Marginal Returns to Public Universities

Jack Mountjoy

University of Chicago Booth School of Business and NBER

This paper studies the causal impacts of public universities on the outcomes of their marginally admitted students. I use administrative admission records spanning all 35 public universities in Texas, which collectively enroll 10 percent of American public university students, to systematically identify and employ decentralized cutoffs in SAT/ACT scores that generate discontinuities in admission and enrollment. The typical marginally admitted student completes an additional year of education in the four-year sector, is 12 percentage points more likely to earn a bachelor's degree, and eventually earns 5-10 percent more than their marginally rejected but otherwise identical counterpart. Marginally admitted students pay no additional tuition costs thanks to offsetting grant aid; cost-benefit calculations show internal rates of return of 19-23 percent for the marginal students themselves, 10-12 percent for society (which must pay for the additional education), and 3-4 percent for the government budget. Finally, I develop a method to disentangle separate effects for students on the extensive margin of the four-year sector versus those who would fall back to another four-year school if rejected. Substantially larger extensive margin effects drive the results.

Marginal Returns to Public Universities

Jack Mountjoy[*]

March 27, 2024

Abstract

This paper studies the causal impacts of public universities on the outcomes of their marginally admitted students. I use administrative admission records spanning all 35 public universities in Texas, which collectively enroll 10 percent of American public university students, to systematically identify and employ decentralized cutoffs in SAT/ACT scores that generate discontinuities in admission and enrollment. The typical marginally admitted student completes an additional year of education in the four-year sector, is 12 percentage points more likely to earn a bachelor's degree, and eventually earns 5-10 percent more than their marginally rejected but otherwise identical counterpart. Marginally admitted students pay no additional tuition costs thanks to offsetting grant aid; cost-benefit calculations show internal rates of return of 19-23 percent for the marginal students themselves, 10-12 percent for society (which must pay for the additional education), and 3-4 percent for the government budget. Finally, I develop a method to disentangle separate effects for students on the extensive margin of the four-year sector versus those who would fall back to another four-year school if rejected. Substantially larger extensive margin effects drive the results.

Overall, the scarcity of credible evidence regarding the causal effect of college on earnings is striking given the voluminous literature on the returns to schooling more generally.

-Barrow and Malamud (2015), "Is College a Worthwhile Investment?"

1 Introduction

Is college worth it? American survey respondents are increasingly pessimistic, with a majority now declaring a four-year education "not worth the cost" (Belkin, 2023). Compelling causal evidence, meanwhile, remains surprisingly rare and limited in scope.[1] The vast majority of evidence is correlational, comparing the earnings outcomes of individuals with different levels of college attainment while controlling to various degrees (or often not at all) for observable confounders like academic ability and family background. These ubiquitous comparisons typically suggest large returns, but the specter of selection bias looms: those who end up with more college education may have had more advantages from the outset, confounding any causal impacts of college with systematic selection into it. Furthermore, even if college does boost earnings on average, students, taxpayers, and donors must pay for the privilege, and the policy-relevant net returns to enrolling *marginal* students may diverge substantially from the average (Carneiro et al., 2011; Zimmerman, 2014). Is the marginal American college student a good investment?

To make progress on this question, I assemble a large and previously untapped collection of admission cutoffs employed by a wide diversity of American public universities. I start with administrative admission records spanning 13 cohorts of applicants to each of the 35 public universities in Texas. Together, these universities enroll over 10 percent of all American public university students (National Center for Education Statistics, 2022). Using the individual-level test scores and admission decisions recorded in this data, I systematically identify hundreds of decentralized cutoffs in SAT and ACT scores, varying across schools and sometimes within schools across years, that generate abrupt discontinuities in admission and enrollment. I then link the marginal applicants around these cutoffs backward in time to their individual high school academic records and demographics to study their pre-college backgrounds, and forward in time to study their outcome trajectories of postsecondary enrollment, credit accumulation, degree completion, major choice, tuition costs, financial aid, student loan accumulation, and labor market earnings. Together, these data linkages and discontinuities enable a fuzzy regression discontinuity research design that transparently documents how student outcomes change discontinuously across the cutoffs, and attributes those changes in outcomes to discontinuous changes in admission and enrollment, justified by smooth densities of applicants and their pre-college characteristics through the cutoffs.

The marginal students around these admission cutoffs are an important population to study for at least three reasons. First, by construction, they straddle clear policy levers and help an-

---

[1]See Barrow and Malamud (2015) for a review, along with recent contributions by Smith et al. (2020), Bleemer (2021), and Kozakowski (2023). A larger literature has examined the causal earnings effects of compulsory schooling reforms, but these typically increase schooling at primary and secondary levels with little effect on college-going. See, for example, the recent contribution by Clark (2023) and the citations therein.

swer the question of whether public universities should expand or contract along their current admission margins. The answer is deeply uncertain without credible estimates of both the benefits and the costs generated by marginally admitted students, which this paper aims to provide. Second, marginally admitted students have weak academic preparation relative to their peers, and therefore have especially ambiguous ex ante returns to enrolling. On one hand, they may benefit disproportionately from the opportunity due to limited outside options; on the other, they may incur substantial costs to themselves and to society that outweigh any benefits from the attempt, which has a high likelihood of ending in dropout. Finally, in contrast to the limited number of existing studies in the U.S. that use isolated admission cutoffs at a handful of institutions, this paper marshals hundreds of cutoffs spanning nearly the entire public university sector of the second largest state over more than a decade. With this substantially larger, more diverse, and more recent sample of marginal applicants and target institutions, I contribute broadly applicable estimates of not only the returns to enrolling the typical marginal public university applicant, but also, as I detail below, the economically distinct contributions of the extensive margin of attending any four-year college versus the "intensive" margin of attending a more selective four-year college.

The paper proceeds as follows. After describing the data sources and linkages, Section 2 introduces the admission cutoffs in SAT/ACT scores and the regression discontinuity research design they enable. On average, among applicants to a given university in a given year, scoring just above rather than just below an admission cutoff causes the probability of admission to jump abruptly by 27 percentage points, leading to a precisely estimated 14 percentage point first stage in the probability of enrolling at that institution. The density of applicants and their observable characteristics are smooth through these cutoffs, justifying the use of local cutoff-crossing as an exogenous instrument for enrollment. The causal effects identified by this strategy pertain to compliers who enroll in a target institution if and only if they barely cross that institution's admission cutoff; I show that the observable characteristics of these cutoff compliers are similar to all marginal applicants around the cutoffs (i.e., including always-takers who are below the cutoff but enroll and never-takers who are above the cutoff and do not). As expected, the average cutoff complier is significantly more disadvantaged than the average college applicant in terms of academic ability and family background, and more comparable to the average high school graduate.

The main causal estimates in Section 3 begin by documenting that the applicants who are marginal to a cutoff at a given public university have a range of fallback options if they don't get in. Notably, about half of cutoff compliers are able to fall back to another Texas four-year institution if rejected. The pooled results that follow are therefore a nearly equally weighted mix of both "intensive margin" effects of starting at a more-preferred versus less-preferred four-year institution and "extensive margin" effects of starting at any four-year institution. Section 5 develops a method to disentangle the causal contributions of these two margins; the results until then keep them pooled as co-contributors to the policy-relevant return to enrolling the "average marginal" public university applicant relative to their next-best alternative. Cutoff-crossing induces large changes in traditional measures of the educational quality of the institutions student attend, including peer test scores,

gross tuition, peer BA completion rates, and peer mean earnings. Interestingly, only 6 percent of all cutoff compliers forego higher education altogether if rejected; the vast majority of compliers on the extensive margin of attending any four-year college have a two-year community college as their next-best alternative. This is a noteworthy result in itself: the empirically relevant extensive margin among marginal university applicants is between the four-year sector and the two-year sector, rather than no college at all.

The remainder of Section 3 explores longer-run consequences on educational attainment, costs, and earnings. Cutoff-crossing substantially alters the cumulative educational attainment of marginal students, with cutoff compliers ultimately completing one full year's worth of additional credits in the four-year sector and becoming 12 percentage points more likely to ever earn a bachelor's degree from any institution. Some of this comes at the expense of reduced attainment in the two-year sector, with about half a year's worth of fewer credits at two-year schools and a 7 percentage point reduction in associate's degree or certificate completion, but overall cutoff compliers end up with meaningfully more postsecondary education in terms of both quantity and quality.

Calculating the cumulative costs of these additional educational investments, I find that marginally admitted students actually pay no additional net tuition: they incur an additional $4,000 in gross tuition charges but receive a roughly equal amount of additional grant aid to nullify any incremental net cost. Marginal students do end up taking out an additional $5,000 in total student loans, likely to finance room and board charges and other consumption during college. From society's perspective, of course, the additional educational investments in marginal students are not free; I estimate that the average cutoff complier generates over $10,000 of additional costs of educating students in the four-year sector, but reduces costs in the two-year sector by around $3,000. I return to these results below as inputs into the formal cost-benefit calculations in Section 4.

Turning to labor market trajectories, I first show that marginal students experience no change in their likelihood of appearing in the Texas earnings data up to 14 years out from the initial application, which assuages concerns about differential attrition induced by cutoff crossing (Foote and Stange, 2022). I then trace out the earnings trajectories of cutoff compliers in three different ways, measuring annual Texas earnings inclusive of zeros, annualized positive earnings, and earnings rank within each statewide high school graduating cohort. All three of these earnings measures deliver a similar pattern of dynamic effects. Initially, admitted compliers earn less than their rejected counterparts in the first four years after application, as they are much more likely to be actively enrolled in a four-year program. In years five through seven, admitted compliers catch up to their rejected counterparts but do not yet experience any earnings premium. An earnings premium does eventually emerge, however, starting eight years out from application and persisting steadily thereafter, though with some statistical imprecision across individual years. When averaging each earnings measure within each individual across 8-14 years out from application, the estimates imply proportional gains between 5-10 percent and 2.5 percentiles in the rank distribution. Since cutoff compliers complete one additional year of education in the four-year sector and half a year less in the two-year sector, a back-of-the-envelope calculation would imply a gross return to a year of

3

college ranging from 5-10 percent if two-year credits are not valued at all in the labor market, up to 10-20 percent if two-year and four-year credits are equally valued. The final results in Section 3 show the robustness of the RD estimates across a battery of alternative specifications, including a wide range of bandwidths, control sets, and local polynomial functional forms.

In Section 4, I gather these results into a cost-benefit analysis of the private and social returns to enrolling marginal public university students, taking into account the dynamics of the cost and benefit flows and the cost externalities from diverting students from two-year colleges and cheaper four-year institutions. The undiscounted benefits of enrolling marginal students eventually surpass the costs, but at different horizons for students, society, and taxpayers. For students, the initial earnings foregone by increased enrollment keep their benefits below their (zero) costs until 8 years out from application. After 8 years, the net gain is positive and steadily grows with a roughly linear trajectory. Society must wait longer for these benefits to outweigh the resource costs of educating the marginal students, which occurs between years 12 and 13.[2] For taxpayers, the additional tax revenue generated from the earnings gains surpasses the costs after 25 years. Under the conservative assumption that the annual earnings effect will remain at the pooled years 8-14 estimate of $2,800 from year 15 until retirement (instead of growing proportionally with higher earnings levels), the undiscounted lifetime value of the gross earnings benefit is roughly $108,000, which then declines quickly as the discount rate increases. The intersections of the present values of benefits and costs across discount rates yield the relevant internal rates of return: a substantial 23 percent for students, 12 percent for society, and 4 percent for the government budget. If I use the annualized positive earnings measure instead of the annual Texas earnings measure, the IRRs are similar but slightly lower at 19 percent for students, 10 percent for society, and 3 percent for taxpayers.

Finally, in Section 5, I return to the fact that nearly half of cutoff compliers would fall back to another four-year institution if barely rejected from the target university, while the other half would initially fall out of the four-year sector. Since these distinct complier types on the "intensive" versus "extensive" margins of the four-year sector may experience substantially different effects of enrolling in the target university, and these effects are distinct economic parameters of interest, I develop a method to disentangle their separate contributions to the "average marginal" effects summarized above. Intensive and extensive margin compliers are not directly distinguishable in the data, since the counterfactual fallback option for each marginal applicant is nowhere recorded. However, the fact that I observe an applicant's entire portfolio of admissions to any of the 35 Texas public universities enables a powerful stratification of marginal applicants into two observable groups: those who have at least one admission offer from *another* Texas public university, and those who have none. This stratification does not perfectly separate intensive and extensive margin compliers, but it comes fairly close, and far more so than other observable stratifiers like the identity of the target school or an applicant's pre-college covariates. Thus, a seemingly straightforward strategy to identify separate treatment effects for (mostly) intensive margin compliers versus

---

[2]Defining the social gross benefit as simply the pre-tax earnings gain would overstate the true social benefit to the extent that some of the private earnings gain is pure signaling (Aryal et al., 2022), but understate it in the presence of benefits beyond earnings like less crime and better health (Oreopoulos and Salvanes, 2011).

(almost entirely) extensive margin compliers would simply divide the sample into applicants with and without another four-year admission offer and run the main fuzzy RD specification in each. Unfortunately, the other-admission stratifier itself is somewhat endogenous to cutoff crossing, likely caused by the availability of rolling admissions and spring admission cycles at many Texas public universities; marginal applicants who are barely rejected at their target university often still have time to secure an admission to a fallback school within the same application year. The question I answer in Section 5, then, is what can we learn about intensive versus extensive margin treatment effects with a strong but endogenous stratifier?

The answer, it turns out, is still a great deal. I first show that several, but not all, of the mean potential outcome ingredients that go into the intensive and extensive margin treatment effects of interest are separately identified by adapting the complier-describing method of Abadie (2002) to appropriately-defined "response types" with respect to the effect of cutoff crossing on having any other admissions. This adaptation amounts to running a series of fuzzy RD regressions that replace the outcome with interactions among the outcome, the enrollment treatment, and the other-admission stratifier. I then use the relative rankings of the known mean potential outcomes across the response types, along with the relative rankings of their known mean covariate characteristics, to inform an intuitive rank assumption on the unknown mean potential outcomes that immediately delivers upper and lower bounds on each of the intensive and extensive margin effects of interest. These bounds end up being very tight, nearly point identifying the separate effects. The implied effects for extensive margin compliers are substantially larger than those for intensive margin compliers across all the main outcomes. The pooled effects of enrolling "average marginal" public university applicants are therefore driven overwhelmingly by extensive margin compliers who would not initially enroll in any four-year college if rejected, with much smaller contributions from the intensive margin compliers who would fall back to a less-preferred four-year school.[3]

This paper advances the remarkably small handful of prior studies using exogenous admission variation to study causal impacts of American colleges on educational and labor market outcomes.[4] Closest to this paper is Kozakowski (2023), which studies returns to admission into the least selective Massachusetts state universities using statewide minimum SAT and GPA requirements, estimated on a sample restricted to low-income and minority applicants. This paper studies a massively

---

[3]To the extent there are "ripple effects" in admissions (Gandil, 2024), where admitting an applicant to school A out of her selective fallback school B induces another applicant into her vacated slot at school B, such effects would only exist on the intensive margin, since the fallbacks on the extensive margin are completely non-selective (community colleges and non-enrollment). A full analysis of such ripple effects is infeasible given limited information about applicant rankings over schools; if anything, they would likely modestly increase the overall gains to admitting marginal students, given that some of the intensive margin compliers may induce other students to switch into their vacated four-year slots, who themselves would come from a mix of intensive and extensive margin fallbacks.

[4]A rapidly growing literature exploits admission cutoffs in other countries, many of which combine the choices of institution and field of study in the initial application. See, for example, Saavedra (2009); Hastings et al. (2014); Kirkeboen et al. (2016); Canaan and Mouganie (2018); Anelli (2020); Zimmerman (2019); Sekhri (2020); Jia and Li (2021); Duryea et al. (2023), and Lovenheim and Smith (2023) for a recent review. In the United States, a few papers have studied earnings impacts of statewide changes in admissions policies using difference-in-differences research designs, including Bleemer (2022) and Black et al. (2023). A growing literature uses financial aid eligibility cutoffs to study impacts of aid programs on college student outcomes; see Dynarski et al. (2023) for a recent review along with Galperin (2023) and Londoño-Vélez et al. (2023) for recent contributions.

larger and more diverse public university sector, both in terms of applicants and target institutions, allowing for more statistical precision, student diversity, and institutional breadth in estimating the marginal returns to American public universities. I also develop methods to disentangle the returns reaped by students on the intensive versus extensive margins of the four-year college sector, helping to understand their distinct contributions to the "average marginal" return. Another related paper is Bleemer (2021), which studies educational and earnings impacts of admission to four selective institutions in the University of California system via a "top four percent" policy based on high school class rank. The results in Bleemer (2021) speak primarily to the intensive margin impacts of shifting high-GPA but low-SAT students across different institutions within the four-year sector.[5] This paper marshals broad admissions variation across applicants, target institutions, and the intensive and extensive margins of 4-year enrollment; develops methods for quantifying the distinct contributions of these margins; and conducts formal cost-benefit analyses from the perspectives of students, taxpayers, and society, allowing me to directly answer fundamental questions about marginal returns to American public universities.

The large scope and detailed data in this paper help advance earlier work using admission cutoffs involving more limited sets of institutions and student outcomes. Hoekstra (2009) and Zimmerman (2014) both use RD designs to study earnings returns to college admission, but only to one particular institution—an unnamed state flagship university in Hoekstra (2009), and Florida International University in Zimmerman (2014), the least selective member of Florida's state university system. Goodman et al. (2017) use statewide admission requirements that apply across several public universities in Georgia to study enrollment impacts on educational but not labor market outcomes, while Smith et al. (2020) use the same admission cutoffs merged with credit reports to estimate enrollment impacts on credit scores, a predicted earnings measure, and other measures of financial well-being, but not actual earnings. Daugherty et al. (2014) use high school GPA data from one large Texas school district to study the impacts of marginally qualifying for the state's Top Ten Percent automatic admissions program, but only on enrollment outcomes within a four year window. Altmejd et al. (2021) make use of admission cutoffs in several countries, including the U.S., to study sibling spillovers in college and major choice, but not labor market returns.

More broadly, the exogenous admission variation in this paper contributes credible causal estimates of the returns to education to a literature that has long been concerned about "ability bias" and other confounds in observational comparisons (Noyes, 1945; Becker, 1964; Griliches, 1977; Willis, 1986; Angrist and Krueger, 1991; Card, 2001; Heckman et al., 2018). I am able to empirically verify that a rich set of pre-college covariates are balanced across the admission cutoffs, including direct measures of student ability like high school test scores and advanced coursework, proxies for "non-cognitive" skills like attendance and disciplinary infractions, and demographics like

---

[5]Relatedly, Cohodes and Goodman (2014) study the educational impacts of crossing a Massachusetts merit scholarship eligibility cutoff that shifts high-achieving students across different four-year institutions. A longstanding literature uses various observational research designs to estimate the earnings effects of attending higher "quality" U.S. colleges on the intensive margin (e.g., Brewer and Ehrenberg, 1996; Dale and Krueger, 2002; Dillon and Smith, 2020; Ge et al., 2022; Chetty et al., 2023), with a few recent papers unbundling these effects into college-specific value-added estimates (Cunha and Miller, 2014; Hoxby, 2019; Chetty et al., 2020; Mountjoy and Hickman, 2021).

race, gender, and family income. I also contribute a formal cost-benefit analysis to a literature that often only considers gross treatment effects, advancing the small subset of studies that explicitly estimate internal rates of return to educational investments.[6] Finally, my bounding approach to separately identifying impacts for students on the intensive versus extensive margins of the four-year sector contributes to the small but growing literature that highlights the multiple margins of educational choices and develops methods to disentangle their distinct causal contributions.[7]

2 Data and Research Design

2.1 Setting and Data Sources

The data for this analysis come from linking multiple administrative registries that span the entire state of Texas, maintained by the University of Texas at Dallas Education Research Center (UTD-ERC, 2022). Texas is the second largest U.S. state by land area and population (30 million) and the fastest growing state in population level (nearly 500,000 net increase annually). If it were its own country, Texas would comprise the 9th largest economy in the world ($2.4 trillion GDP).

The analysis sample begins with the universe of students who graduate from a Texas public high school between 2004 and 2016. The 2004 cohort is the first to have SAT and ACT scores recorded in the college admission records, and I stop at the 2016 cohort to observe at least six years of post-application outcomes for all sample members. I link several student registries maintained by the Texas Education Agency (TEA) to assemble pre-college data on these students' demographics, standardized test scores, high school coursework, attendance, disciplinary infractions, and high school campus.

I then link these high school graduates to administrative application and admission records from all 35 Texas public 4-year universities, maintained by the Texas Higher Education Coordinating Board (THECB).[8] These application-level records include the admission decision and the SAT/ACT score used in that decision, which together enable the regression discontinuity research design described below.[9]

I follow these students longitudinally through additional THECB administrative registries of college enrollment spells, credit accumulation, degree completion, and financial aid spanning all public and private non-profit postsecondary institutions in the state. Importantly, these data allow me to observe educational outcomes regardless of transfer across institutions. For the 2008-

---

[6]E.g., Becker (1964); Heckman et al. (2006, 2008); Zimmerman (2014); Barrow and Malamud (2015); Bhuller et al. (2017); Hoxby (2018); Ost et al. (2018); Kozakowski (2023).

[7]Rouse (1995); Miller (2007); Heckman and Urzua (2010); Brand et al. (2014); Feller et al. (2016); Kline and Walters (2016); Kirkeboen et al. (2016); Hull (2018); Mountjoy (2022); Galperin (2023); Kamat (2023); Lee and Salanié (2023).

[8]The THECB does not collect student-level data on applications or admissions to private Texas colleges, which enroll roughly 17% of four-year college-goers in Texas, but I do observe the universe of enrollments and degree completions at these schools, allowing me to track public university applicants who end up enrolling in them.

[9]For the minority of applicants who submit both an SAT score and an ACT score, I convert the SAT score to an ACT score using the concordance table published by the College Board (2009) and use the test with the higher of the two values, which seems to align with how admissions offices treated these applications.

2016 cohorts, National Student Clearinghouse records are available that track the initial college enrollments of all Texas public high school graduates across nearly all colleges in the United States, allowing me to distinguish between not enrolling in any Texas college and not enrolling in any college anywhere. I also observe each student's annual financial aid package, which I augment with annual institution-level data from IPEDS (National Center for Education Statistics, 2022) to construct student-year-level cost measures of gross tuition, grant aid, net tuition, loan accumulation, and colleges' educational expenditures.

Finally, to study earnings trajectories, I merge in quarterly earnings records from the Texas Workforce Commission (TWC) that cover all Texas employees subject to the state unemployment insurance system.[10] Importantly, I show below that crossing the admission cutoffs has precisely zero effect on the probability of appearing in the earnings data, assuaging concerns about endogenous attrition (Foote and Stange, 2022).

2.2 Admission Cutoffs

Nearly all of the public universities in Texas post "assured" or "guaranteed" admissions criteria for first-year applicants on their websites, typically involving a minimum SAT/ACT score for each quartile of high school class GPA rank.[11,12] One might worry that applicants could systematically sort themselves around these publicly advertised admission cutoffs through their test-taking strategies and application decisions, leading to potential violations of the smoothness assumptions underlying the regression discontinuity research design.

In the admissions data, however, large discontinuities in the probability of admission to a given school in a given year often occur at test score values well below the advertised criteria, rendering the "assured" criteria often far from necessary for admission. Using each university's publicly posted criteria to define the cutoffs for the regression discontinuity design would therefore miss many of the actual discontinuities at lower test scores, while also inviting manipulation. I also do not have complete historical data on the criteria posted by each university in each year.

For these reasons, I infer admission cutoffs from the data rather than defining them ex ante, using a procedure similar to Hoekstra (2009), Andrews et al. (2017), Carneiro et al. (2019), Altmejd et al. (2021), Brunner et al. (2023), and Miller (2023), among others. First, I define an application

---

[10]Stevens (2007) estimates that 90% of the civilian labor force is captured in state UI records; excluded are the self-employed, independent contractors, some federal employees including military personnel, and workers in the informal sector.

