist of 2015 ENEM takers, and as such several individuals are still enrolled in college which makes it too early to find them in the labor market. Instead, we create a measure of predicted income that takes into account the academic trajectories between the 2016 and 2019 academic years. We argue that using the predicted income has two important advantages. First, we use this income as a currency that allow us to compare individuals with different trajectories that would be otherwise impossible weigh up against each other. Second, we can use trajectories as surrogates, and under some assumptions interpret these results as the long term effects of the policy (Athey et al., 2019).

Table D.2 illustrates the problem. Assume 0 is the outside option, and 1 and 2 are degree programs in public and private institutions respectively. For the 2015 ENEM cohort we observe trajectories represented by columns (1)-(4), but we do not observe income. Instead for past cohorts we observe both trajectories together with student's income several years after taking ENEM. We use past cohorts to create a mapping from trajectories to income, and then apply this mapping to the 2015 ENEM cohort to recover their predicted income.

Table D.2: Examples of Academic Trajectories

	Degree in $t = 1$	Degree in $t = 2$	Degree in $t = 3$	Degree in $t = 4$	Income in $t = 8$
Student A	1	1	1	1	400
Student B	1	0	0	0	100
Student C	1	2	2	2	300
Student D	0	0	0	2	150
Student E	2	2	2	2	350

Our working sample consists of all ENEM takers between years 2009-2012. We have one snapshot of income in 2017 from the administrative matched employer employee records (RAIS). Ideally, we would not parametrically match individuals with identical academic trajectories as a way to recover predicted income for the 2015 ENEM sample. Although this exact matching is hypothetically possible, it is unfeasible in our data due to the large number of degrees and the many possible combinations of degrees that define a trajectory.

As an alternative approach, we summarize the academic trajectory using degree attainment in year 1 and year 4. We provide two different specification which differ on the set of individual level controls that we add together with the trajectory fixed effects. In specification (1) we add a set of controls X_i^1 which include gender, and the test scores in each of the 5 components of the ENEM test. In specification (2) we include the same set of covariates but also add variables accounting for the traits targeted by the AA regulation (i.e. racial and high-school dummies). Specifically, these

equations are:

$$\begin{aligned}
\text{Specification 1:} \quad y_{i,T} &= \phi_t^1 + X_i^1 \pi^1 + \delta^1_{J(i,t+1)} + \delta^1_{J(i,t+4)} + \epsilon_i^1 \\
\text{Specification 2:} \quad y_{i,T} &= \phi_t^2 + X_i^2 \pi^2 + \delta^2_{J(i,t+1)} + \delta^2_{J(i,t+4)} + \epsilon_i^2,
\end{aligned}$$

where α_t are dummies indicating the year of ENEM, T denotes the year in which we observe income. Degree fixed effects are captured by δ, and $J(i,t)$ is a function indicating the degree that student i attends in period t. We use these coefficients to predict earnings of 2015 ENEM takers (SISU 2016).

E Additional balance checks

E.1 Grader Assignment

In columns (1)-(3) of Table E.1 we present the OLS first-stage relationship between the average score shifter and test scores across all subjects. The regression coefficient is of 0.721 when using the full sample of ENEM takers. Although our sample is sufficiently large, this coefficient is different than 1 because of the introduction of the third grader when there is disagreement among the two initially assigned graders. In column (2), we show the regression estimates when we restrict the sample to individuals participating in SISU 2016, our relevant sample. The coefficient remains very similar to that of the full sample, and drops slightly in magnitude to 0.692.

Given that the allocation of graders is random, our score shifter measure should be uncorrelated with student's performance in the other components of the test and any other important confounders. In column (3) we show that the first stage coefficient is remarkably stable to adding a rich host of student level controls.[1] Consistent with a satisfactory randomized assignment, this specification suggests no correlation between the score shifter z_i and any relevant student characteristics. In columns (4)-(6) we maintain the same set of specifications, but change the dependent variable to the average in score across the other four multiple choice tests. We find a precise estimate of 0 when we use the full sample of ENEM takers. When we restrict the sample to SISU applicants we find a slight imbalance, but almost negligible in magnitude. In column (6) we see that the coefficient stays stable after including several control variables suggesting the minor imbalance is not correlated with any relevant student characteristic.