[11]Applicants in the top GPA decile of their high school class are guaranteed admission to all Texas public universities (except UT-Austin, which has a stricter threshold), regardless of their SAT/ACT scores, thanks to the state's Top Ten Percent law. Neither the TEA nor the THECB collect data on students' exact high school GPAs, precluding a statewide RD analysis of the impacts of Top Ten Percent eligibility; see Daugherty et al. (2014) for an analysis using GPA data collected from one large Texas school district. In the admission files, however, the THECB does record an indicator for whether an applicant was automatically admitted due to their Top Ten Percent status. Since admission cutoffs in SAT/ACT scores are irrelevant for Top Ten Percent students, I drop them from the RD analysis sample; Appendix Figure A.1 shows that Top Ten Percent applicants are relatively rare around the admission cutoffs and are distributed smoothly through them.

[12]After the end of my sample, in the wake of the COVID-19 pandemic, many of these universities have gone "test optional" and now list separate guaranteed admission criteria for applicants with and without SAT/ACT scores.

cell as the combination of target school, application year, top quartile high school GPA status,[13] and submitted test (SAT or ACT). I exclude cells in which virtually all applicants are admitted, as these cells do not have meaningful cutoffs to find, as well as a small number of cells with incomplete admissions data. I also exclude all cells at UT-Austin, as the highly selective flagship's "holistic" admission process does not employ simple cutoffs in SAT/ACT scores in any of my sample years. Within each of the 828 remaining application cells, I estimate a series of local linear RDs centered at each distinct test score value, and then define the cutoff for that cell as the test score value with the largest discontinuity in admission and enrollment. Appendix B describes this procedure in more detail. Porter and Yu (2015) show that it delivers a superconsistent estimator of the true cutoff in each cell, leaving the asymptotic distribution of the second-stage RD estimator unaffected by the fact that the cutoffs are estimated from the data rather than known ex ante. Roughly one-fourth of the candidate cutoffs exhibit statistically significant ($t > 1.96$) discontinuities in the log density of applications or (more rarely) the covariate-predicted BA completion or earnings measures described in the next subsection, likely because these cutoffs actually corresponded to a university's publicly posted criteria in a given application year. I disqualify these potentially publicly known cutoffs from consideration in the search algorithm, but since many individual application cells lack much statistical power to detect manipulation, this exclusion does not mechanically ensure balance when pooling across all cutoffs below.

Figure 1 plots the resulting distribution of admission cutoff locations, separately for SAT submitters and ACT submitters. For context, the lightly shaded histogram in the background shows the population distribution of test scores among all applicants. In the darker shaded foreground, each observation in the histogram is the location of the cutoff within a given application cell, weighted by the number of applications immediately around the cutoff multiplied by that cutoff's estimated first-stage discontinuity in enrollment. Interestingly, some of the most empirically relevant SAT cutoffs are at the round numbers of 700, 800, 900, and 1000; these values are not prioritized in any way by the cutoff location algorithm, but rather reveal the simple rules actually used by many admissions offices to ration admission offers.

2.3   RD Diagnostics

Figures 2, 3, and 4 conduct three important diagnostics on the suitability of these admission cutoffs as the basis of a fuzzy regression discontinuity research design. First, Figure 2 examines the relevance of cutoff-crossing for admission and enrollment at the application's target school. The figure plots the nonparametric probabilities of these outcomes conditional on each unique value of the running variable, defined as the applicant's test score minus the admission cutoff she faces given her admission cell. Large discontinuities emerge clearly at the cutoffs for both SAT and ACT submitters. The bottom panel pools both types of submitters by concording SAT scores to ACT

---

[13]Several schools automatically admit applicants from the top GPA quartile of their high school class, and others more generally set distinct test score cutoffs for top quartile applicants versus applicants outside the top quartile. The admissions data do not universally record the top quartile status of every applicant, but I am able to logically infer it for most applicants, and predict it for the remainder, using the procedure described in Appendix B.

Figure 1: Distribution of SAT and ACT Admission Cutoffs
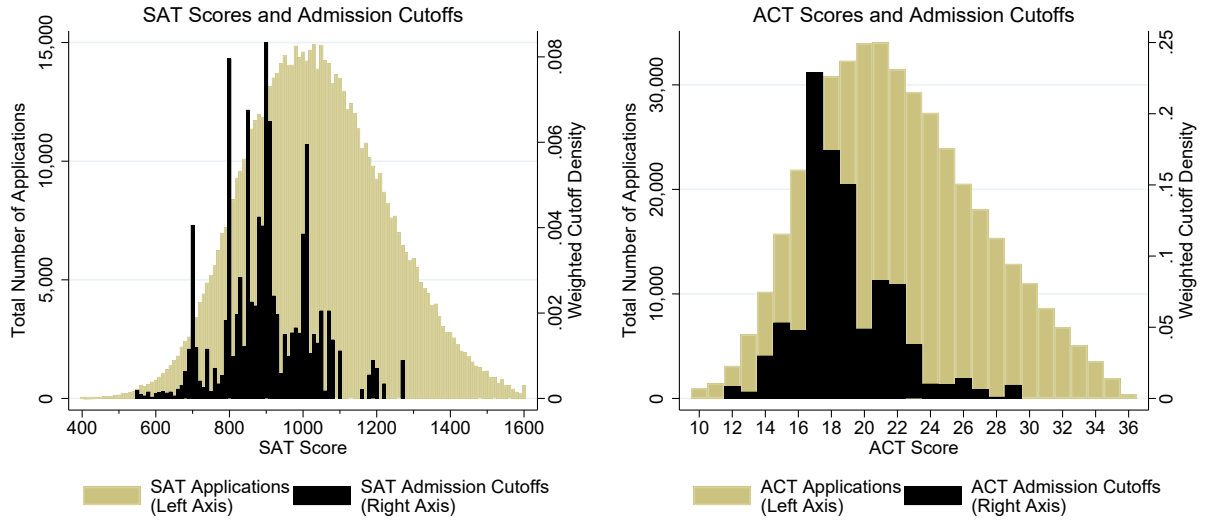
Notes: The lightly shaded histogram in the background shows the population distribution of test scores among all applicants. In the darker shaded foreground, each observation in the histogram is the location of the cutoff within a given application cell, weighted by the number of applications immediately around the cutoff multiplied by that cutoff's estimated first-stage discontinuity in enrollment.

scores and uses the main RD specification, detailed in the next subsection, to estimate a precise enrollment first stage of 14.3 percentage points, with a standard error of 0.3 percentage points and an F-statistic of 2,393.

Second, Figure 3 plots the nonparametric histogram of applications at each unique value of the running variable, separately for SAT and ACT submitters. There are no discontinuities in application density at the cutoffs, formalized by a null McCrary (2008) test in the bottom panel. The frequency counts on the vertical axes of the top two plots also convey the large sample sizes available for this study, and the location of the cutoff within the applicant distribution gives a sense of where the typical marginal applicant tends to rank in the applicant pool.

Finally, Figure 4 examines whether pre-determined student characteristics are balanced across the cutoff. To summarize these characteristics, I first estimate a logit regression of bachelor's degree completion within 8 years on a suite of pre-college covariates.[14] I then use this logit model to construct a predicted probability of BA completion for each applicant, and nonparametrically plot its conditional expectation for each unique value of the running variable. I conduct an analogous exercise for covariate-predicted earnings averaged over 8-14 years out from application. Figure 4 shows that both predicted BA completion and predicted earnings increase strongly across the support of the running variable, highlighting the predictive power of these pre-college covariates, but

---

[14]The pre-college covariates are categorical indicators for gender, race, free or reduced price lunch eligibility, gifted program participation, and ever being classified as at risk of dropping out of high school, as well as cubic polynomials in 10th grade math and English standardized test scores, high school graduation year, the number of advanced courses taken in high school, the percentage of days absent, and the number of days suspended for disciplinary infractions.

Figure 2: RD Diagnostics: First Stage Discontinuities in Admission and Enrollment
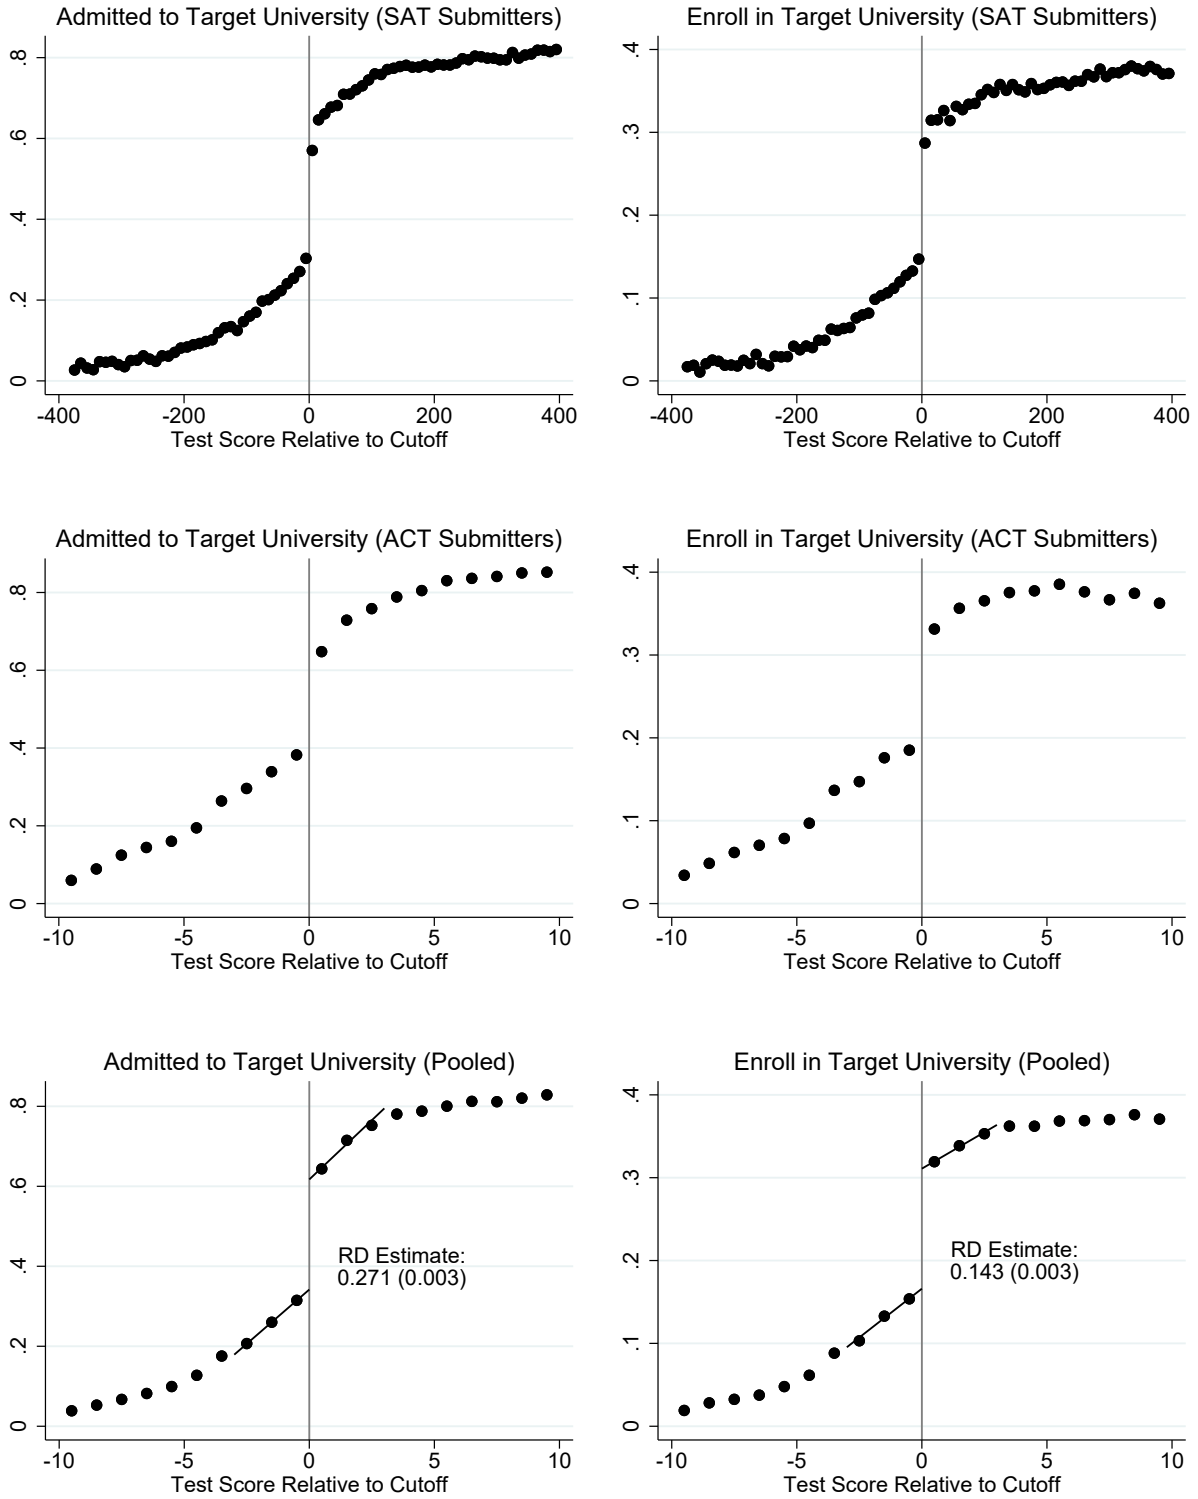

Notes: This figure plots the probability of admission and enrollment for each unique value of the running variable for each test, defined as the applicant's test score minus the admission cutoff she faces given her application cell. The bottom panel pools SAT and ACT submitters by dividing each SAT submitter's running variable by 40 and grouping with the nearest ACT running variable value. The RD estimates come from the main specification described in Section 2.4.

11

Figure 3: RD Diagnostics: Density of Applications
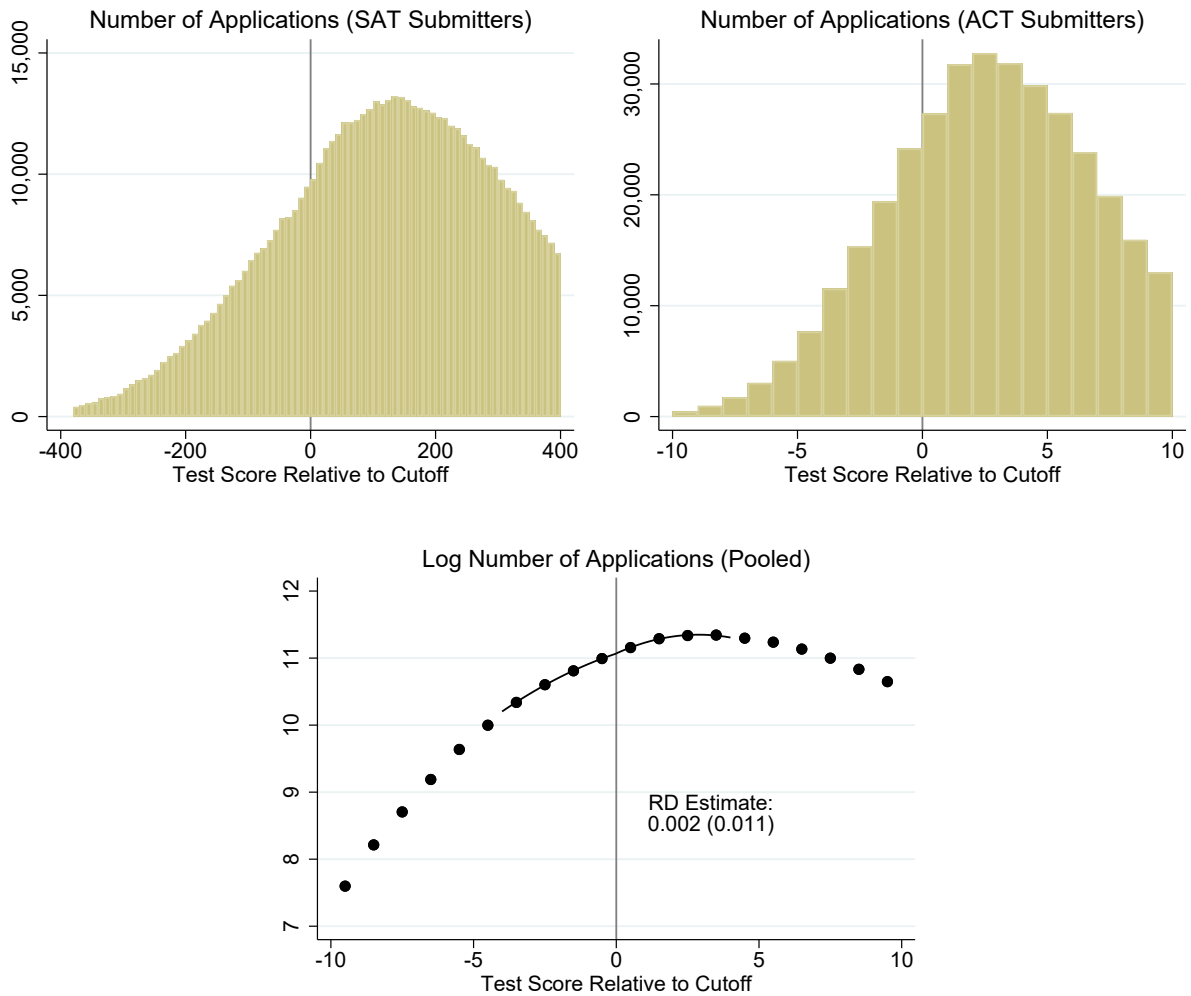
Notes: The top panel plots the nonparametric density of applications at each unique value of the running variable, separately for SAT and ACT submitters. The bottom panel conducts a McCrary (2008) test of manipulation by taking the log number of applications at each value of the concorded running variable and estimating the discontinuity at the cutoff with separate local quadratic functions on each side of the cutoff and a bandwidth of four concorded ACT points.

Figure 4: RD Diagnostics: Covariate Balance
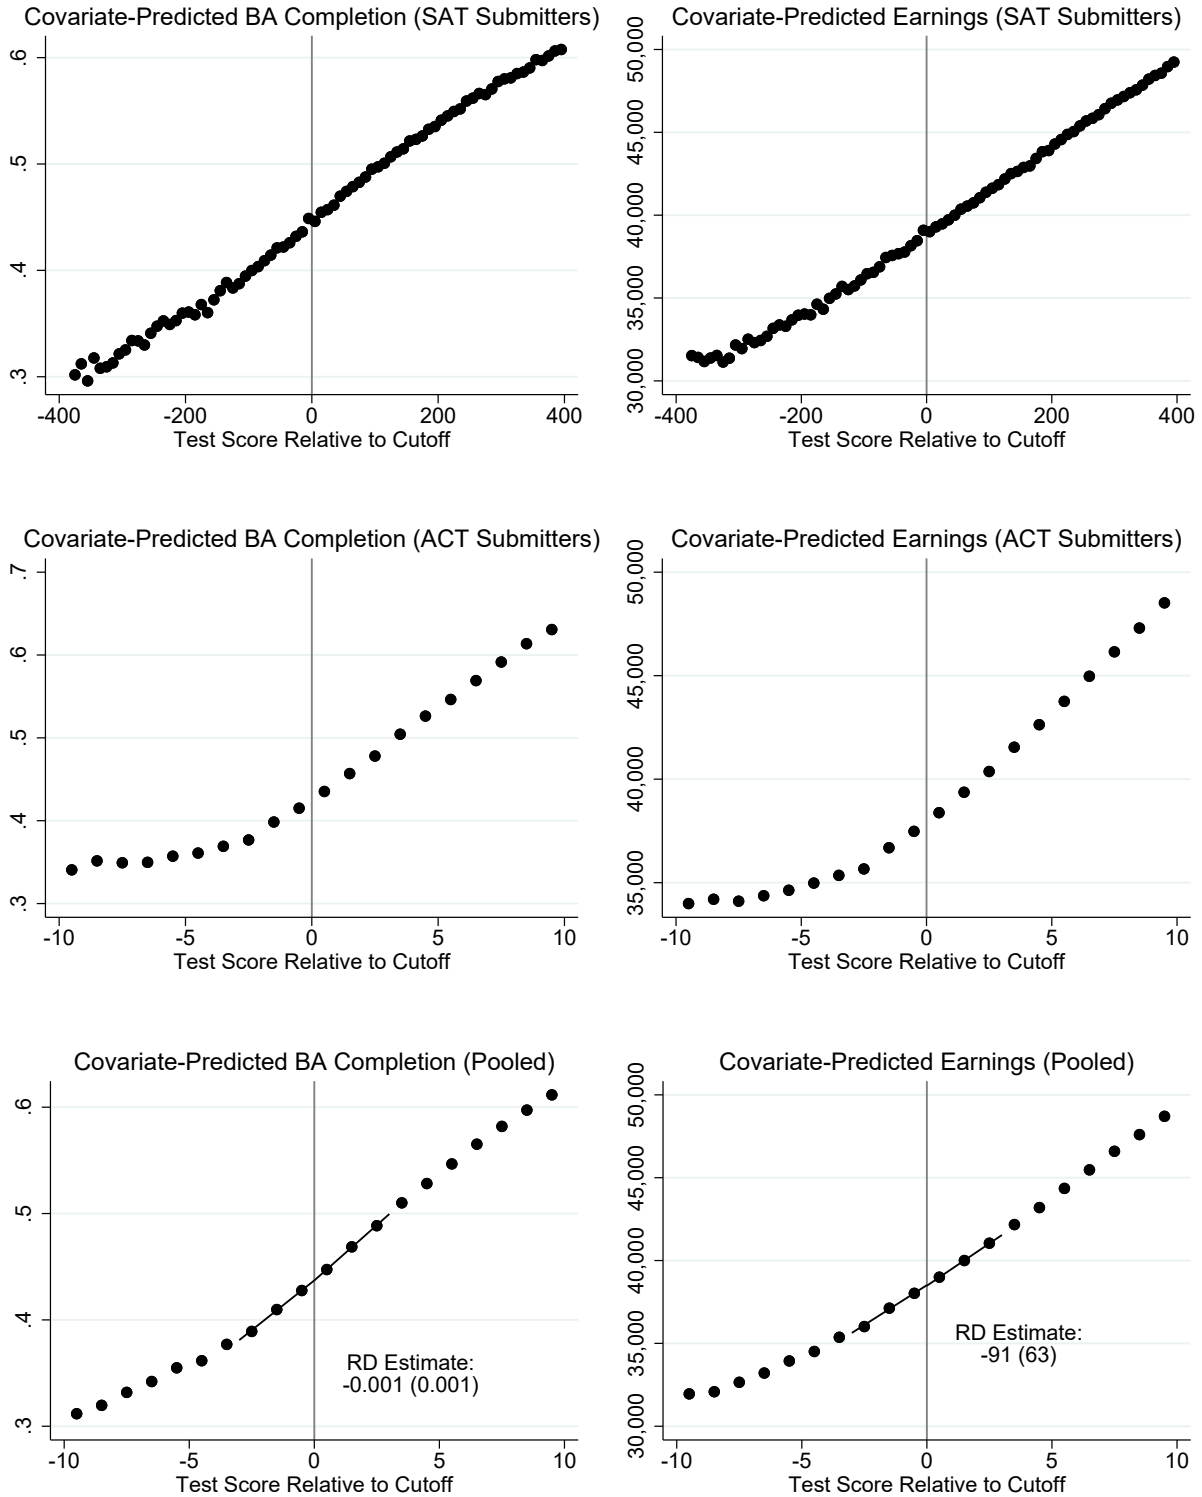
Notes: The left panels plot the covariate-predicted probability of bachelor's degree completion within 8 years of application, predicted via a logit regression using the covariates described in footnote 14. The right panels plot covariate-predicted earnings averaged over 8-14 years after application, predicted via linear regression using the same covariates as BA completion. The bottom panel pools SAT and ACT submitters by dividing each SAT submitter's running variable by 40 and grouping with the nearest ACT running variable value. The RD estimates come from the main specification described in Section 2.4.

13

with no discontinuities at the admission cutoff. Appendix Figure A.2 verifies that this smoothness persists at the level of the individual covariates underlying these predictions. Another important balance consideration is whether students who cross the admission cutoff become discontinuously more or less likely to have observable earnings outcomes; I investigate this below in the context of the earnings effect estimates, and find a precise zero effect of cutoff crossing on appearing in the earnings data (Figure 10), assuaging concerns about endogenous attrition.

2.4 Target Parameters, Identification, and Estimation

To describe and implement the fuzzy regression discontinuity design motivated by the diagnostics above, let D indicate the binary treatment of whether a given applicant to a given target university ends up enrolling at that university. The potential treatments $D_1(r)$ and $D_0(r)$ indicate whether the applicant would enroll if the university's admission cutoff, c, were set exogenously below or above her test score running variable value $R = r$, respectively. Marginal applicants are those who have test scores equal to the target university's cutoff, i.e. $R = c$.[15] An applicant's potential outcome is Y_1 if she enrolls at the target university and Y_0 if she does not; her observed outcome is then $Y = Y_0 + (Y_1 - Y_0)D$. With this notation in hand, the causal parameter of interest is $E[Y_1 - Y_0 | R = c, D_1(c) = 1, D_0(c) = 0]$, the local average treatment effect (LATE) of enrolling at the target university among cutoff compliers, i.e. marginal applicants induced to enroll in the target university by barely crossing its admission cutoff.

This parameter is identified by the fuzzy RD estimand,

$$\frac{\lim_{r \downarrow c} E[Y|R=r] - \lim_{r \uparrow c} E[Y|R=r]}{\lim_{r \downarrow c} E[D|R=r] - \lim_{r \uparrow c} E[D|R=r]}, \tag{1}$$

under the following set of standard assumptions (Hahn et al., 2001; Dong, 2018). First, cutoff-crossing is a relevant instrument for enrolling in the target university; the discontinuities in Figure 2 clearly show the existence of such a first stage. Second, the conditional expectations of the unobservables—potential outcomes Y_1 and Y_0 and potential treatments D_1 and D_0—as functions of the running variable are continuous through the cutoff. While this assumption cannot be tested definitively, continuity of unobservables is supported by the density of applicants and their rich set of observable characteristics both running smoothly through the cutoffs in Figures 3 and 4, and the institutional setting is consistent with the exclusion restriction that crossing the admission cutoff at a target university only affects outcomes through its effect on initial enrollment at that university. The direction of that effect, moreover, is likely to be weakly positive for all marginal applicants, justifying the final assumption of instrument monotonicity.

To maximize precision and summarize the marginal returns to public universities, I pool across all of the application cells described in Section 2.2, with the running variable normalized to zero at the cutoff within each cell and measured in concorded ACT points.[16] Cattaneo et al. (2016) show

---

[15]Only 1 percent of applicants in the RD sample are marginal to more than one university's cutoff, as measured by being within 1 ACT point or 10 SAT points; dropping these applicants has little effect on the results.

[16]I define the cutoff as the midpoint between the two discrete running variable values that straddle the first stage

that this pooled RD estimand identifies a well-defined and clearly interpretable weighted average of cell-specific LATEs, with more weight placed on cells where the applicants at the cutoff are more numerous and more likely to be compliers (i.e. exhibiting a larger first stage discontinuity). My main estimates therefore aggregate across several potentially interesting dimensions of heterogeneity, but these dimensions are not ignored: Section 5 explores heterogeneity across compliers who are on the intensive versus extensive margins of the four-year sector, and a separate paper in progress explores heterogeneity across individual target universities. Until then, the parameter of interest is the overall return to enrolling the "average marginal" public university applicant.

I estimate (1) with local linear approximations of $E[Y|R]$ and $E[D|R]$, differing arbitrarily on either side of the cutoff, within a narrow bandwidth of three concorded ACT points (120 SAT points) and weighted with a triangular kernel. A narrower bandwidth of only two concorded ACT points would have no degrees of freedom for each side's linear fit among ACT submitters, but I show robustness to this narrower bandwidth (which does have degrees of freedom among the finer-grained scores of SAT submitters) in Section 3.5. Because ACT scores are discretely distributed (and technically SAT scores as well, though more finely so), data-driven methods of optimal bandwidth determination and inference that require a continuous running variable may be inappropriate (e.g., Imbens and Kalyanaraman, 2012; Calonico et al., 2014). In practice, I show in Section 3.5 below that the bandwidths selected for each outcome by the method in Calonico et al. (2014), acting as if the running variable were continuous, are always quite close to the bandwidth of three concorded ACT points that I use in the main specification, and the results are similar across a very wide range of alternative bandwidths. The baseline specification requires no additional control variables; I show below that the estimates are very similar when adding detailed controls for pre-college covariates and application cells. Standard errors are clustered at the applicant level, following the reasoning of Kolesár and Rothe (2018) against clustering on the discrete running variable.

2.5   Describing the Sample and the Compliers

With the definition of cutoff compliers and baseline RD specification in hand, Table 1 describes the telescoping populations involved in this study to contextualize the main results that follow. The population begins with the 3.3 million graduates from Texas public high schools from 2004-2016, described in Column 1. Roughly one-third of them apply to a Texas public university during their senior year of high school (Column 2), and just over 300,000 of those applicants qualify for my baseline RD sample (Column 3) by having test scores within three concorded ACT points of a cutoff. Columns 4 and 5 use the baseline RD specification, replacing outcomes Y with pre-determined covariates X, to estimate the characteristics of all marginal applicants immediately at the cutoff (including compliers, always-takers, and never-takers) and the subset of marginal applicants who are cutoff compliers, i.e. enroll in the target university if and only if they barely cross the cutoff. Given the first stage estimate in the top left panel of Figure 5, compliers comprise roughly 14

discontinuity. For SAT submitters, I divide their SAT running variable by 40 to convert it to ACT units.

Table 1: Describing the Sample and the Compliers

	TX Public High School Graduates	Applicants to TX Public Universities	RD Sample: Applicants in Bandwidth	Marginal Applicants at the Cutoff	Enrollment Compliers
Female	.50	.55	.56	.57	.53
White	.42	.45	.35	.32	.29
Hispanic	.40	.34	.37	.38	.33
Black	.14	.13	.22	.24	.31
Asian	.04	.07	.05	.05	.07
FRPL	.49	.37	.46	.48	.46
At-risk	.56	.31	.48	.51	.51
Gifted	.12	.23	.09	.08	.10
HS math (std.)	.10	.62	.11	.06	.03
HS English (std.)	.11	.58	.21	.18	.19
SAT score (1600)	-	1035	919	904	892
ACT score (36)	-	21.8	18.7	18.3	17.8
N	3,303,833	1,089,581	302,238		

Notes: Each column is a subset of the preceding column. The RD sample in Column 3 is comprised of all applicants who face an admission cutoff, have a concorded ACT score within three points of the cutoff in their cell, and are outside the automatically admitted top decile of their high school GPA distribution (see footnote 11). The means of marginal applicants in Column 4 are estimated as the intercept in the reduced form RD specification described in Section 2.4. The means of enrollment compliers in Column 5 are estimated from the main IV specification using the method of Abadie (2002).

percent of marginal applicants, leaving the potential for compliers to differ substantially from marginal applicants more broadly. Comparing Columns 4 and 5, however, shows that compliers are roughly representative of the broader population of marginal applicants. Compliers are more disadvantaged than the average public university applicant in Column 2, as expected given their marginal positions in the applicant pool; they are more comparable to the average high school graduate in Column 1 in terms of academic preparation and family income.

3 Causal Impacts of Enrolling Marginal Applicants

3.1 Institutional Characteristics

The first set of causal estimates show that cutoff compliers enroll in substantially different types of colleges as a result of barely crossing the admission cutoff of the public university at which they are marginal. Figure 5 visualizes the reduced-form effect of cutoff crossing on the sector of the initial college attended in the first academic year after application. Each plot also reports the local average treatment effect (LATE) estimate among cutoff compliers that results from dividing the reduced form discontinuity by the first stage enrollment discontinuity.

In the top left panel of Figure 5, the outcome is enrolling in the target university, so the reduced form discontinuity of 14 percentage points is the size of the enrollment first stage, and the

Figure 5: Impacts on Enrolling in the Target University versus Next-Best Alternatives
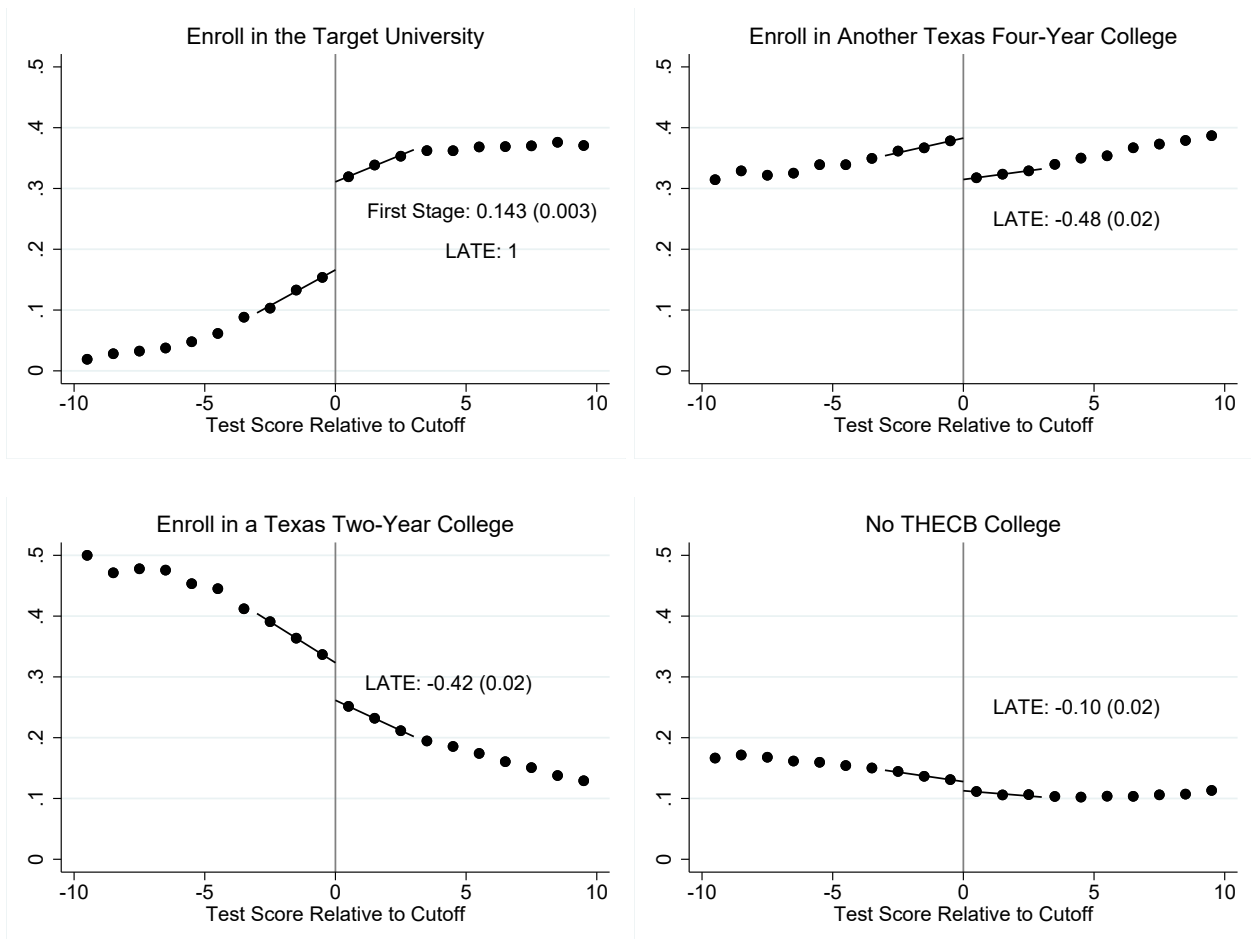
Notes: These sectors correspond to the applicant's first enrollment in the academic year after application. The LATE estimates come from the main fuzzy RD specification described in Section 2.4. The LATE of enrolling in the target university is 1 by construction, since that is the treatment variable. "No THECB College" means not enrolling in any institution in the Texas Higher Education Coordinating Board data, which span all public and non-profit private colleges in Texas. The right panel of Appendix Figure A.3 plots college enrollment outside of the THECB data universe but recorded in the National Student Clearinghouse, available for the younger two-thirds of the sample.

corresponding LATE is 1 by construction: compliers switch from zero probability of enrollment to probability one as a result of barely crossing the cutoff. The LATEs in the other panels of Figure 5 can then be interpreted as a mutually exclusive and exhaustive decomposition of the complier population into types defined by their next-best alternative to enrolling in the target university. The top right panel shows that 48 percent of compliers would fall back to a different Texas four-year college, including public and private schools; in the bottom left, 42 percent would fall back to a Texas two-year community college; and in the bottom right, the remaining 10 percent would not enroll in any Texas public or private college covered by the THECB data. Appendix Figure A.3 uses the younger two-thirds of the sample with National Student Clearinghouse data to show that the majority of that final fallback category is truly attending no college (6 out of the 10 percent), with the remaining 4 percent of untreated compliers attending a college outside of the THECB data

universe but recorded in the National Student Clearinghouse. Appendix Figure A.3 also shows that only 4 out of the 48 percent of fallbacks to another Texas four-year college are to a private institution. Thus, an important takeaway from these results is that cutoff compliers have two main next-best alternatives to enrolling in the target university—other Texas public universities, and two-year community colleges—with far fewer compliers falling back to a private college, going out-of-state, or abandoning higher education altogether. Section 5 develops methods to identify separate effects of enrollment for these distinct complier types.

Figure 6 shows that cutoff compliers experience substantial changes in the characteristics of their peers and popular institutional "quality" measures by enrolling in the target university instead of their next-best alternative.[17] The top panel of Figure 6 shows that the average high school math score of a complier's peers increases by 41 percent of a standard deviation, as measured among the entire population of Texas high school standardized test takers, and those peers are 12 percentage points less likely to have been low-income in high school, as measured by eligibility for free or reduced-price lunch. The middle panel shows that cutoff compliers are propelled into institutions that charge $2,500 more in gross tuition, which is roughly a 45 percent increase relative to the untreated complier mean of $5,500. Those institutions also spend $2,900 more per year educating each student, a 40 percent increase over the untreated complier mean. The bottom panel of Figure 6 turns to average peer outcomes: cutoff compliers experience a dramatic 26 percentage point increase in their peers' 6-year BA completion rate and $6,700 higher peer mean earnings measured 8-14 years after application.

3.2 Enrollment Dynamics, Credit Accumulation, and Degree Completion

The previous results show that marginally admitted students experience large changes in their entry points into higher education. The next set of results show that these initial impacts persist into divergent long-run educational trajectories. In several of the figures that follow, I plot the outcome trajectories of treated compliers (those who fall just above the admission cutoff and therefore enroll) versus untreated compliers (those who fall just below the cutoff and therefore do not enroll) such that the vertical distance between them is the LATE of enrolling at the target university, measured at a given number of years since the initial application. In the Appendix, I show the reduced form RD plots corresponding to each outcome measured 8 years out from application.

The plots in the left column of Figure 7 show that cutoff-crossing has a decisive influence on compliers' long-term engagement with the target university at which they are marginal.[18] Over the span of ten years from the initial application, the top left plot shows that only 10 percent of untreated compliers ever manage to enroll in the target university, meaning that cutoff-crossing in the initial application is nearly a sufficient indicator for whether a complier will ever enroll at that institution. The middle left plot turns to credit accumulation as a fine-grained measure of

---

[17]Applicants who end up enrolling nowhere are included in this analysis by assigning them the mean value of the dependent variable among Texas high school graduates who do not enroll in college.

[18]See Appendix Figure A.4 for the reduced form RD plots corresponding to the outcomes in Figure 7.

Figure 6: Impacts on Peer and Institutional Characteristics
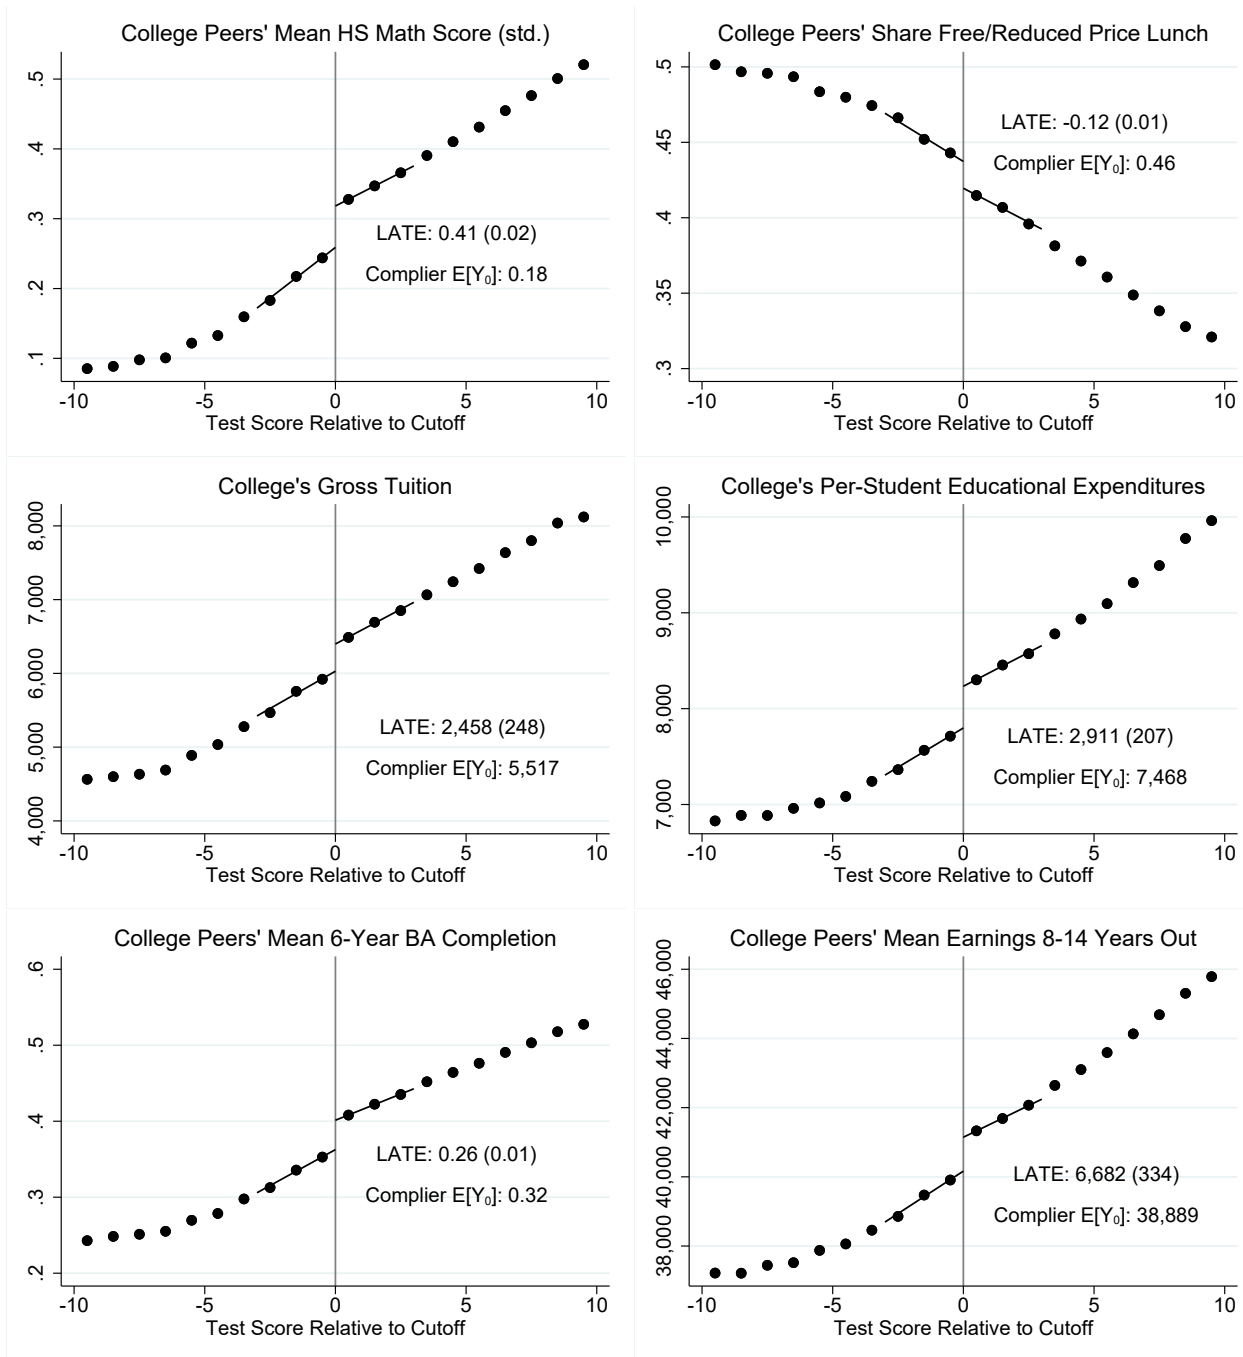
Notes: These college-level characteristics correspond to the applicant's first enrollment in the academic year after application. Applicants who end up enrolling nowhere are included in this analysis by assigning them the mean value of the dependent variable among Texas high school graduates who do not enroll in college. The LATE and untreated complier mean estimates come from the main fuzzy RD specification described in Section 2.4.