Table E.1: First Stage Regressions

Dependent Variable	Essay			Multiple Choice		
	(1)	(2)	(3)	(4)	(5)	(6)
Score shifter (z_i)	0.721	0.693	0.696	0.001	-0.012	-0.010
	(0.002)	(0.003)	(0.003)	(0.001)	(0.002)	(0.001)
Sample	All	SISU	SISU	All	SISU	SISU
Controls	No	No	Yes	No	No	Yes
Mean	542.9	576.9	576.9	503.9	520.7	520.7
Observations	4,742,468	2,704,388	2,704,367	4,742,468	2,704,388	2,704,367
R-squared	0.023	0.021	0.170	0.000	0.000	0.282

Notes: The table presents the coefficients from OLS of ENEM scores on the score shifter z_i. Columns (1)-(3) use the test in the essay as main dependent variable, and columns (4)-(6) use the average score in the other four multiple choice components of the test. Specifications (1) and (4) use the full sample of ENEM takers. All other specifications use a restricted sample of ENEM takers who apply to the centralized system SISU. Columns (3) and (6) include a host of student level controls including: the type of high-school attended, dummies for different races, the state of residency, gender, year of high school graduation, age and marital status.

[1] These controls include the type of high-school attended, dummies for different races, the state of residency, gender, year of high school graduation, age and marital status.

E.2 Assignment of examination booklets

Next, we assess the credibility of the randomness of the assignment of booklets to students. In order to reduce the dimensionality of our covariates and to avoid multi-testing issues, we assess balance by checking the stability of the OLS coefficient when including a comprehensive set of covariates. Specifically, we regress test scores in a given component of the test against booklet indicator dummies. We presents the results in Table E.2. To account for the fact that booklets are assigned within high schools, we include high school fixed effects in our main specification (odd columns). In a second specification, presented in the even columns, we add a rich vector of students covariates, including age, gender, race dummies, parental education, household income, and answers from an individual questionnaire taken from ENEM that characterize student's SES and educational plans. We find that coefficients are remarkably stable across specifications.

Table E.2: Booklet Assignment

	History		Science		Math		Language	
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)
Yellow	0.59	0.55	1.13	1.11	0.14	0.36	-1.53	-1.60
	(0.15)	(0.14)	(0.14)	(0.14)	(0.21)	(0.20)	(0.14)	(0.14)
Red	-0.34	-0.20	-1.64	-1.46	4.57	4.48	-0.68	-0.69
	(0.15)	(0.14)	(0.14)	(0.14)	(0.21)	(0.20)	(0.15)	(0.14)
Blue	0.32	0.39	2.19	2.36	2.76	2.93	-3.49	-3.53
	(0.15)	(0.14)	(0.14)	(0.14)	(0.21)	(0.20)	(0.15)	(0.14)
High school FE	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Controls	No	Yes	No	Yes	No	Yes	No	Yes

Notes: This table shows the estimates of a regression of test scores on booklets. The green booklet is the omitted category. The odd columns include only school fixed effects accounting for the fact that booklets are assigned within schools. The even columns include a rich vector of students covariates including: age, gender, race dummies, parental education, household income, high-school fixed effects, and answers from an individual questionnaire taken from ENEM that characterize student's SES and educational plans.

F Model Fit

To verify that our parameters can recover the observed allocation, we simulate the admission cutoffs predicted by the model and contrast them against those observed in the data. To recover the admission cutoffs we follow steps 1 to 5 described in Appendix H, and simulate 20 different admission cutoffs for each of the different degree programs.

Panel (a) and (b) in Figure F show the model fit for the admission cutoffs of open and reserved seats, respectively. In the horizontal axis, we plot the observed admission cutoff scores, while in the vertical axis, we plot the average simulated admission cutoff score across all simulations, and the 95% coverage interval. It is important to note that in our data we observe four different admission cutoffs for the reserved spots: one for each affirmative action type. To construct the observed cutoff we take a weighted average of the four admission cutoffs where the weights correspond to the share of students of each affirmative action type applying on SISU. We observe that admission cutoffs predicted by our model are very similar to those observed in the data. The correlation coefficient of the slope is virtually 1 for both open and reserved seats.