Figure 7: Impacts on Long-Run Educational Trajectories
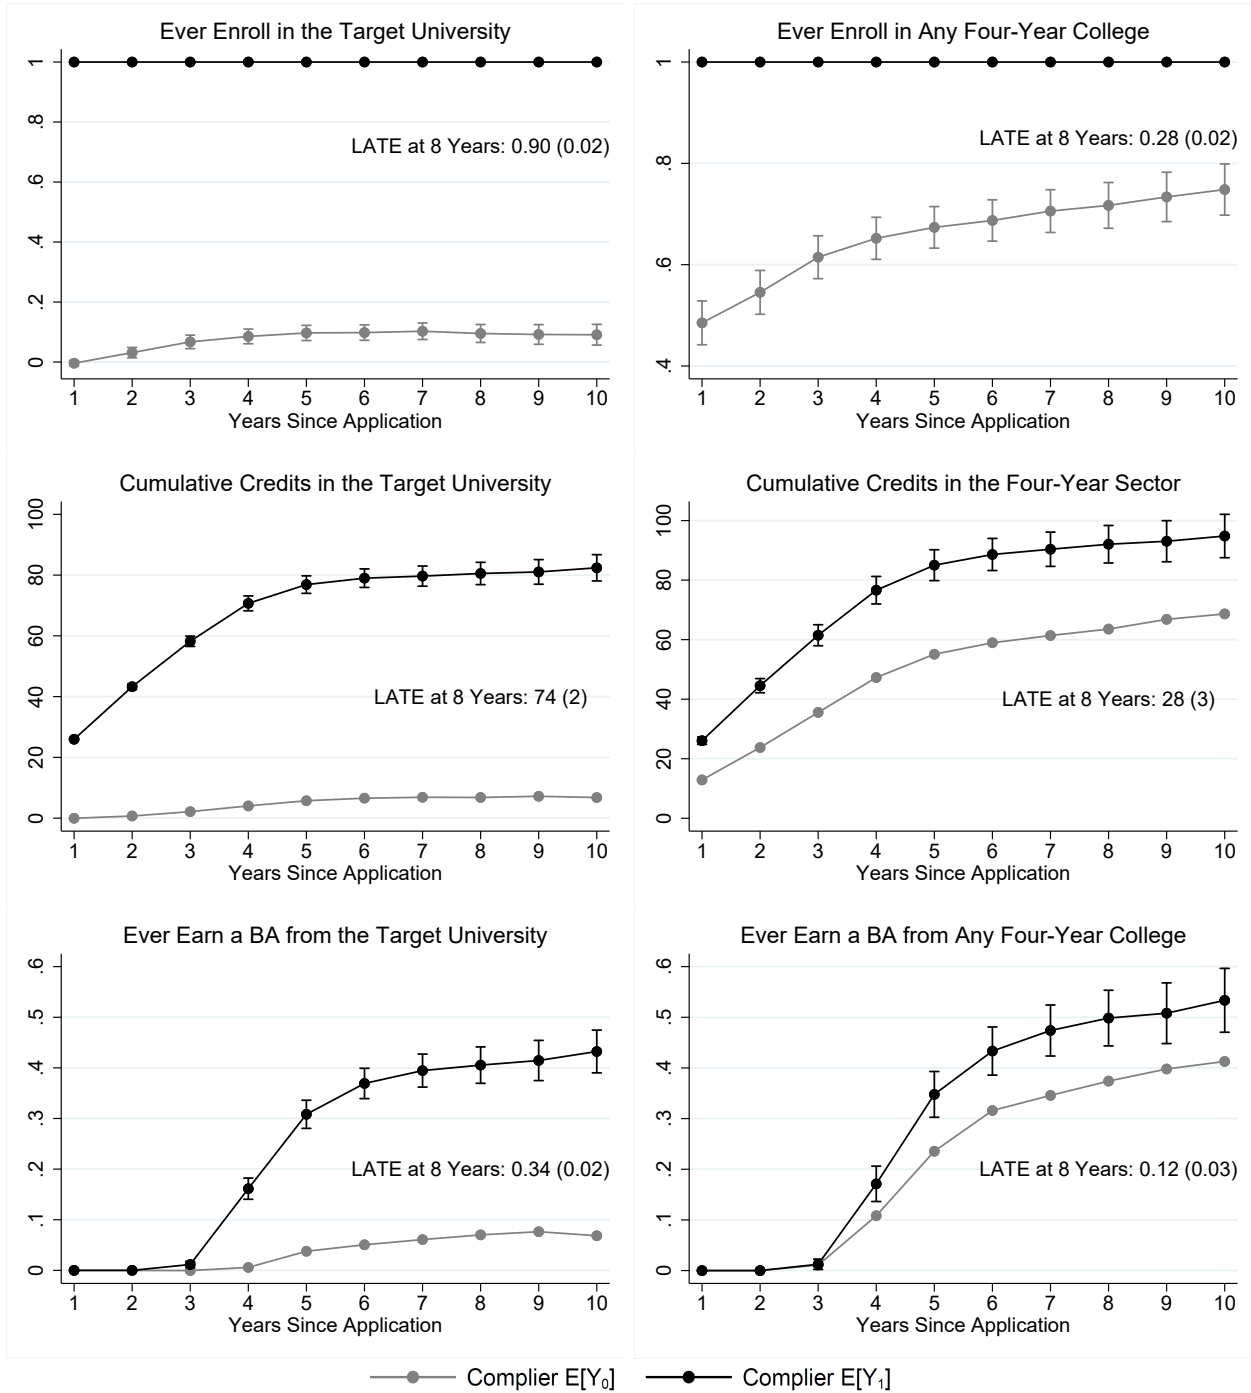
Notes: The gray dots in each plot show the mean outcome of compliers who fall just below the admission cutoff (untreated) at each year since application, and the black dots show the mean outcome of compliers who fall just above the cutoff (treated), which is the untreated mean plus the LATE. The estimates come from the main fuzzy RD specification described in Section 2.4, estimated separately for each year since application. See Appendix Figure A.4 for the reduced form RD plot corresponding to each outcome measured at 8 years.
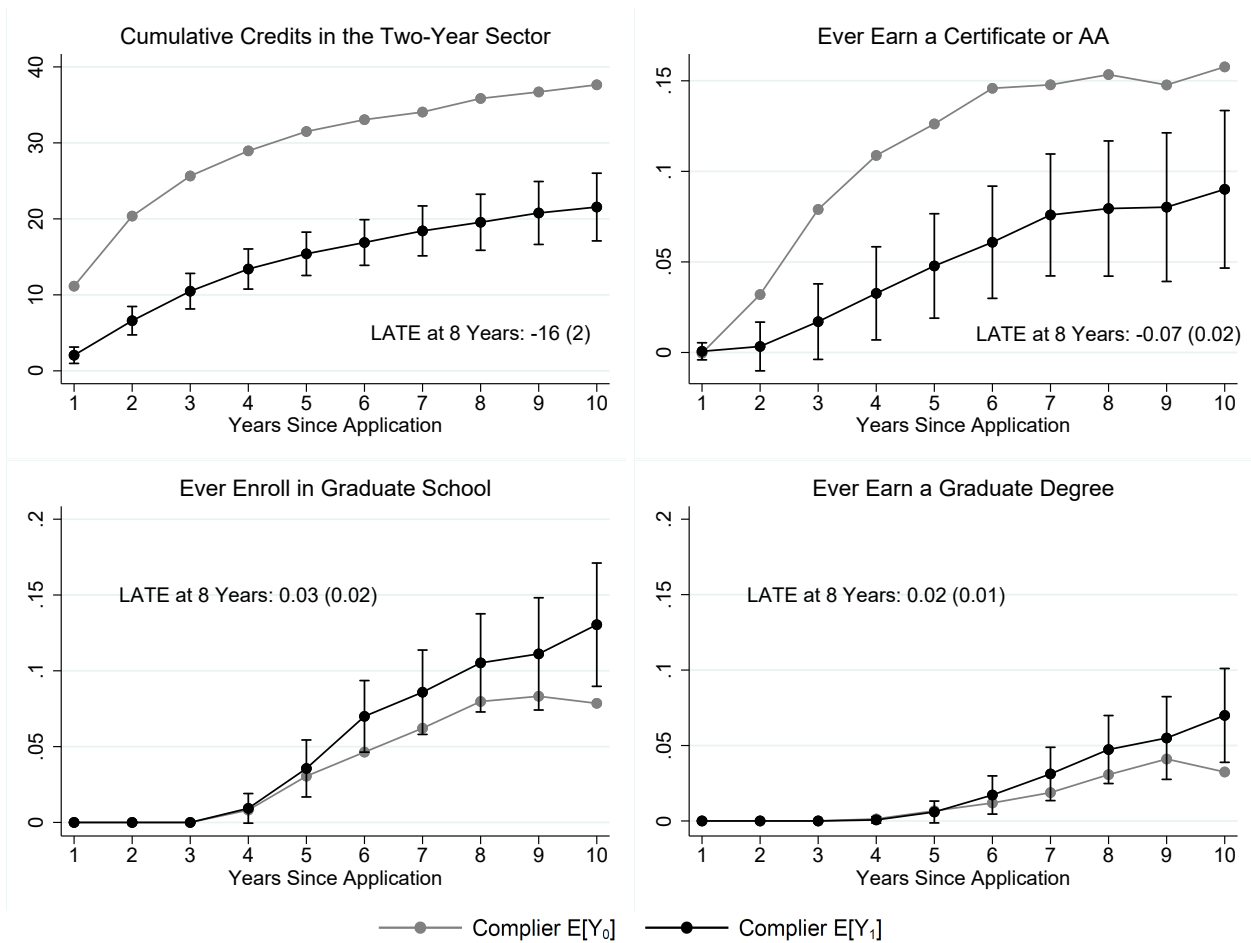
educational attainment, and shows that cutoff-crossing causes compliers to eventually complete 74 more credits at the target university. This is roughly equivalent to 2.5 years of a four-year degree, which requires 120 credits. The bottom left plot shows that cutoff-crossing increases the probability of completing a bachelor's degree at the target university by a dramatic 34 percentage points. In terms of dynamics, only around half of this long-run BA effect appears at the on-time benchmark of four years out, with large gains in years five and six and stabilization around year seven.

The plots in the right column of Figure 7 show that cutoff-crossing also has a large influence on compliers' long-term engagement with the four-year college sector more broadly. The top right plot shows that cutoff-crossing leads to a substantial 28 percentage point increase in the probability that compliers ever enroll in *any* four-year institution. The LATE in the first year after application corresponds to the 52 percent of compliers in Figure 5 who would initially fall back to a community college or no college if they fell just short of the cutoff; the subsequent dynamics of the untreated complier mean show that some of them eventually gain access to a four-year institution, but over one-fourth of compliers *never* set foot in the four-year sector when initially rejected. The middle right plot shows that cutoff-crossing causes compliers to eventually complete 28 more credits at any four-year institution, roughly equivalent to one full year of a four-year program. Comparing the middle left and middle right plots shows that treated compliers complete the vast majority of their four-year credits at the initial target institution, whereas untreated compliers complete the vast majority of their (fewer) credits at other institutions, which is a natural consequence of the sharp enrollment divergence in the top left plot.

The bottom right plot of Figure 7 shows that compliers become 12 percentage points more likely to complete a bachelor's degree from any four-year institution as a result of marginal cutoff crossing. Only a fraction of this effect appears at the on-time benchmark of four years out; the majority of both treated and untreated compliers who ever complete a bachelor's degree do so well after the four-year mark, with the difference (the treatment effect) stabilizing around year six. The levels of the trajectories show the low overall completion rates among this academically marginal population: untreated compliers have only around a 40 percent chance of ever earning a bachelor's degree from any institution, with cutoff-crossing increasing that chance substantially but ultimately to a level just above 50 percent. In terms of majors, Appendix Figure A.5 shows that STEM bachelor's degrees are relatively rare among both treated and untreated compliers, with treated compliers experiencing a small imprecise reduction in the likelihood of STEM BA completion. The overall 12 percentage point gain in bachelor's degrees in the bottom right plot of Figure 7 is therefore driven entirely by additional degrees in non-STEM fields.

The previous results focused on educational trajectories in the four-year undergraduate sector; Figure 8 explores substitution away from two-year community colleges and tests for impacts on graduate education. The top left panel shows that cutoff-crossing causes compliers to complete 16 fewer credits in the two-year sector. Thus, roughly half of the additional credits completed in the four-year sector in Figure 7 are cannibalized from the two-year sector, with the other half comprising net gains in total postsecondary attainment. The top right panel of Figure 8 shows

Figure 8: Impacts on Long-Run Educational Trajectories in the Two-Year and Graduate Sectors

Cumulative Credits in the Two-Year Sector   Ever Earn a Certificate or AA

LATE at 8 Years: -16 (2)   LATE at 8 Years: -0.07 (0.02)

Ever Enroll in Graduate School   Ever Earn a Graduate Degree

LATE at 8 Years: 0.03 (0.02)   LATE at 8 Years: 0.02 (0.01)

Years Since Application

Complier E[Y_0]   Complier E[Y_1]

Notes: The gray dots in each plot show the mean outcome of compliers who fall just below the admission cutoff (untreated) at each year since application, and the black dots show the mean outcome of compliers who fall just above the cutoff (treated), which is the untreated mean plus the LATE. The estimates come from the main fuzzy RD specification described in Section 2.4, estimated separately for each year since application. See Appendix Figure A.6 for the reduced form RD plot corresponding to each outcome measured at 8 years.

that about 15 percent of untreated compliers would eventually complete an associate's degree or certificate from a community college, and cutoff-crossing reduces that rate by 7 percentage points. Thus, some of the gains in four-year bachelor's degrees in Figure 7 come at the expense of shorter degrees, but the substitution is only partial. Finally, the bottom two plots show weakly positive effects on graduate school enrollment (left) and graduate degree completion (right), neither of which are very common among this academically marginal population.

3.3 Costs

How much do these additional educational investments cost the marginal student, and society? The top left panel of Figure 9 begins by showing that cutoff crossing ultimately adds around $4,000 to the average complier's gross tuition charges, ignoring financial aid, accumulated across all full and

Figure 9: Impacts on Long-Run Educational Costs
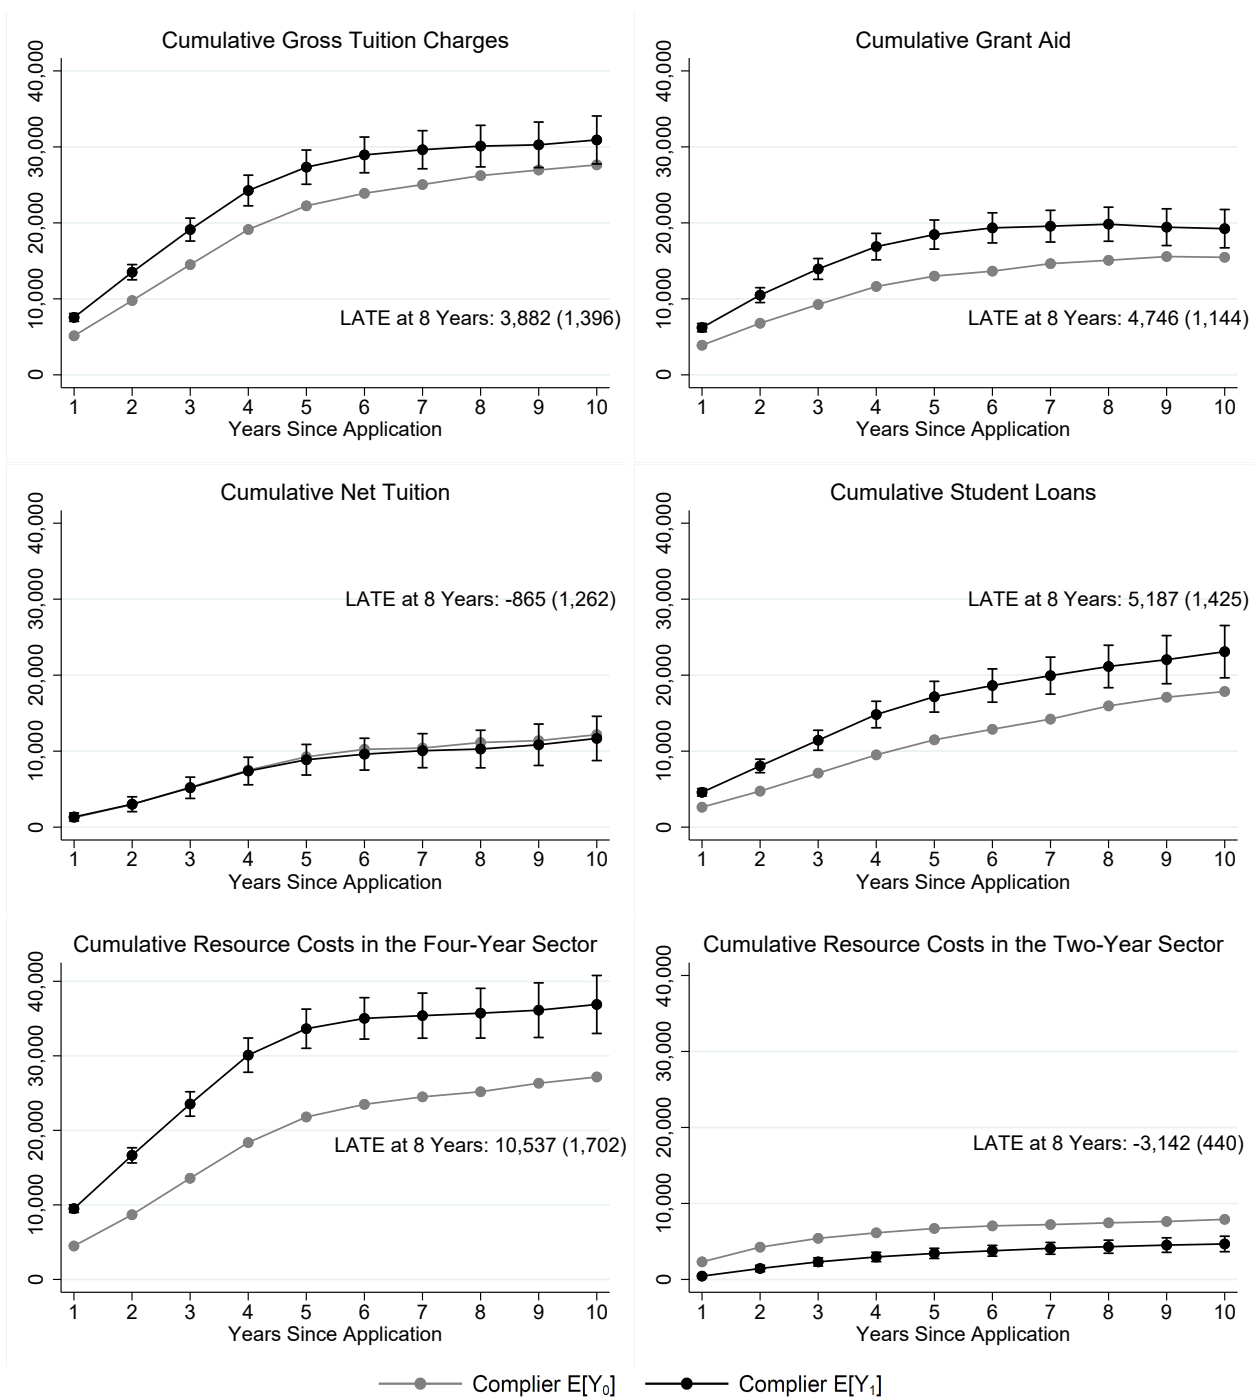
Notes: The gray dots in each plot show the mean outcome of compliers who fall just below the admission cutoff (untreated) at each year since application, and the black dots show the mean outcome of compliers who fall just above the cutoff (treated), which is the untreated mean plus the LATE. The estimates come from the main fuzzy RD specification described in Section 2.4, estimated separately for each year since application. The gross tuition and resource cost measures are pro-rated each semester according to the individual student's enrollment intensity relative to full-time. Gross tuition charges include mandatory fees and the college's estimated cost of required books and supplies, but exclude room and board. All students are assumed to pay in-state tuition rates at public institutions, and community college students are assumed to pay in-district rates. See Appendix Figure A.7 for the reduced form RD plot corresponding to each outcome measured at 8 years.

23

partial semesters enrolled.[19] Many students receive grant aid to offset these gross tuition costs, however, and the top right panel shows that the additional accumulation of grants actually fully offsets the additional tuition charges, leading to no detectable increase in cumulative net tuition in the middle left plot. Remarkably, then, none of the additional tuition cost induced by cutoff crossing is borne by the average marginal student herself. Compliers do end up taking out roughly $5,000 in additional loans, as shown in the middle right panel, which are likely used to finance room and board charges and other consumption during college.

From society's perspective, of course, these investments are not free; the additional resources required to educate marginally admitted students could have been used for other purposes. At the same time, the results in Figure 7 imply that some of the additional costs generated in the four-year sector are offset by reduced costs in the two-year sector. The bottom two panels of Figure 9 quantify these increments to society's ledger: the average cutoff complier ultimately generates over $10,000 of additional resource costs in the four-year sector, but reduces costs in the two-year sector by around $3,000, leading to an individual-level average net increase of around $8,000 in postsecondary resource costs. These estimates assume that the marginal cost of educating an additional student at a given school in a given year at a given enrollment intensity is equal to the observed average cost of doing so, which is a strong but common assumption, given the difficulty of accurately measuring marginal costs. In Section 4, I gather these cost estimates together with the earnings results below to conduct formal cost-benefit calculations from the perspectives of the marginal students themselves, taxpayers, and society.
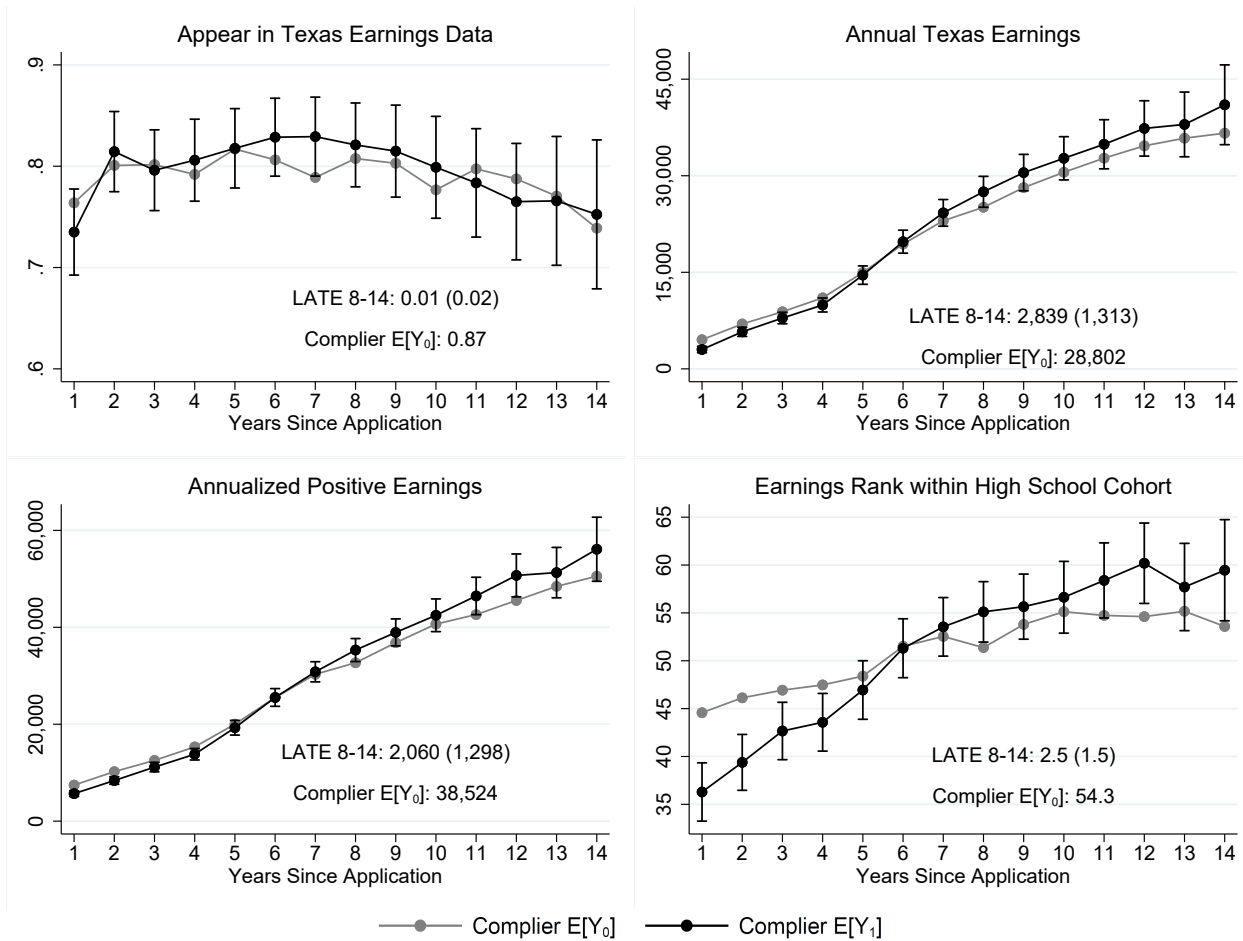
3.4 Earnings Trajectories

Do the divergent educational trajectories induced by cutoff crossing translate into divergent earnings trajectories? The top left panel of Figure 10 begins by showing no detectable difference in the probability of appearing in the Texas earnings data across treated versus untreated compliers; roughly 80 percent of both appear in the earnings data in any given year, and 87 percent appear at some point between 8-14 years out from application. As is common in state-level earnings analyses, individuals with missing earnings in a given year could either be working outside of Texas or not working anywhere, so it is difficult to distinguish true missings from true zeros. The remaining plots in Figure 10 therefore present three different ways of measuring complier earnings trajectories. The top right plot considers annual Texas earnings, which takes the state's perspective and simply adds up the real earnings observed for each individual each year, assigning a value of zero to any quarter in which no Texas earnings are recorded. Since annual Texas earnings may underestimate some individuals' total earnings, the bottom left plot drops the quarters with no Texas earnings, averages the remaining quarters of real positive Texas earnings within a year, and multiplies by four to get an annualized measure of positive earnings. Finally, the bottom right plot computes

---

[19]The gross tuition and resource cost measures are pro-rated each semester according to the individual student's enrollment intensity relative to full-time. Gross tuition charges include mandatory fees and the college's estimated cost of required books and supplies, but exclude room and board. All students are assumed to pay in-state tuition rates, and community college students are assumed to pay in-district rates.

Figure 10: Impacts on Earnings Trajectories

Appear in Texas Earnings Data

LATE 8-14: 0.01 (0.02)

Complier E[Y_0]: 0.87

Years Since Application

Annual Texas Earnings

LATE 8-14: 2,839 (1,313)

Complier E[Y_0]: 28,802

Years Since Application

Annualized Positive Earnings

LATE 8-14: 2,060 (1,298)

Complier E[Y_0]: 38,524

Years Since Application

Earnings Rank within High School Cohort

LATE 8-14: 2.5 (1.5)

Complier E[Y_0]: 54.3

Years Since Application

Complier E[Y_0] ——•—— Complier E[Y_1]

Notes: The gray dots in each plot show the mean outcome of compliers who fall just below the admission cutoff (untreated) at each year since application, and the black dots show the mean outcome of compliers who fall just above the cutoff (treated), which is the untreated mean plus the LATE. The estimates come from the main fuzzy RD specification described in Section 2.4, estimated separately for each year since application. Annual Texas earnings add up the real earnings observed for each individual each year, assigning a value of zero to any quarter in which no Texas earnings are recorded. Annualized positive earnings drop the quarters with no Texas earnings, average the remaining quarters of real positive Texas earnings within a year, and multiply by four to get an annualized measure. Earnings ranks correspond to each individual's percentile rank in the annualized positive earnings distribution of their statewide high school graduating cohort. Earnings are winsorized at the 99th percentile and measured in real 2015 dollars. See Appendix Figure A.8 for the reduced form RD plot corresponding to each outcome averaged over 8-14 years.

each individual's rank in the annualized positive earnings distribution of their statewide high school graduating cohort.

All three of these earnings measures deliver a similar pattern of dynamic effects for cutoff compliers. Initially, admitted compliers earn less their rejected counterparts in the first four years after application, as they are much more likely to be actively enrolled in a four-year program. In years five through seven, as many of the admitted compliers are finishing their college education and entering the workforce full-time, they do not yet experience any earnings premium over their rejected

counterparts. An earnings premium does eventually emerge, however, starting at eight years out from application and persisting steadily thereafter, though with some statistical imprecision across individual years. When averaging each earnings measure within each individual across 8-14 years out from application, the LATEs and untreated complier mean potential outcome estimates imply that admitted compliers earn somewhere between 5-10 percent more than their rejected counterparts, and the rank measure shows a gain of 2.5 percentiles in the relative distribution. Appendix Figure A.8 shows the reduced form RD plot corresponding to each outcome averaged over 8-14 years.

A few back-of-the-envelope calculations help to interpret these earnings effects vis-a-vis the gains in educational attainment. Cutoff compliers gain one full year of education in the four-year sector; if this is the exclusive driver of the 5-10 percent earnings gain, that would imply a 5-10 percent gross return to a year of college. Cutoff compliers do complete half a year less of education in the two-year sector, however, so the implied gross return to a year of college would range from a lower bound of 5-10 percent if two-year credits are not valued at all in the labor market up to an upper bound of 10-20 percent if two-year and four-year credits are equally valued, as in that case the total gain in education is half a year. More daringly, if the exclusive driver of the 5-10 percent earnings gain is the 12 percentage point gain in BA completion, the implied gross return to a bachelor's degree would be in the ballpark of 40-80 percent (dividing 5-10 percent by 0.12). Compliers complete fewer associate's degrees and certificates, however, so the implied gross return to a BA would be higher to the extent that the foregone sub-baccalaureate degrees would have had positive returns. These simple back-of-the envelope calculations ignore a host of potential exclusion violations, like increases in educational quality conditional on years of attainment or degree completion, but they turn out to fall within the range of prior estimates using other research designs (e.g., Card, 2001; Barrow and Malamud, 2015).

3.5   Specification Checks

Appendix Figures A.9 and A.10 investigate the robustness of these results across a battery of alternative specifications. Figure A.9 plots the main estimates and their 95 percent confidence intervals for the baseline bandwidth of three concorded ACT points, along with the estimates resulting from a wide range of alternative bandwidths. I also mark the bandwidth selected for each outcome by the data-driven method in Calonico et al. (2014), acting as if the discrete running variable were continuous; this bandwidth is always close to the main specification's bandwidth of three concorded ACT points. The point estimates are similar across this wide range of bandwidths, with expectedly less precision at smaller bandwidths and more precision at larger ones.

Figure A.10 explores several alternative ways of structuring the regression specification and the control set. The baseline specification solely controls for linear slopes in the running variable on either side of the cutoff. The estimates are virtually unchanged when adding the full suite of pre-college covariates described in footnote 14. The results are also very similar when adding fixed effects for each of the 828 application cells, defined as the unique combinations of target school,

application year, top quartile high school GPA status, and submitted test (SAT or ACT). The next specification adds saturated interactions between the application cell indicators and the running variable slopes, allowing for arbitrarily different local linear parameters across each application cell, and again finds very similar results. Finally, I change the local polynomial functional form from linear to quadratic in the running variable on either side of the cutoff and find similar results, including across bandwidths of five, four, and three concorded ACT points, with three points offering no degrees of freedom for ACT submitters and more imprecise estimates.

4 Cost-Benefit Analysis

The results in the previous section show that marginally admitted students eventually reap positive gross earnings returns, but they take many years to materialize. And while the marginal students themselves tend to pay no additional tuition after grant aid, they generate thousands of dollars of additional education costs for society and the government budget that could have been used for other purposes. In this section, I use the causal cost and earnings estimates as inputs into a cost-benefit analysis to calculate the net returns to enrolling marginal students from the perspectives of the individual students themselves, taxpayers, and society.

The top panel of Figure 11 begins by plotting the dynamics of the undiscounted cumulative social benefit of enrolling the average cutoff complier, which is simply the running accumulation of the annual pre-tax earnings effects in Figure 10.[20] This definition of the social gross benefit would overstate the true social benefit to the extent that some of the private earnings gain is pure signaling (Aryal et al., 2022), but understate it in the presence of benefits beyond earnings like less crime and better health (Oreopoulos and Salvanes, 2011), along with any consumption value of the college experience itself (Gong et al., 2021; Aucejo et al., 2023). Those gross social benefits are then decomposed into the student benefit (the post-tax earnings gain) and the taxpayer benefit (the tax revenue gain), assuming a 20 percent tax and transfer rate on the earnings increments (Hendren and Sprung-Keyser, 2020; Angrist et al., 2022).[21]

On the other side of the ledger, the top panel of Figure 11 also shows the two relevant cost measures: the cumulative private cost to the student, which is the (roughly zero) additional net tuition paid by marginally admitted cutoff compliers,[22,23] and the cumulative social/taxpayer cost,

---

[20]I use the annual Texas earnings measure from the top right panel of Figure 10, but all of the calculations are similar if I use the annualized positive earnings measure instead, as shown in Appendix Figure A.12.

[21]Since the net tuition cost to students is essentially zero, a constant proportional tax on earnings does not affect the student-perspective crossing point and internal rate of return calculations that follow, but a progressive marginal tax would (Heckman et al., 1998; Bhuller et al., 2017).

[22]This purely monetary private cost measure does not include any psychic or other non-monetary costs of college education, which may be substantial and help explain why some individuals do not enroll despite large financial returns (Heckman et al., 2006).

[23]I do not include room and board charges in the cost-benefit analysis due to the fact that students must consume housing and food regardless of their enrollment status, and any increase in such costs generated by cutoff-crossing may primarily represent increased consumption during college rather than educational investment per se, e.g. living in a more expensive on-campus dorm and paying for a more expensive on-campus meal plan relative to off-campus alternatives. Similarly, I do not include student loans in the cost-benefit analysis, given that mandatory net tuition

Figure 11: Cost-Benefit Calculations
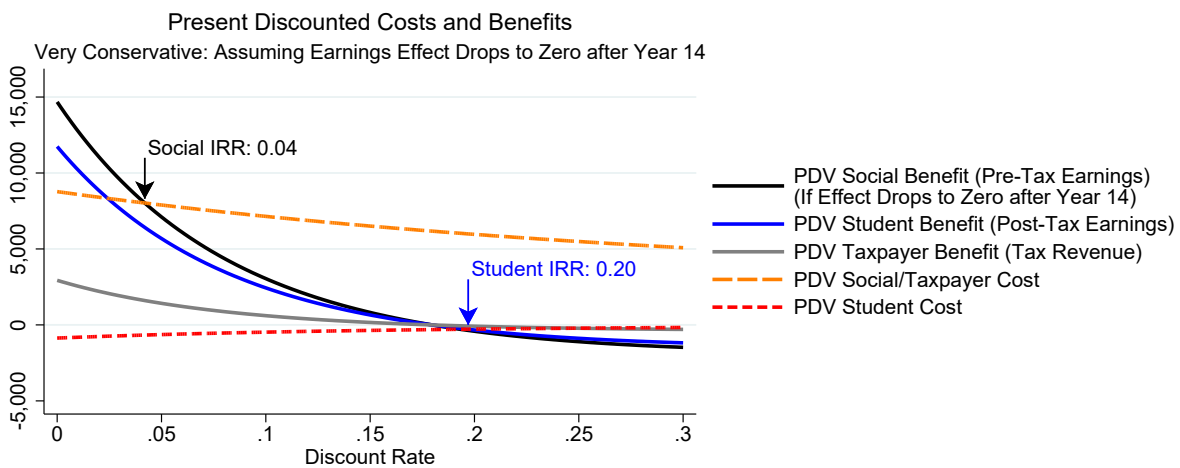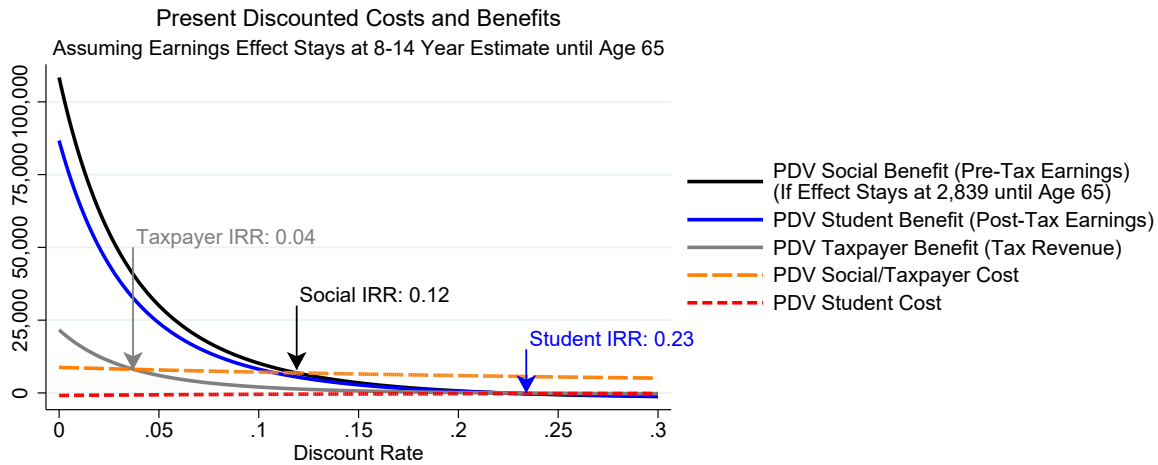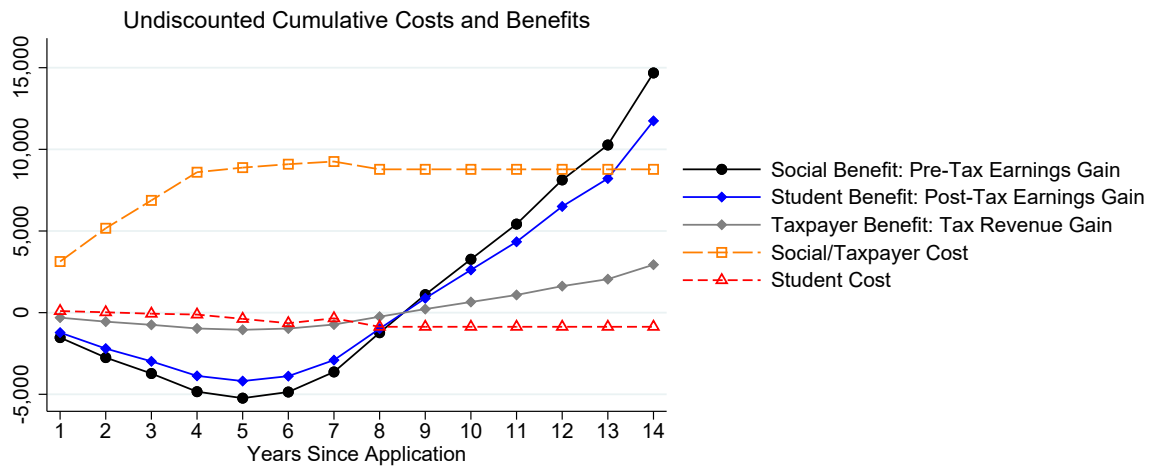
Notes: The top panel plots the accumulation of the annual Texas earnings effects in Figure 10 (see Figure A.12 for a version that uses the annualized positive earnings effects instead). Those (pre-tax) social benefits are then decomposed into the student benefit (post-tax earnings gain) and the taxpayer benefit (tax revenue gain), assuming a 20 percent tax and transfer rate. Those benefits are then compared to the cumulative student cost (net tuition) and social/taxpayer cost (educational resources). The middle and bottom panels show how the present discounted values of benefits and costs vary with the discount rate. The middle panel assumes that the annual earnings effect will remain at the pooled years 8-14 estimate of $2,839 (from the top right panel of Figure 10) from year 15 until retirement at age 65. The bottom panel assumes that the unobserved earnings effects beyond year 14 drop to zero, allowing only for the earnings benefits that are observed up until year 14.

which is the increase in educational resource costs inclusive of the large increase in the four-year sector and the small decrease in the two-year sector documented in the bottom panels of Figure 9. Informed by the dynamics in Figure 9, I assume that these costs stabilize at eight years after application.

Together, the dynamics of the cumulative costs and benefits in the top panel of Figure 11 show that the benefits of enrolling marginal students eventually surpass the costs, but at different horizons for students, society, and taxpayers. For students, the initial earnings foregone by increased college enrollment keep their benefits below their (zero) costs up until 8 years out from application. After 8 years, the net gain is positive and steadily grows at with a roughly linear trajectory. Society must wait even longer for the benefits to outweigh the costs, which occurs between years 12 and 13. For taxpayers, the gross revenue gains are increasing but still below costs as of year 14, so some extrapolation is required to estimate the eventual crossing point; Appendix Figure A.11 shows this point to be at 25 years out from application under the assumption about lifecycle earnings effects described in the next paragraph.

The middle and bottom panels of Figure 11 add discounting to the analysis to calculate net present values and internal rates of return. Instead of imposing one discount rate for the calculations, I plot a wide range of discount rates on the horizontal axis and show how the present discounted values of benefits and costs vary with the discount rate. In the middle panel, I assume that the annual pre-tax earnings effect will remain at the pooled years 8-14 estimate of $2,839 (from the top right panel of Figure 10) from year 15, when the typical student is age 33, until retirement at age 65, at which point the earnings effect drops to zero. Under this assumption, the undiscounted gross present value of the pre-tax earnings benefit is roughly $108,000, which then declines quickly as the discount rate increases. The present values of the student and taxpayer benefits are proportional to the social benefit by the assumption of a constant tax and transfer rate of 20 percent. Since the student cost is essentially zero, it remains flat at zero regardless of the discount rate; the social cost declines modestly given that all of the costs are incurred in the first few years of the timeline. The vertical distances between the costs and benefits show the net present value of enrolling the average marginal student given a particular discount rate, and their intersections yield the relevant internal rates of return: a substantial 23 percent for students, 12 percent for society, and 4 percent for taxpayers. As shown in Appendix Figure A.12, if I use the annualized positive earnings measure instead of the annual Texas earnings measure, the returns are similar but slightly lower at 19 percent for students, 10 percent for society, and 3 percent for taxpayers. The positive IRR for taxpayers implies that for discount rates no greater than 3-4 percent, the marginal value of public funds (MVPF) (e.g., Mayshar, 1990; Hendren, 2016) is infinite: the marginal enrollment pays for itself from the taxpayer's perspective thanks to increased future tax revenues (eventually)

---

charges do not increase; the additional loans accumulated in Figure 9 are likely used by marginal students primarily to finance additional room and board charges and other consumption during college at the expense of consumption after college, to the extent that the loan balances and financing charges are eventually paid back. To the extent that increased living expenses during college and the cost of financing them should be included as educational investment costs rather than intertemporal consumption smoothing, they would increase the cumulative costs plotted in the middle panel of Figure 11 and thus decrease the internal rates of return.

outweighing the upfront expenditures.

The bottom panel of Figure 11 makes the much more conservative assumption that the unobserved earnings effects beyond year 14 drop to zero, allowing only for the earnings benefits that are observed up until year 14. Even under this extremely conservative assumption, the student's internal rate of return is still 20 percent, but society's return drops to 4 percent, and the marginal enrollment no longer pays for itself from the taxpayer's perspective. In reality, even the middle panel's assumption of a constant dollar effect from year 14 through retirement may be conservative, as the earnings gains may grow proportionally with increasing lifecycle earnings levels (Neal, 2018).

5 Disentangling the Intensive and Extensive Margins

Figure 5 showed that 48 percent of cutoff compliers would initially fall back to another four-year institution if barely rejected from the target university, while the other 52 percent would initially fall out of the four-year sector, primarily to a two-year community college. Since these distinct complier types on the "intensive" versus "extensive" margins of the four-year sector may experience substantially different effects of enrolling in the target university, in this section I develop methods to disentangle their distinct contributions to the "average marginal" effects estimated above. How much of those effects are driven by the intensive margin of enrolling in a more-preferred four-year institution over a less-preferred one, versus the extensive margin of enrolling in any four-year institution?

Separate treatment effects for these two margins are not readily recovered by the fuzzy RD design laid out in Section 2.4 thanks to a fundamental underidentification problem: there are two treatment effects of interest, but only one available instrument (cutoff-crossing). Intensive and extensive margin compliers are not directly distinguishable in the data, since the counterfactual fallback option for each marginal applicant is nowhere recorded. However, the fact that I observe an applicant's entire portfolio of admissions to any of the 35 Texas public universities enables a powerful stratification of marginal applicants into two observable groups: those who have at least one admission offer from *another* Texas public university, and those who have none. This stratification does not perfectly separate intensive and extensive margin compliers, but it comes fairly close, and far more so than other observable stratifiers like the identity of the target school or an applicant's pre-college covariates: the complier-describing logic of Abadie (2002) shows that 72 percent of cutoff compliers with at least one other admission are on the intensive margin of the four-year sector (i.e., would fall back to that other available four-year option), while only 9 percent of cutoff compliers with no other admissions are on the intensive margin, i.e. fully 91 percent of them are on the extensive margin and would initially fall out of the four-year sector if rejected.

Thus, a seemingly straightforward strategy to identify separate treatment effects for (mostly) intensive margin compliers versus (almost entirely) extensive margin compliers would simply divide the sample into applicants with and without another four-year admission offer and run the main fuzzy RD specification in each. Unfortunately, there is one wrinkle: the other-admission stratifier

itself is somewhat endogenous to cutoff crossing. As shown in Appendix Figure A.13, crossing the admission cutoff causes the reduced form share of applicants with another admission to drop by a precisely estimated 4 percentage points. This phenomenon is likely caused by the availability of rolling admissions and spring admission cycles at many Texas public universities; marginal applicants who are barely rejected at their target university often still have time to secure an admission to a fallback school within the same application year. Since I do not observe any dates associated with applications or admissions within a given year, I cannot restrict the stratification to other admissions that were secured prior to the admission decision at the target university. Such a stratification would also likely reduce the power of the stratification itself, as some of the applicants with no other admission at the time of the rejection would later secure one.

The question to answer in this section, then, is what can we learn about intensive versus extensive margin treatment effects with a strong but endogenous stratifier? The answer, it turns out, is still a great deal. To get there, I begin with some useful notation. Let $A \in \{0, 1\}$ denote the observable stratifier of whether a given applicant to a given target university has any other admission offers from Texas public universities. Similar to the potential treatments D_1 and D_0 introduced in Section 2.4, let A_1 and A_0 indicate whether the applicant would have another admission offer if her test score running variable fell just above ($Z = 1$) or just below ($Z = 0$) the target university's admission cutoff, respectively.[24] Finally, as before, the applicant's potential outcome is Y_1 if she enrolls at the target university and Y_0 if she does not. With this notation in hand, the decomposition of interest is

$$E[Y_1 - Y_0 | D_0 = 0, D_1 = 1]$$
$$= \omega E[Y_1 - Y_0 | D_0 = 0, D_1 = 1, A_0 = 1] \tag{2}$$
$$+ (1 - \omega) E[Y_1 - Y_0 | D_0 = 0, D_1 = 1, A_0 = 0],$$

where the top treatment effect is the pooled complier LATE studied up until this section; the middle treatment effect is the LATE for the (mostly) intensive margin compliers who would have another four-year admission offer if barely below the target university's cutoff ($A_0 = 1$); the bottom treatment effect is the LATE for the (almost entirely) extensive margin compliers who would have no other four-year admission offer if barely below the target university's cutoff ($A_0 = 0$); and $\omega \equiv \frac{Pr[D_0=0,D_1=1,A_0=1]}{Pr[D_0=0,D_1=1]}$ is the share of cutoff compliers who would have another four-year admission if barely below the cutoff.

The identification challenge is that A_0 is a latent type, unobservable for marginal applicants who fall just above the cutoff, and thus I cannot condition on it directly to identify the separate treatment effects in (2). To see this clearly, expand the stratification of $A_0 = 1$ versus $A_0 = 0$ into

---

mutually exclusive and exhaustive latent response types with respect to A:

$A_0 = 0, A_1 = 0$: A-never-taker (never has another admission)

$A_0 = 1, A_1 = 1$: A-always-taker (always has another admission) $\qquad$ (3)

$A_0 = 1, A_1 = 0$: A-complier (secures another admission if and only if below cutoff)

$A_0 = 0, A_1 = 1$: A-defier (secures another admission if and only if above cutoff)

As indicated by the gray text, I assume away the presence of A-defiers; it would be somewhat odd behavior for a marginal applicant to scramble for another admission as a result of getting *admitted* to the target university.

Assumption 1 (No A-defiers): $Pr(A_0 = 0, A_1 = 1) = 0.$

I can then rewrite the decomposition in (2) as

$$
\begin{aligned}
& E[Y_1 - Y_0 | D\text{-complier}] \\
&= \omega E[Y_1 - Y_0 | D\text{-complier and } A\text{-never-taker}] \\
&+ (1 - \omega) E[Y_1 - Y_0 | D\text{-complier and } (A\text{-complier or } A\text{-always-taker})],
\end{aligned} \qquad (4)
$$

where "D-complier" refers to the pooled cutoff compliers studied up until this point, defined by the conditioning set $\{D_0 = 0, D_1 = 1\}$, and the three response types with respect to A refer to those defined in (3). The identification challenge arises entirely due to the presence of A-compliers, as they are the unobservable response type whose behavior renders the other admission stratifier A endogenous to crossing the admission cutoff at the target university.

A second useful assumption for the identification results that follow is that A-compliers are a subset of D-compliers:

Assumption 2 (All A-compliers are D-compliers): $Pr(D_0 = 0, D_1 = 1 | A_0 = 1, A_1 = 0) = 1.$

If a marginal applicant is a D-always-taker or a D-never-taker, then crossing the admission cutoff at the target university has no effect on their enrollment choice. Assumption 2 says that those students also experience no effect on their other admissions A, which is natural; if cutoff-crossing at the target university does not cause a student to change their enrollment at the target university, then it likely does not cause them to change their portfolio of other admissions either.