Figure F.1: Model fit

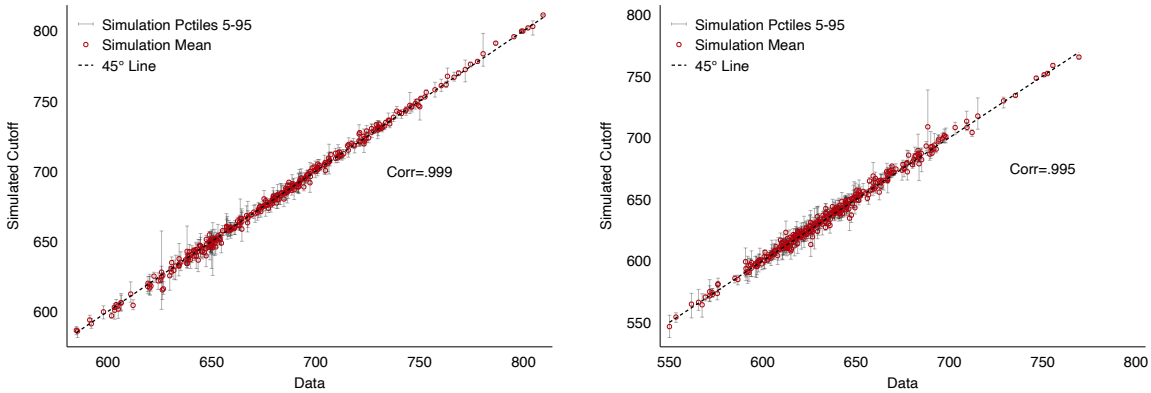

(a) Open

(b) Reserved

Notes: This figure shows the in-sample model fit for the admission cutoffs of open and reserved seats. In the horizontal axis, we plot the observed cutoff score, while in the vertical axis, we plot the average cutoff score across all simulations and the 95% coverage interval.

G Empirical Bayes

In this Appendix we describe the procedure to compute empirical Bayes posterior means for our parameters of interest. Parameters α_j^t and β_j^t of the potential outcome equation govern the structural relationship between predicted income, and degree attendance and student characteristics. Our model estimation produce a vector of estimates for each degree program and AA tuple, $\hat{\varphi}_j^t = \left[\hat{\alpha}_j^t, \hat{\beta}_j^t\right]$. These coefficients are unbiased but noisy estimates of the underlying outcome parameters φ_j^t. We follow Abdulkadiroglu et al. (2020) and use a hierarchical Bayesian model to improve our parameter estimates. Specifically,

$$\begin{aligned} \hat{\varphi}_j^t | \varphi_j^t &\sim N(\varphi_j^t, \Psi_j^t), \\ \varphi_j^t &\sim N(\mu_{\varphi^t}, \Sigma_{\varphi^t}), \end{aligned}$$

where Ψ_j^t is the sampling variance of the estimator $\hat{\varphi}_j^t$, and $\mu_{\varphi_j^t}$ and $\Sigma_{\varphi_j^t}$ are hyperparameters that govern a prior distribution for φ_j^t. We estimate $\mu_{\varphi_j^t}$ and $\Sigma_{\varphi_j^t}$ using the sample average and sample variance of $\hat{\varphi}_j^t$, respectively.

The empirical Bayes posterior mean for φ_j is given by

$$\varphi_j^{t*} = \left((\hat{\Psi}_j^t)^{-1} + (\hat{\Sigma}_{\varphi^t})^{-1}\right)^{-1} \left((\hat{\Psi}_j^t)^{-1}\hat{\varphi}_j^t + (\hat{\Sigma}_{\varphi^t})^{-1}\hat{\mu}_{\varphi^t}\right)$$

where $\hat{\Psi}_j^t$, $\hat{\mu}_{\varphi^t}$, and $\hat{\Sigma}_{\varphi^t}$ are estimates of Ψ_j^t, μ_{φ^t}, and Σ_{φ^t}.

H Simulation procedure

In this Appendix we describe the procedure used to simulate the counterfactuals. Let M denote the number of Monte Carlo simulations. The mth simulations works as follows:

1. Simulate a vector of unobserved tastes $\eta_{ij}^m \sim EVT1$
2. Compute preferences $\hat{\succ}_i^m$ reflecting indirect utilities $\hat{u}_{ij}^m$ using preferences estimates $(\hat{\delta}_j^t, \hat{\gamma}_j^t, \hat{\kappa}^t)$ together with η_{ij}^m
3. Construct the set of student types as $\hat{\Theta}^m = \bigcup_i \hat{\theta}_i^m$, where $\hat{\theta}_i^m = (\hat{\succ}_i^m, s_i, t_i)$
4. Compute the matching function $\varphi(\hat{\Theta}^m, q, \omega) = \hat{\mu}^m$ based on mechanism φ's allocation rules.
5. Calculate the cutoff scores $c_{jt}^m(\hat{\mu}^m)$ that are consistent with the equilibrium
6. Use Equation (9) to compute the predicted potential outcome as $\hat{Y}_{ij} = \hat{\alpha}_j + X_{ij}'\hat{\beta}_j + \sum_{k=1}^J \hat{\psi}_k \cdot (\hat{\eta}_{ik}^m - \mu_\eta) + \hat{\rho} \cdot (\hat{\eta}_{ij}^m - \mu_\eta)$
7. Compute the expect potential outcome for the corresponding matching function $\hat{Y}_i(\hat{\mu}^m) = \sum_j \mathbb{1}\{\hat{\mu}^m(\theta_i) = j\} \cdot \hat{Y}_{ij}$

Thus, the expected potential outcome for individual i under affirmative action schedule ω is:

$$\bar{\hat{Y}}_i(\omega) = \frac{1}{M} \sum_m \hat{Y}_i(\hat{\mu}^m)$$