Under Assumptions 1 and 2 (in addition to the fuzzy RD assumptions discussed in Section 2.4), several, but not all, of the ingredients that go into (4) are separately identified by adapting the complier-describing logic of Abadie (2002) to the A response types. This adaptation amounts to a series of fuzzy RD estimands that replace the outcome Y with interactions among Y, D, and A, which identify several mean potential outcome levels of the A response types. As I show in

Appendix C,

$$\frac{E[YDA|Z=1] - E[YDA|Z=0]}{E[DA|Z=1] - E[DA|Z=0]} = E[Y_1|D\text{-complier and }A\text{-always-taker}],$$

where the estimand on the left side amounts to a fuzzy RD regression of the "outcome" YDA on the "treatment" DA, instrumenting for DA with cutoff-crossing Z. I repeat this across the other combinations of D and A to yield

$$\frac{E[Y(1-D)(1-A)|Z=1] - E[Y(1-D)(1-A)|Z=0]}{E[(1-D)(1-A)|Z=1] - E[(1-D)(1-A)|Z=0]} = E[Y_0|D\text{-complier and }A\text{-never-taker}]$$

$$\frac{E[YD(1-A)|Z=1] - E[YD(1-A)|Z=0]}{E[D(1-A)|Z=1] - E[(D(1-A)|Z=0]} = E[Y_1|D\text{-complier and }(A\text{-never-taker or }A\text{-complier})]$$

$$\frac{E[Y(1-D)A|Z=1] - E[Y(1-D)A|Z=0]}{E[(1-D)A|Z=1] - E[(1-D)A|Z=0]} = E[Y_0|D\text{-complier and }(A\text{-always-taker or }A\text{-complier})].$$

None of these identified mean potential outcome levels can be immediately differenced to form a well-defined treatment effect, since the conditioning set of A-response types differs across each. But the patchwork they provide will reveal some clear strategies for learning about the separate treatment effects of interest. To see this, the top panel of Figure 12 first plots the estimates of these four mean potential outcomes for BA completion using the main fuzzy RD specification from Section 2.4. The known quantities in black solid dots include the untreated (Y_0) mean potential outcomes among both subgroups of interest: the A-never-takers, who would have no other admissions if rejected from the target school ($A_0 = 0$) and are 91 percent on the extensive margin, and the combined A-compliers and A-always-takers, who would have at least one other admission if rejected from the target school ($A_0 = 1$) and are 72 percent on the intensive margin. The shares of these response types are also identified, by the denominators of the modified fuzzy RD estimands above: 38 percent of D-compliers are A-never-takers ($A_0 = 0$), and the other 62 percent are A-compliers (27 percent) or A-always-takers (35 percent), both of whom have $A_0 = 1$.

The identified quantities also include two treated (Y_1) mean potential outcomes (in gray solid dots), but they are for subgroups that are not exactly the subgroups of interest: a composite of A-never-takers and A-compliers, and isolated A-always-takers. The hollow red circles with question marks, in contrast, represent the unknown quantities of interest: the treated (Y_1) potential outcomes among A-never-takers by themselves, and the composite of A-compliers and A-always-takers. If we knew the location of these red circles, we could immediately form a mean treatment effect $Y_1 - Y_0$ among the (almost entirely) extensive margin compliers with $A_0 = 0$ (no other admissions if rejected from the target university) and a separate mean treatment effect for the (mostly) intensive margin compliers with $A_0 = 1$ (at least one other admission if rejected from the target university).

Without further assumptions, the two unknown mean Y_1's in the top panel of Figure 12 are not

Figure 12: Disentangling the Intensive and Extensive Margins: BA Completion
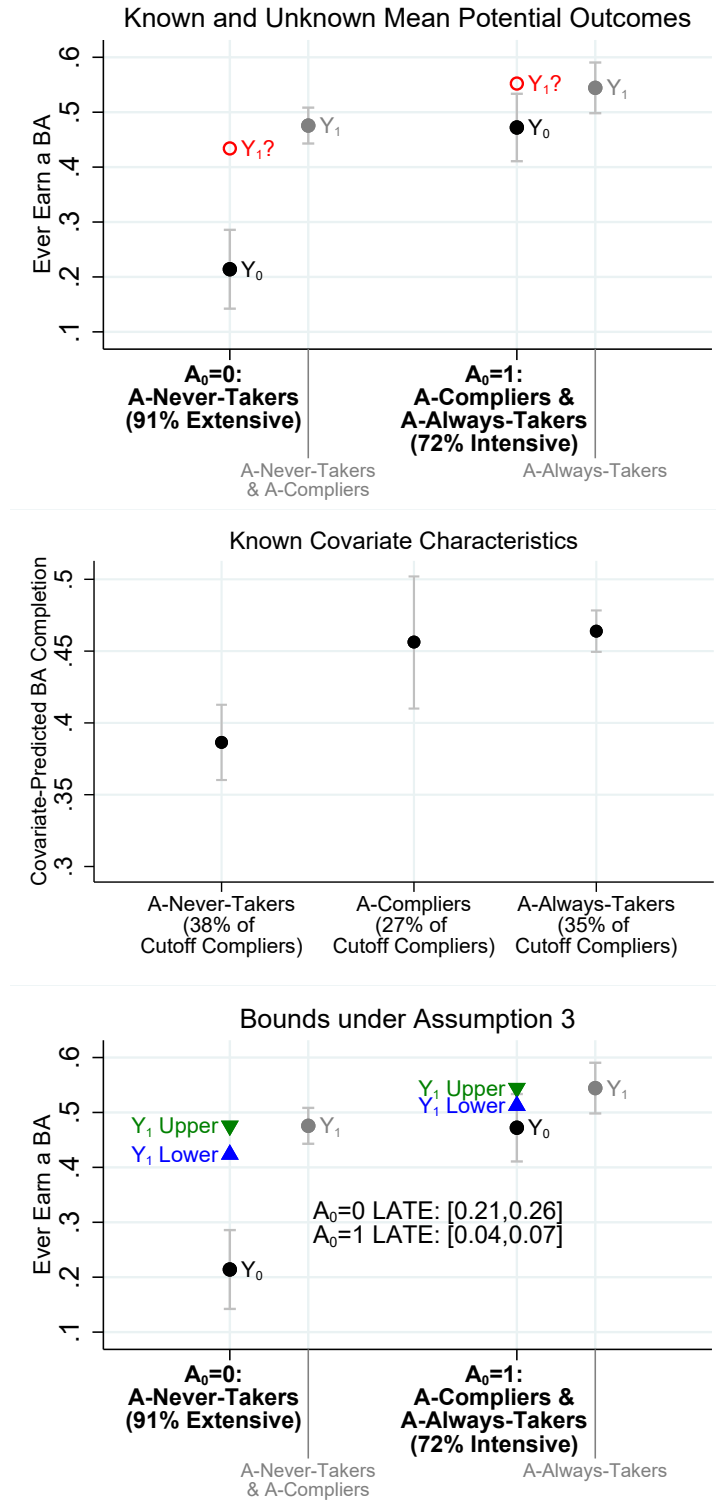
Notes: The top panel plots the estimates of the separate mean potential outcomes identified in Section 5, estimated using the main RD specification described in Section 2.4. The middle panel plots the covariate characteristics (summarized as covariate-predicted BA completion) of the three response types with respect to other admissions A, which are identified by replacing Y with X in the identification arguments of Section 5. The bottom panel shows the upper and lower bounds on the mean Y_1's of interest implied by Assumption 3.

point-identified, but we can use the values of the two known Y_1's to try to learn about them. We could first consider agnostic bounds in the spirit of Horowitz and Manski (2000) and Lee (2009), but knowing that the mean Y_1 for the composite of A-never-takers and A-compliers is roughly 0.5 puts little restriction on the possible mean Y_1 for A-never-takers alone, given the relative masses of the two groups.

I propose a stronger but still somewhat agnostic approach that uses the relative rankings of the identified mean potential outcomes, along with the relative rankings of the identified mean characteristics of these response types, to inform a rank assumption involving the unknown mean potential outcomes.[25] First, notice in the top panel of Figure 12 that the identified untreated mean Y_0's increase when moving from left to right on the plot across the response types, i.e. from A-never-takers to the composite of A-compliers and A-always-takers. Second, notice a similar increasing pattern in the identified treated mean Y_1's when moving from the composite of A-never-takers and A-compliers to the isolated A-always-takers. Third, the middle panel of Figure 12 describes the mean pre-college covariates (summarized as covariate-predicted BA completion) of the three response types, which are separately identified by replacing the outcome Y with the covariate X (in this case, covariate-predicted BA completion) in the identification arguments above. Importantly, A-never-takers have significantly lower covariate-predicted BA completion compared to A-compliers and A-always takers, who themselves are indistinguishable from each other.

Informed by these rankings, consider the assumption that the mean potential outcome of enrolling in the target university (Y_1) weakly increases across the response types. Letting "DC" denote shorthand for D-complier:

Assumption 3 (Weakly increasing types):

$E[Y_1|DC$ and A-never-taker$] \leq E[Y_1|DC$ and A-complier$] \leq E[Y_1|DC$ and A-always-taker$]$.

Assumption 3 is both intuitive and empirically grounded. If all three A response types were to enroll in the target university, it is natural to assume that A-always-takers, who always have at least one other admission, would tend to have weakly better outcomes than A-compliers, who scramble for another admission if rejected from the target university, who in turn would tend to have weakly better outcomes than A-never-takers, who never have any other admissions. The identified covariate characteristics of these three groups support such a ranking, as shown in the middle panel of Figure 12, as well as the rankings within the identified Y_1's and within the identified Y_0's in the top panel of Figure 12.

The bottom panel of Figure 12 visualizes the implications of Assumption 3. It immediately identifies an upper bound on the mean Y_1 for A-never-takers (green triangle), which is simply the originally identified mean Y_1 for the composite of A-never-takers and A-compliers. Likewise, Assumption 3 implies that the originally identified mean Y_1 for A-always-takers is an upper bound

---

[25]See Lee and Salanié (2023) and Galperin (2023) for different but related approaches applied to Head Start enrollment with multiple fallback options and grant aid with selected enrollment data, respectively.

for the mean Y_1 of the composite of A-compliers and A-always-takers. Each of these upper bounds, in turn, imply lower bounds for the other stratum (blue triangles), since the mean Y_1's across the two strata ($A_0 = 0$ and $A_0 = 1$) must average up to the known mean Y_1 among all cutoff compliers.

These bounds end up being quite informative about the separate contributions of intensive and extensive margin compliers. The implied LATE among the (almost entirely) extensive margin compliers with $A_0 = 0$ is substantially larger than the small LATE among the (mostly) intensive margin compliers with $A_0 = 1$. In fact, the small LATE for the $A_0 = 1$ stratum likely overstates the true intensive margin effect, due to the fact that the A_0 stratification does not perfectly isolate intensive margin compliers; recall that 72% of D-compliers with $A_0 = 1$ are on the intensive margin, meaning 28% of them are actually on the extensive margin, i.e. they have another four-year admission but do not excercise the option, falling out of the four-year sector instead. If these extensive margin applicants with $A_0 = 1$ reap large treatment effects akin to the 91% extensive margin compliers with $A_0 = 0$, then the truly intensive margin treatment effects would be even smaller than the tightly bounded LATE for the $A_0 = 1$ stratum. Thus, the vast majority, and perhaps nearly the entirety, of the pooled complier effects on BA completion estimated in Section 3 are driven by extensive margin compliers who would not enroll in any four-year college if rejected, with a much smaller contribution, if any, from the intensive margin compliers who would fall back to a less-preferred four-year option.

Figure 13 conducts the same exercise for inferring separate intensive and extensive margin effects on college peer quality, cumulative four-year credits, and earnings. The top two panels show relatively larger increases in college peer quality (as measured by high school math scores and future earnings) along the (almost entirely) extensive margin, but still substantively large increases along the (mostly) intensive margin as well. Importantly, then, the treatment-control contrast along the intensive margin still involves a substantial increase in traditional measures of 4-year college selectivity. The middle left panel shows dramatically larger extensive versus intensive margin effects on cumulative credits in the four-year sector. The final three panels decompose effects on the three different earnings measures. The bounds for earnings rank (bottom right) are very tight, while the bounds for annual Texas earnings (middle right) and annualized positive earnings (bottom left) very slightly cross, due to the fact that for both outcomes the estimated mean Y_1 for A-never-takers and A-compliers is actually slightly below the estimated mean Y_1 for A-always-takers, though not substantively so. All three earnings measures show substantially larger effects for extensive versus intensive margin compliers, implying a dominant role for the extensive margin in driving the returns to enrolling "average marginal" public university applicants. This conclusion is further strengthened when remembering that 28% of the $A_0 = 1$ stratum is actually on the extensive margin; as noted above, if these extensive margin applicants with $A_0 = 1$ reap similar treatment effects as the $A_0 = 0$ stratum, who are 91% extensive margin compliers, then the truly intensive margin treatment effects are even smaller.

36

Figure 13: Disentangling the Intensive and Extensive Margins: Peer Quality, Credits, and Earnings
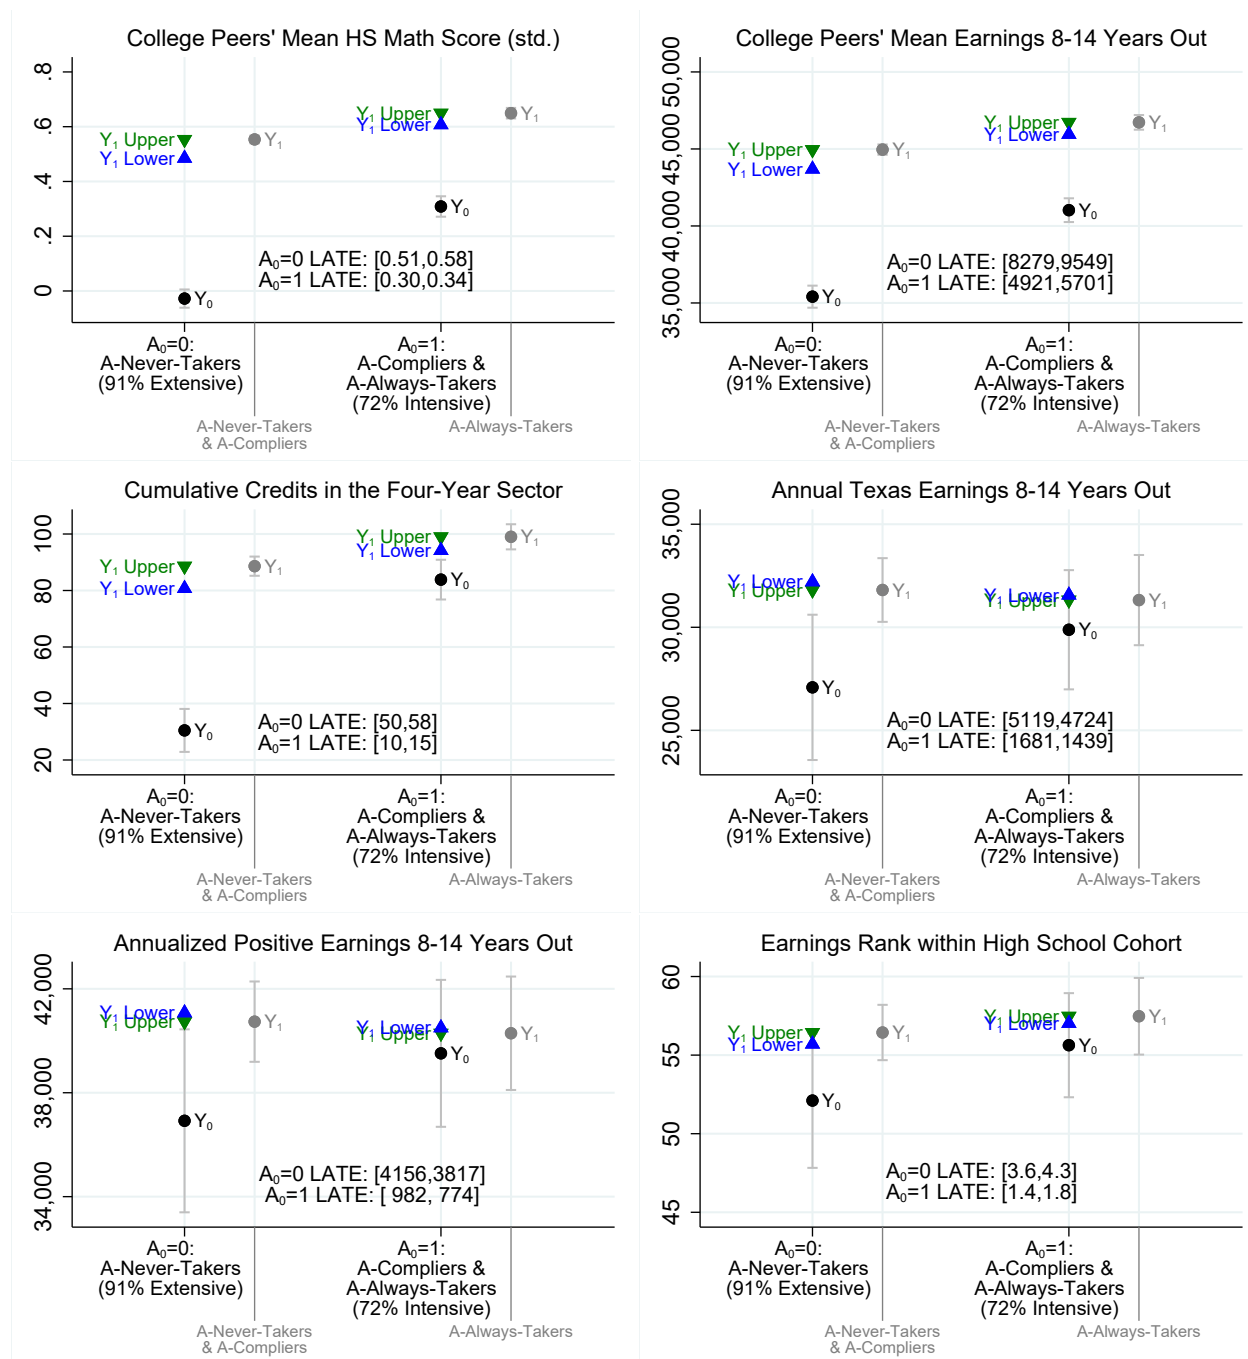
Notes: This figure plots the estimates of the separate mean potential outcomes identified in Section 5, estimated using the main RD specification described in Section 2.4, along with the upper and lower bounds on the mean Y_1's of interest implied by Assumption 3.

6　Conclusion

In this paper, I studied the causal impacts of public universities on the outcomes of their marginally admitted students. I used administrative admission records spanning all 35 public universities in Texas, which collectively enroll over 10 percent of American public university students, to systematically identify and employ decentralized cutoffs in SAT and ACT scores that generate discontinuities in admission and enrollment. I linked marginal applicants around these cutoffs back to their pre-college characteristics and forward to their educational and labor market outcome trajectories in administrative data. Together, these data linkages and discontinuities enabled a fuzzy regression discontinuity research design that transparently attributes jumps in applicant outcomes around the cutoffs to jumps in admission and enrollment, justified by smooth densities of applicants and their baseline characteristics through the cutoffs.

Marginally admitted students complete an additional year of four-year education, are 12 percentage points more likely to earn a bachelor's degree, and eventually earn 5-10 percent more than their marginally rejected but otherwise identical counterparts. Marginally admitted students pay none of the additional monetary costs of these investments thanks to offsetting grant aid; I conducted formal cost-benefit calculations that show internal rates of return of 19-23 percent for the marginal students themselves, 10-12 percent for society, and 3-4 percent for the government budget. Finally, I developed a method to derive tight bounds on separate effects for students on the extensive margin of the four-year sector versus those who would fall back to a less-preferred four-year school if rejected. These results revealed a dominant role for the extensive margin in driving the average marginal returns.

References

ABADIE, A. (2002): "Bootstrap Tests for Distributional Treatment Effects in Instrumental Variable Models," *Journal of the American Statistical Association*, 97, 284–292.

ALTMEJD, A., A. BARRIOS-FERNÁNDEZ, M. DRLJE, J. GOODMAN, M. HURWITZ, D. KOVAC, C. MULHERN, C. NEILSON, AND J. SMITH (2021): "O Brother, Where Start Thou? Sibling Spillovers on College and Major Choice in Four Countries," *The Quarterly Journal of Economics*, 136, 1831–1886.

ANDREWS, R. J., S. A. IMBERMAN, AND M. F. LOVENHEIM (2017): "Risky Business? The Effect of Majoring in Business on Earnings and Educational Attainment," NBER Working Paper No. 23575.

ANELLI, M. (2020): "The Returns to Elite University Education: a Quasi-Experimental Analysis," *Journal of the European Economic Association*, 18, 2824–2868.

ANGRIST, J., D. AUTOR, AND A. PALLAIS (2022): "Marginal Effects of Merit Aid for Low-Income Students," *The Quarterly Journal of Economics*, 137, 1039–1090.

ANGRIST, J. D. AND A. B. KRUEGER (1991): "Does Compulsory School Attendance Affect Schooling and Earnings?" *The Quarterly Journal of Economics*, 106, 979–1014.

ARYAL, G., M. BHULLER, AND F. LANGE (2022): "Signaling and Employer Learning with Instruments," *American Economic Review*, 112, 1669–1702.

AUCEJO, E. M., J. FRENCH, AND B. ZAFAR (2023): "Estimating Students' Valuation for College Experiences," *Journal of Public Economics*, 224, 104926.

BARROW, L. AND O. MALAMUD (2015): "Is College a Worthwhile Investment?" *Annual Review of Economics*, 7, 519–555.

BECKER, G. S. (1964): *Human Capital: A Theoretical and Empirical Analysis, with Special Reference to Education*, New York: National Bureau of Economic Research.

BELKIN, D. (2023): "Americans Are Losing Faith in College Education, WSJ-NORC Poll Finds," *The Wall Street Journal*.

BHULLER, M., M. MOGSTAD, AND K. G. SALVANES (2017): "Life-Cycle Earnings, Education Premiums, and Internal Rates of Return," *Journal of Labor Economics*, 35, 993–1030.

BLACK, S. E., J. T. DENNING, AND J. ROTHSTEIN (2023): "Winners and Losers? The Effect of Gaining and Losing Access to Selective Colleges on Education and Labor Market Outcomes," *American Economic Journal: Applied Economics*, 15, 26–67.

BLEEMER, Z. (2021): "Top Percent Policies and the Return to Postsecondary Selectivity," Working paper.

——— (2022): "Affirmative Action, Mismatch, and Economic Mobility after California's Proposition 209," *The Quarterly Journal of Economics*, 137, 115–160.

BRAND, J., F. PFEFFER, AND S. GOLDRICK-RAB (2014): "The Community College Effect Revisited: The Importance of Attending to Heterogeneity and Complex Counterfactuals," *Sociological Science*, 1, 448–465.

BREWER, D. J. AND R. G. EHRENBERG (1996): "Does It Pay to Attend an Elite Private College? Evidence from the Senior High School Class of 1980," *Research in Labor Economics*, 15, 239–271.

BRUNNER, E. J., S. M. DOUGHERTY, AND S. L. ROSS (2023): "The Effects of Career and Technical Education: Evidence from the Connecticut Technical High School System," *The Review of Economics and Statistics*, 1–16.

CALONICO, S., M. D. CATTANEO, AND R. TITIUNIK (2014): "Robust Nonparametric Confidence Intervals for Regression-Discontinuity Designs," *Econometrica*, 82, 2295–2326.

CANAAN, S. AND P. MOUGANIE (2018): "Returns to Education Quality for Low-Skilled Students: Evidence from a Discontinuity," *Journal of Labor Economics*, 36, 395–436.

CARD, D. (2001): "Estimating the Return to Schooling: Progress on Some Persistent Econometric Problems," *Econometrica*.

CARNEIRO, P., E. GALASSO, AND R. GINJA (2019): "Tackling Social Exclusion: Evidence from Chile," *The Economic Journal*, 129, 172–208.

CARNEIRO, P., J. J. HECKMAN, AND E. J. VYTLACIL (2011): "Estimating Marginal Returns to Education," *American Economic Review*, 101, 2754–2781.

CATTANEO, M. D., L. KEELE, R. TITIUNIK, AND G. VAZQUEZ-BARE (2016): "Interpreting Regression Discontinuity Designs with Multiple Cutoffs," *The Journal of Politics*, 78, 1229–1248.

CHETTY, R., D. J. DEMING, AND J. N. FRIEDMAN (2023): "Diversifying Society's Leaders? The Determinants and Causal Effects of Admission to Highly Selective Private Colleges," NBER Working Paper No. 31492.

CHETTY, R., J. N. FRIEDMAN, E. SAEZ, N. TURNER, AND D. YAGAN (2020): "Income Segregation and Intergenerational Mobility Across Colleges in the United States," *The Quarterly Journal of Economics*, 135, 1567–1633.

CLARK, D. (2023): "School Quality and the Return to Schooling in Britain: New Evidence from a Large-Scale Compulsory Schooling Reform," *Journal of Public Economics*, 223, 104902.

COHODES, S. R. AND J. S. GOODMAN (2014): "Merit Aid, College Quality, and College Completion: Massachusetts' Adams Scholarship as an In-Kind Subsidy," *American Economic Journal: Applied Economics*, 6, 251–285.

COLLEGE BOARD (2009): "ACT and SAT Concordance Tables," Tech. rep.

CUNHA, J. M. AND T. MILLER (2014): "Measuring Value-Added in Higher Education: Possibilities and Limitations in the Use of Administrative Data," *Economics of Education Review*, 42, 64–77.

DALE, S. B. AND A. B. KRUEGER (2002): "Estimating the Payoff to Attending a More Selective College: An Application of Selection on Observables and Unobservables," *The Quarterly Journal of Economics*, 117, 1491–1527.

DAUGHERTY, L., P. MARTORELL, AND I. MCFARLIN (2014): "Percent Plans, Automatic Admissions, and College Outcomes," *IZA Journal of Labor Economics*, 3, 10.

DILLON, E. W. AND J. A. SMITH (2020): "The Consequences of Academic Match between Students and Colleges," *Journal of Human Resources*, 55, 767–808.

DONG, Y. (2018): "Alternative Assumptions to Identify LATE in Fuzzy Regression Discontinuity Designs," *Oxford Bulletin of Economics and Statistics*, 80, 1020–1027.

DURYEA, S., R. P. RIBAS, B. SAMPAIO, G. R. SAMPAIO, AND G. TREVISAN (2023): "Who Benefits from Tuition-Free, Top-Quality Universities? Evidence from Brazil," *Economics of Education Review*, 95, 102423.

DYNARSKI, S., L. PAGE, AND J. SCOTT-CLAYTON (2023): "College Costs, Financial Aid, and Student Decisions," in *Handbook of the Economics of Education*, ed. by E. A. Hanushek, S. Machin, and L. Woessmann, Elsevier, vol. 7, chap. 4, 227–285.

FELLER, A., T. GRINDAL, L. MIRATRIX, AND L. C. PAGE (2016): "Compared to What? Variation in the Impacts of Early Childhood Education by Alternative Care Type," *The Annals of Applied Statistics*, 10, 1245–1285.

FOOTE, A. AND K. M. STANGE (2022): "Migration from Sub-National Administrative Data: Problems and Solutions with an Application to Higher Education," NBER Working Paper No. 30232.

GALPERIN, M. (2023): "Who Benefits from College Grant Aid and Why? Evidence from Texas," Working paper.

GANDIL, M. H. (2024): "Trickle Down Education: Ripple Effects in College Admissions," Working paper.

GE, S., E. ISAAC, AND A. MILLER (2022): "Elite Schools and Opting In: Effects of College Selectivity on Career and Family Outcomes," *Journal of Labor Economics*, 40, S383–S427.

GONG, Y., L. LOCHNER, T. STINEBRICKNER, AND R. STINEBRICKNER (2021): "The Consumption Value of College," Working paper.

GOODMAN, J., M. HURWITZ, AND J. SMITH (2017): "Access to Four-Year Public Colleges and Degree Completion," *Journal of Labor Economics*, 35, 829–867.

GRILICHES, Z. (1977): "Estimating the Returns to Schooling: Some Econometric Problems," *Econometrica*, 45, 1–22.

HAHN, J., P. TODD, AND W. VAN DER KLAAUW (2001): "Identification and Estimation of Treatment Effects with a Regression-Discontinuity Design," *Econometrica*, 69, 201–209.

HASTINGS, J. S., C. A. NEILSON, AND S. D. ZIMMERMAN (2014): "Are Some Degrees Worth More than Others? Evidence from College Admission Cutoffs in Chile," NBER Working Paper No. 19241.

HECKMAN, J. J., J. E. HUMPHRIES, AND G. VERAMENDI (2018): "Returns to Education: The Causal Effects of Education on Earnings, Health, and Smoking," *Journal of Political Economy*, 126, S197–S246.

HECKMAN, J. J., L. LOCHNER, AND C. TABER (1998): "Tax Policy and Human-Capital Formation," *The American Economic Review*, 88, 293–297.

HECKMAN, J. J., L. J. LOCHNER, AND P. E. TODD (2006): "Earnings Functions, Rates of Return and Treatment Effects: The Mincer Equation and Beyond," in *Handbook of the Economics of Education*, ed. by E. Hanushek and F. B. Welch, Elsevier, vol. 1, chap. 7, 307–458.

——— (2008): "Earnings Functions and Rates of Return," *Journal of Human Capital*, 2, 1–31.

HECKMAN, J. J. AND S. URZUA (2010): "Comparing IV with Structural Models: What Simple IV Can and Cannot Identify," *Journal of Econometrics*, 156, 27–37.

HENDREN, N. (2016): "The Policy Elasticity," *Tax Policy and the Economy*, 30, 51–89.

HENDREN, N. AND B. SPRUNG-KEYSER (2020): "A Unified Welfare Analysis of Government Policies," *The Quarterly Journal of Economics*, 135, 1209–1318.

HOEKSTRA, M. (2009): "The Effect of Attending the Flagship State University on Earnings: A Discontinuity-Based Approach," *Review of Economics and Statistics*, 91, 717–724.

HOROWITZ, J. L. AND C. F. MANSKI (2000): "Nonparametric Analysis of Randomized Experiments with Missing Covariate and Outcome Data," *Journal of the American Statistical Association*, 95, 77–84.

HOXBY, C. M. (2018): "Online Postsecondary Education and Labor Productivity," in *Education, Skills, and Technical Change: Implications for Future US GDP Growth*, ed. by C. R. Hulten and V. A. Ramey, Chicago: University of Chicago Press, chap. 11, 401–460.

——— (2019): "The Productivity of US Postsecondary Institutions," in *Productivity in Higher Education*, ed. by C. M. Hoxby and K. Stange, University of Chicago Press, 31–66.

HULL, P. (2018): "IsoLATEing: Identifying Counterfactual-Specific Treatment Effects with Cross-Stratum Comparisons," Working paper.

IMBENS, G. AND K. KALYANARAMAN (2012): "Optimal Bandwidth Choice for the Regression Discontinuity Estimator," *The Review of Economic Studies*, 79, 933–959.

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.

KAMAT, V. (2023): "Identification of Program Access Effects with an Application to Head Start," Working paper.

KIRKEBOEN, L. J., E. LEUVEN, AND M. MOGSTAD (2016): "Field of Study, Earnings, and Self-Selection," *The Quarterly Journal of Economics*, 131, 1057–1111.

KLINE, P. AND C. R. WALTERS (2016): "Evaluating Public Programs with Close Substitutes: The Case of Head Start," *The Quarterly Journal of Economics*, 131, 1795–1848.

KOLESÁR, M. AND C. ROTHE (2018): "Inference in Regression Discontinuity Designs with a Discrete Running Variable," *American Economic Review*, 108, 2277–2304.

KOZAKOWSKI, W. (2023): "Are Four-Year Public Colleges Engines for Economic Mobility? Evidence from Statewide Admissions Thresholds," *Annenberg Institute EdWorkingPaper No. 23-727*.

LEE, D. S. (2009): "Training, Wages, and Sample Selection: Estimating Sharp Bounds on Treatment Effects," *The Review of Economic Studies*, 76, 1071–1102.

LEE, S. AND B. SALANIÉ (2023): "Treatment Effects with Targeting Instruments," Working paper.

LONDOÑO-VÉLEZ, J., C. RODRIGUEZ, F. SANCHEZ, AND L. E. ÁLVAREZ-ARANGO (2023): "Financial Aid and Social Mobility: Evidence from Colombia's Ser Pilo Paga," NBER Working Paper No. 31737.

LOVENHEIM, M. AND J. SMITH (2023): "Returns to Different Postsecondary Investments: Institution Type, Academic Programs, and Credentials," in *Handbook of the Economics of Education*, ed. by E. A. Hanushek, S. Machin, and L. Woessmann, Elsevier, vol. 6, chap. 4, 187–318.

MAYSHAR, J. (1990): "On Measures of Excess Burden and Their Application," *Journal of Public Economics*, 43, 263–289.

MCCRARY, J. (2008): "Manipulation of the Running Variable in the Regression Discontinuity Design: A Density Test," *Journal of Econometrics*, 142, 698–714.

MILLER, D. W. (2007): "Isolating the Causal Impact of Community College Enrollment on Educational Attainment and Labor Market Outcomes in Texas," SIEPR Discussion Paper No. 06–33.

MILLER, L. (2023): "Switching Schools: Effects of College Transfers," Working paper.

MOUNTJOY, J. (2022): "Community Colleges and Upward Mobility," *American Economic Review*, 112, 2580–2630.

MOUNTJOY, J. AND B. R. HICKMAN (2021): "The Returns to College(s): Relative Value-Added and Match Effects in Higher Education," *NBER Working Paper No. 29276*.

NATIONAL CENTER FOR EDUCATION STATISTICS (2022): "Integrated Postsecondary Education Data System," https://nces.ed.gov/ipeds.

NEAL, D. A. (2018): *Information, Incentives, and Education Policy*, Harvard University Press.

NOYES, C. R. (1945): "Director's Comment," in *Income from Independent Professional Practice*, ed. by M. Friedman and S. Kuznets, New York: National Bureau of Economic Research, chap. 9, 405–410.

OREOPOULOS, P. AND K. G. SALVANES (2011): "Priceless: The Nonpecuniary Benefits of Schooling," *Journal of Economic Perspectives*, 25, 159–184.

OST, B., W. PAN, AND D. WEBBER (2018): "The Returns to College Persistence for Marginal Students: Regression Discontinuity Evidence from University Dismissal Policies," *Journal of Labor Economics*, 36, 779–805.

PORTER, J. AND P. YU (2015): "Regression Discontinuity Designs with Unknown Discontinuity Points: Testing and Estimation," *Journal of Econometrics*, 189, 132–147.

ROUSE, C. E. (1995): "Democratization or Diversion? The Effect of Community Colleges on Educational Attainment," *Journal of Business and Economic Statistics*, 13, 217–224.

SAAVEDRA, J. (2009): "The Learning and Early Labor Market Effects of College Quality: A Regression Discontinuity Analysis," Working paper.

SEKHRI, S. (2020): "Prestige Matters: Wage Premium and Value Addition in Elite Colleges," *American Economic Journal: Applied Economics*, 12, 207–225.

SMITH, J., J. GOODMAN, AND M. HURWITZ (2020): "The Economic Impact of Access to Public Four-Year Colleges," NBER Working Paper No. 27177.

STEVENS, D. W. (2007): "Employment That Is Not Covered by State Unemployment Insurance Laws," U.S. Census Bureau Technical Paper No. TP–2007–04.

U.S. DEPARTMENT OF HOMELAND SECURITY (2016): "DHS STEM Designated Degree Program List," https://www.ice.gov/sites/default/files/documents/stem-list.pdf.

UTD-ERC (2022): "UT Dallas Education Research Center Data Holdings," https://tsp.utdallas.edu.

WILLIS, R. J. (1986): "Wage Determinants: A Survey and Reinterpretation of Human Capital Earnings Functions," in *Handbook of Labor Economics*, ed. by O. C. Ashenfelter and R. Layard, New York: North-Holland, vol. 1, chap. 10, 525–602.

ZIMMERMAN, S. D. (2014): "The Returns to College Admission for Academically Marginal Students," *Journal of Labor Economics*, 32, 711–754.

——— (2019): "Elite Colleges and Upward Mobility to Top Jobs and Top Incomes," *American Economic Review*, 109, 1–47.

A   Appendix Figures and Tables

Figure A.1: Top Ten Percent Applicants Are Smooth Through the Cutoff
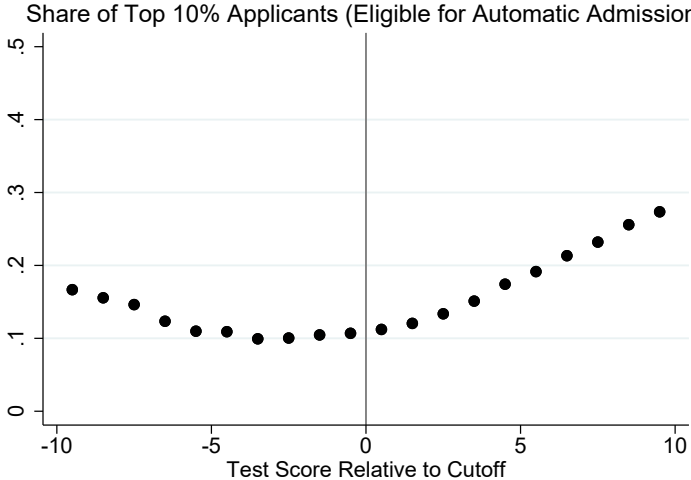
Notes: This figure plots the share of applicants at each of value of the concorded ACT running variable who are in the top ten percent of their high school graduating class, making them eligible for automatic admission at Texas public universities (outside of UT-Austin, which has a stricter threshold) via the Top Ten Percent law.

Figure A.2: Balance across Individual Pre-College Covariates
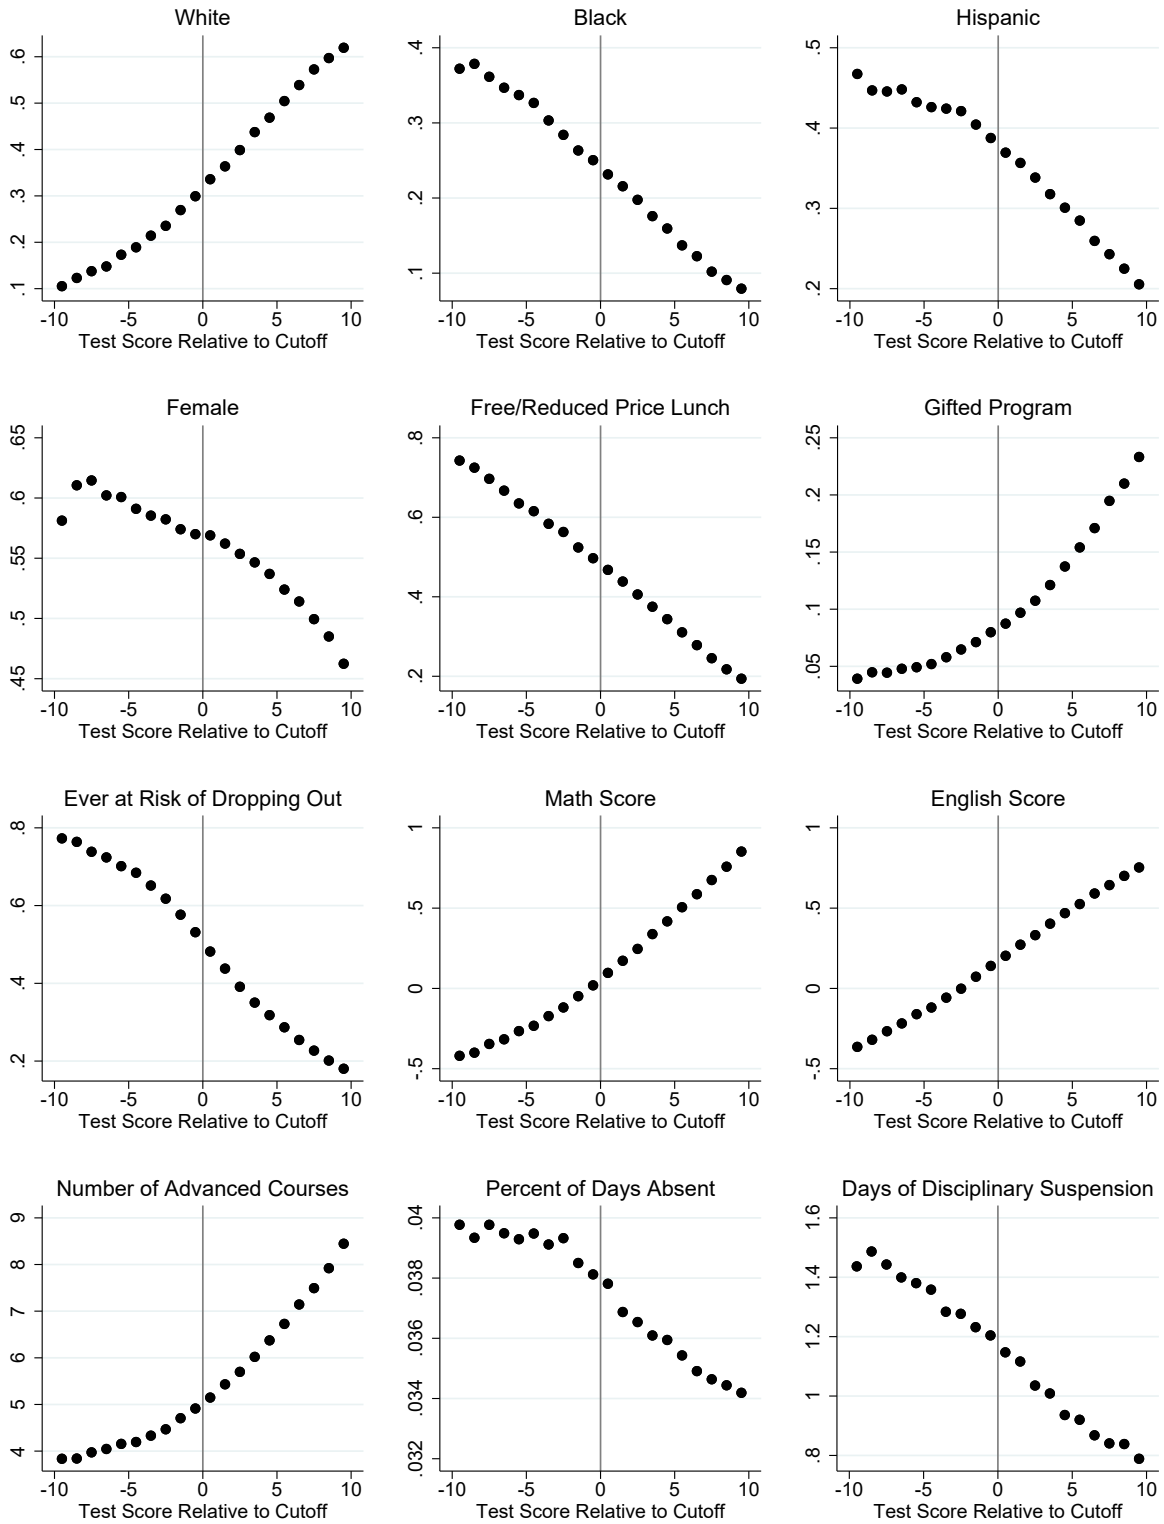
Notes: These plots show balance across the individual pre-college covariates used to predict BA completion and earnings in Figure 4. All covariates are measured during an applicant's observable high school years. Math and English scores are standardized to have mean zero and standard deviation one across all Texas public high school test takers.
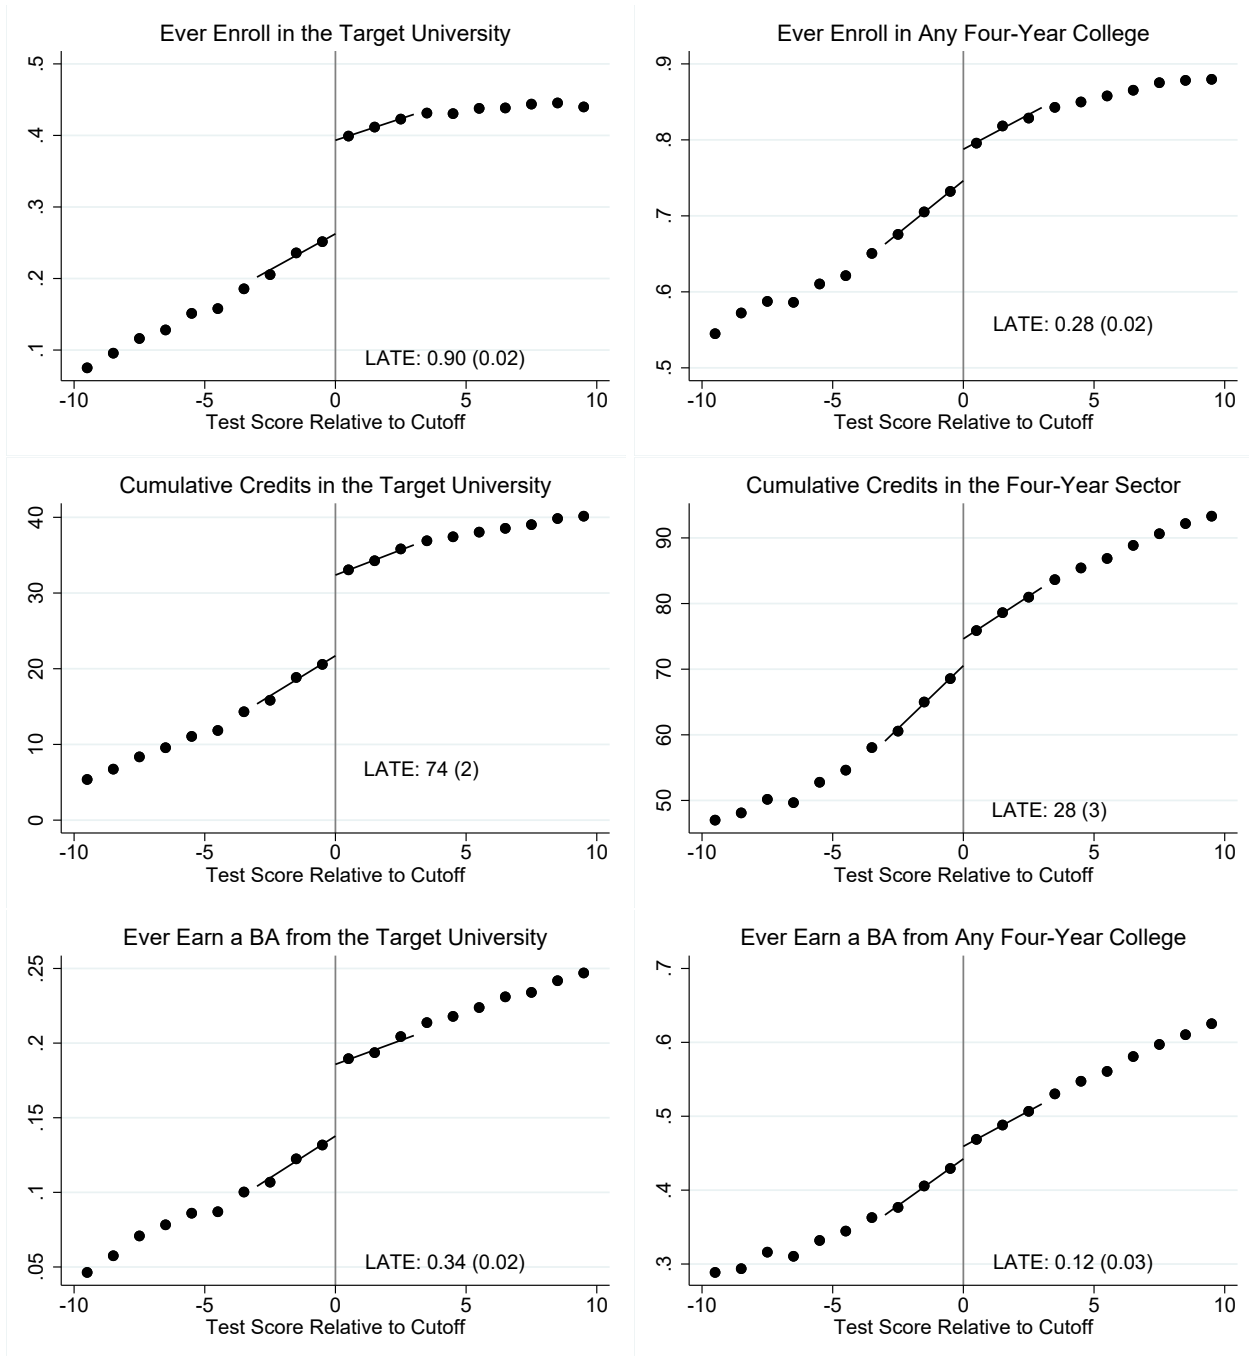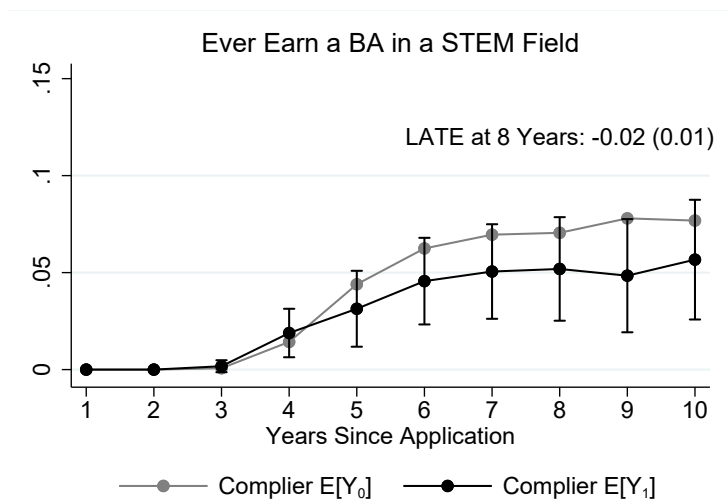
Figure A.3: Private and Out-of-State College Enrollment
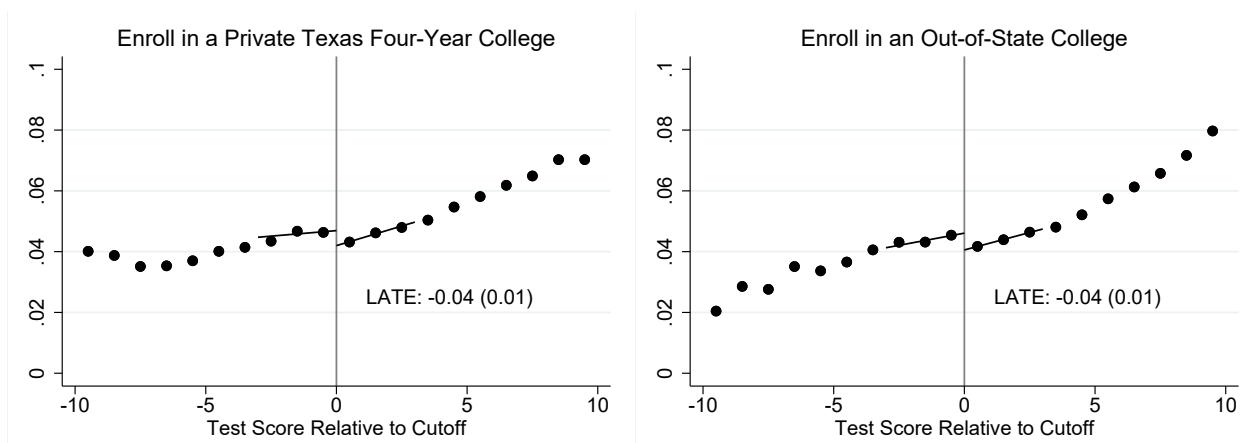
Notes: Enrolling in a private Texas four-year college is observed in the THECB data for all sample cohorts, and included in the outcome "Enroll in Another Texas Four-Year College" in Figure 5. Enrolling in an out-of-state college in the National Student Clearinghouse data is not observed for the 2004-2007 cohorts but observed for the 2008-2016 cohorts.

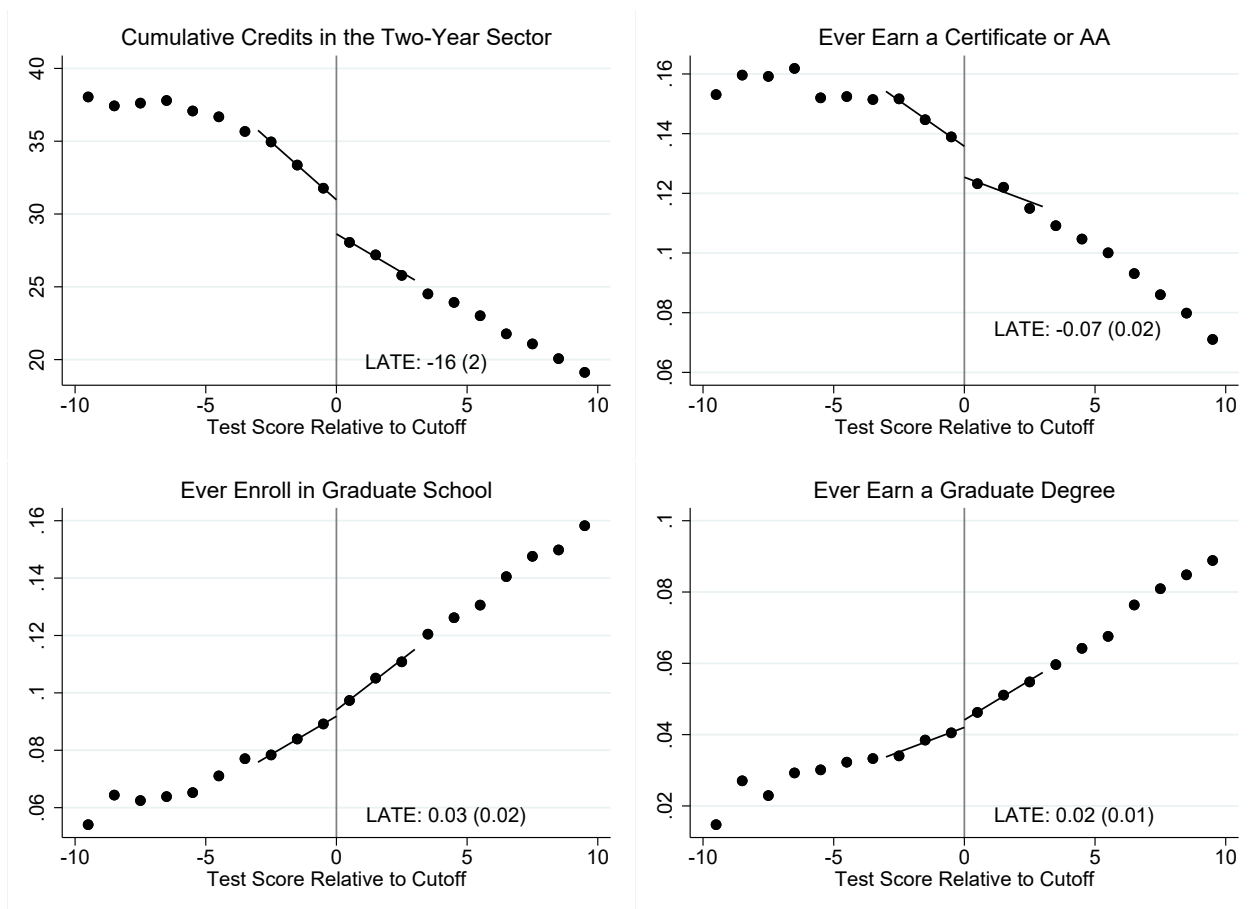Figure A.4: Impacts on Long-Run Educational Outcomes: 8-Year RD Plots

Ever Enroll in the Target University

LATE: 0.90 (0.02)

Test Score Relative to Cutoff

Ever Enroll in Any Four-Year College

LATE: 0.28 (0.02)

Test Score Relative to Cutoff

Cumulative Credits in the Target University

LATE: 74 (2)

Test Score Relative to Cutoff

Cumulative Credits in the Four-Year Sector

LATE: 28 (3)

Test Score Relative to Cutoff

Ever Earn a BA from the Target University

LATE: 0.34 (0.02)

Test Score Relative to Cutoff

Ever Earn a BA from Any Four-Year College

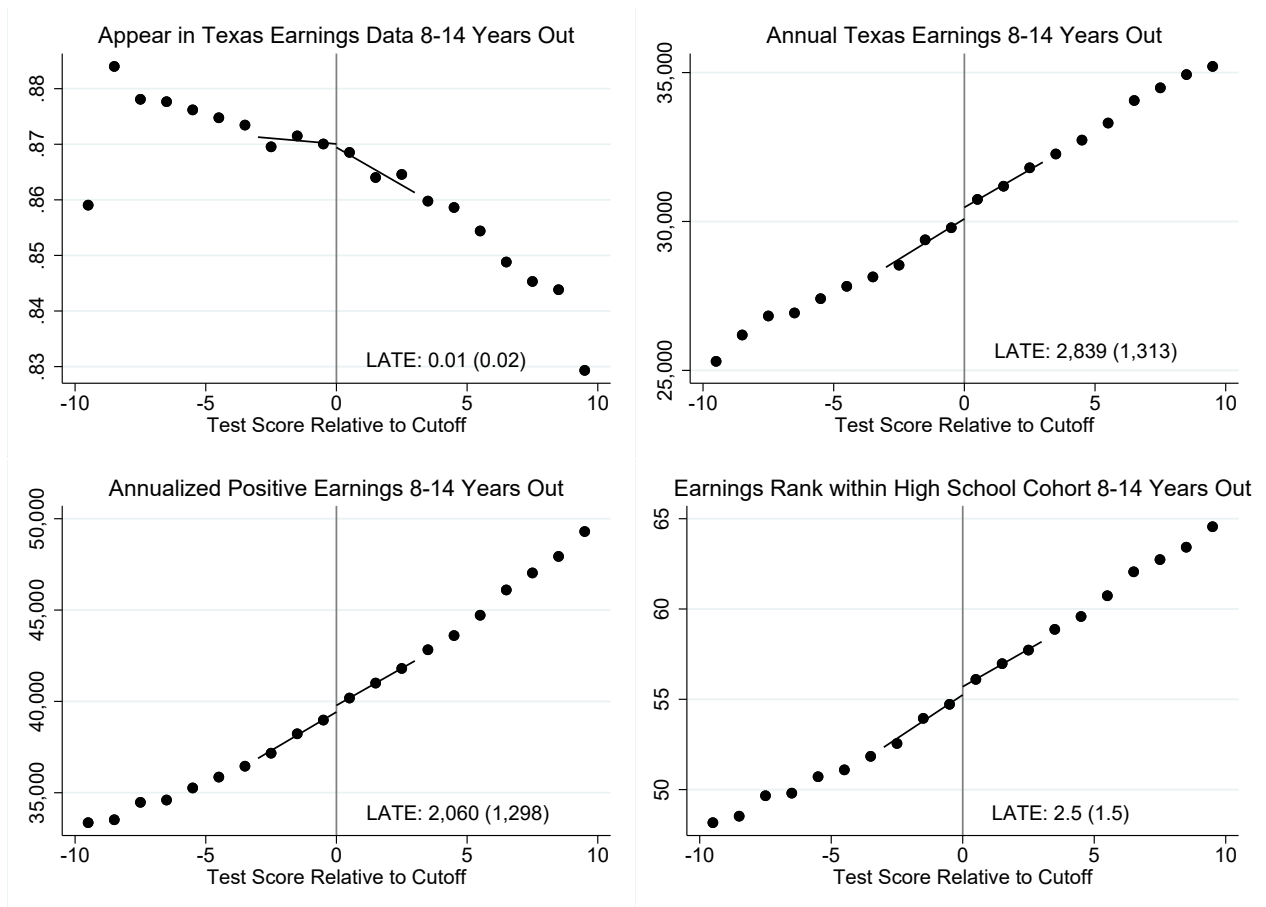LATE: 0.12 (0.03)

Test Score Relative to Cutoff

Notes: These RD plots correspond to the educational outcomes in Figure 7 measured 8 years out from application. The LATE estimates come from the main fuzzy RD specification described in Section 2.4.

48

Figure A.5: Impacts on STEM BA Completion

Ever Earn a BA in a STEM Field

LATE at 8 Years: -0.02 (0.01)

Years Since Application

Complier $E[Y_0]$ Complier $E[Y_1]$

Notes: The gray dots in each plot show the mean outcome of compliers who fall just below the admission cutoff (untreated) at each year since application, and the black dots show the mean outcome of compliers who fall just above the cutoff (treated), which is the untreated mean plus the LATE. The estimates come from the main fuzzy RD specification described in Section 2.4, estimated separately for each year. STEM definition comes from U.S. Department of Homeland Security (2016).

Figure A.6: Impacts on Educational Outcomes in the Two-Year and Graduate Sectors: RD Plots

Notes: These RD plots correspond to the educational outcomes in Figure 8 measured 8 years out from application. The LATE estimates come from the main fuzzy RD specification described in Section 2.4.

Figure A.7: Impacts on Long-Run Educational Costs: 8-Year RD Plots
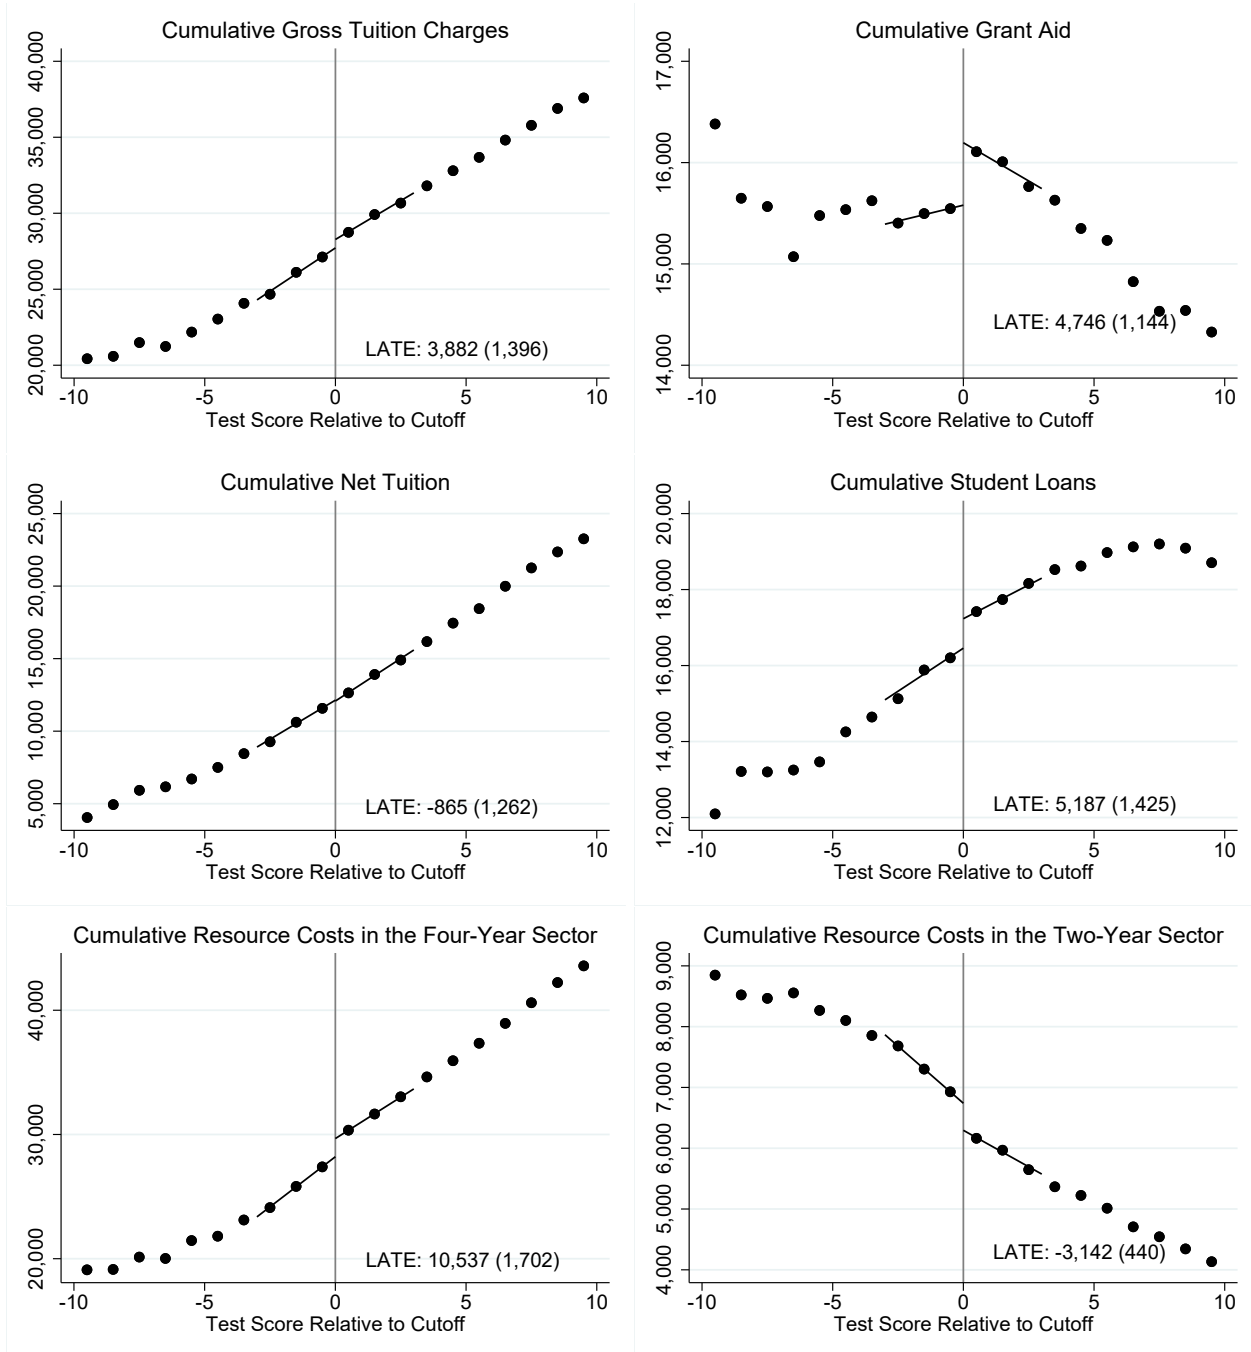
Notes: These RD plots correspond to the cost outcomes in Figure 9 measured 8 years out from application. The LATE estimates come from the main fuzzy RD specification described in Section 2.4.

Figure A.8: Impacts on Earnings: Years 8-14 RD Plots

Notes: These RD plots correspond to the earnings outcomes in Figure 10 measured 8-14 years out from application. The LATE estimates come from the main fuzzy RD specification described in Section 2.4.

Figure A.9: Robustness to Alternative Bandwidths
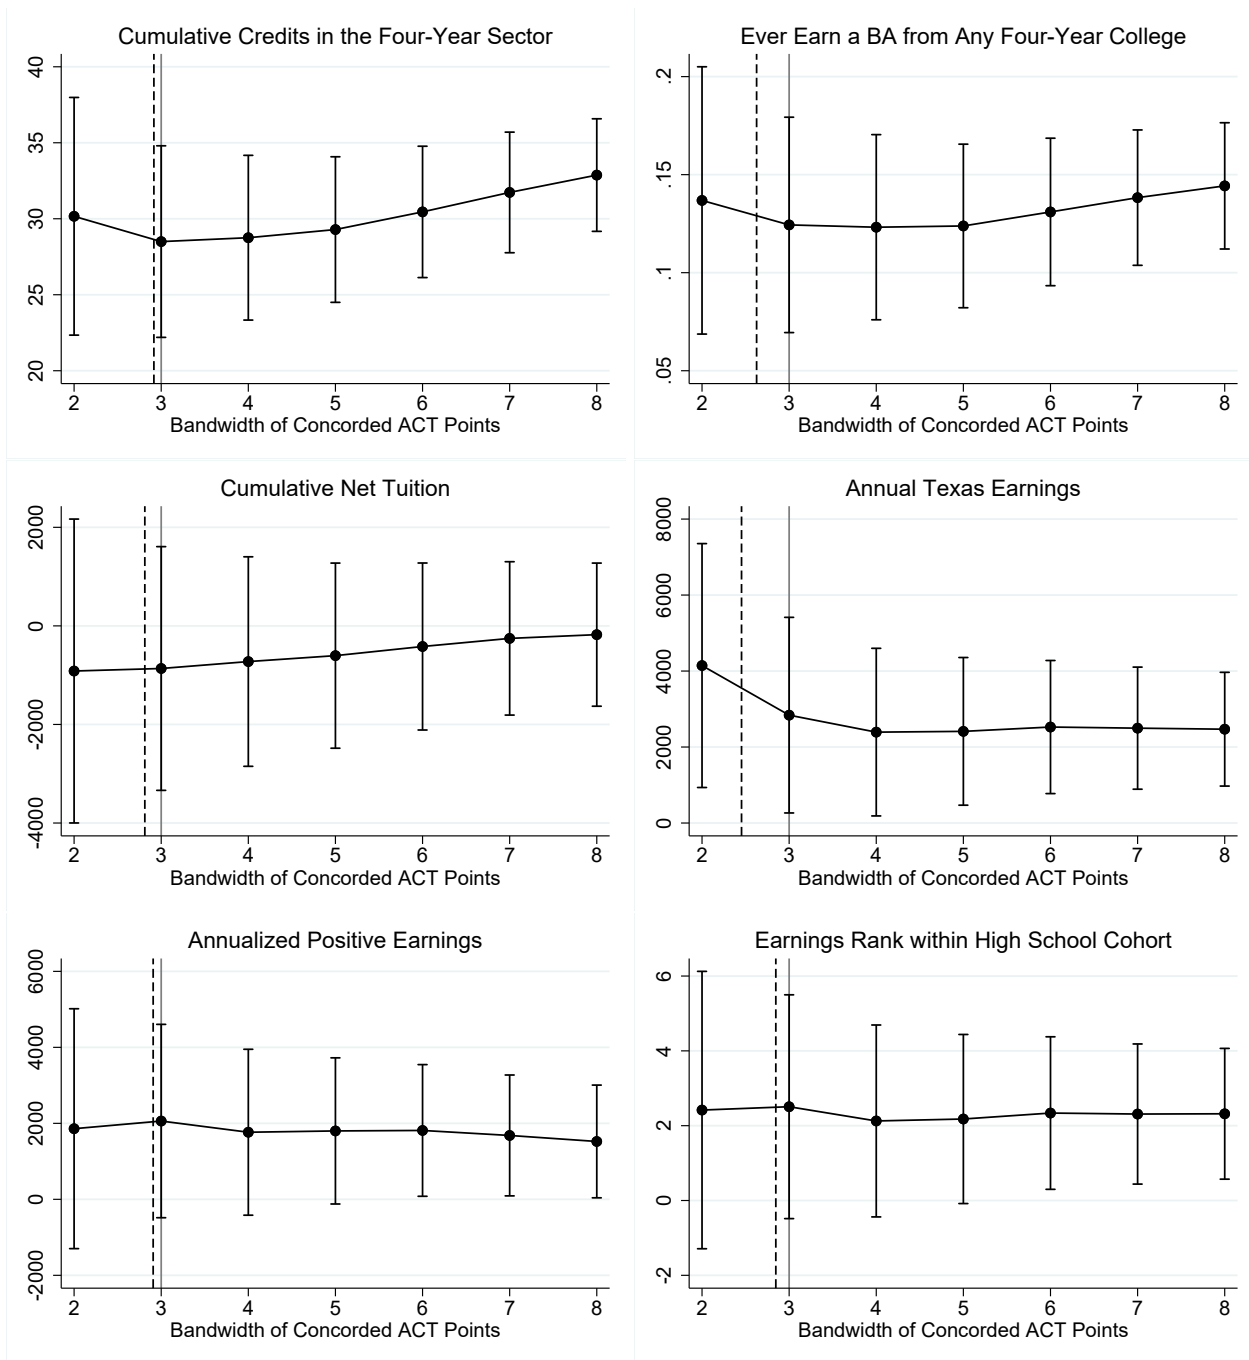
Notes: Each LATE point estimate and 95 percent confidence interval come from the main fuzzy RD specification described in Section 2.4, estimated separately using the different bandwidths of concorded ACT points indicated on the horizontal axis. The solid gray vertical line marks the bandwidth of 3 used in the main specification, and the dashed vertical line marks the bandwidth selected for each outcome by the data-driven method in Calonico et al. (2014), acting as if the discrete running variable were continuous.

Figure A.10: Robustness to Alternative Specifications
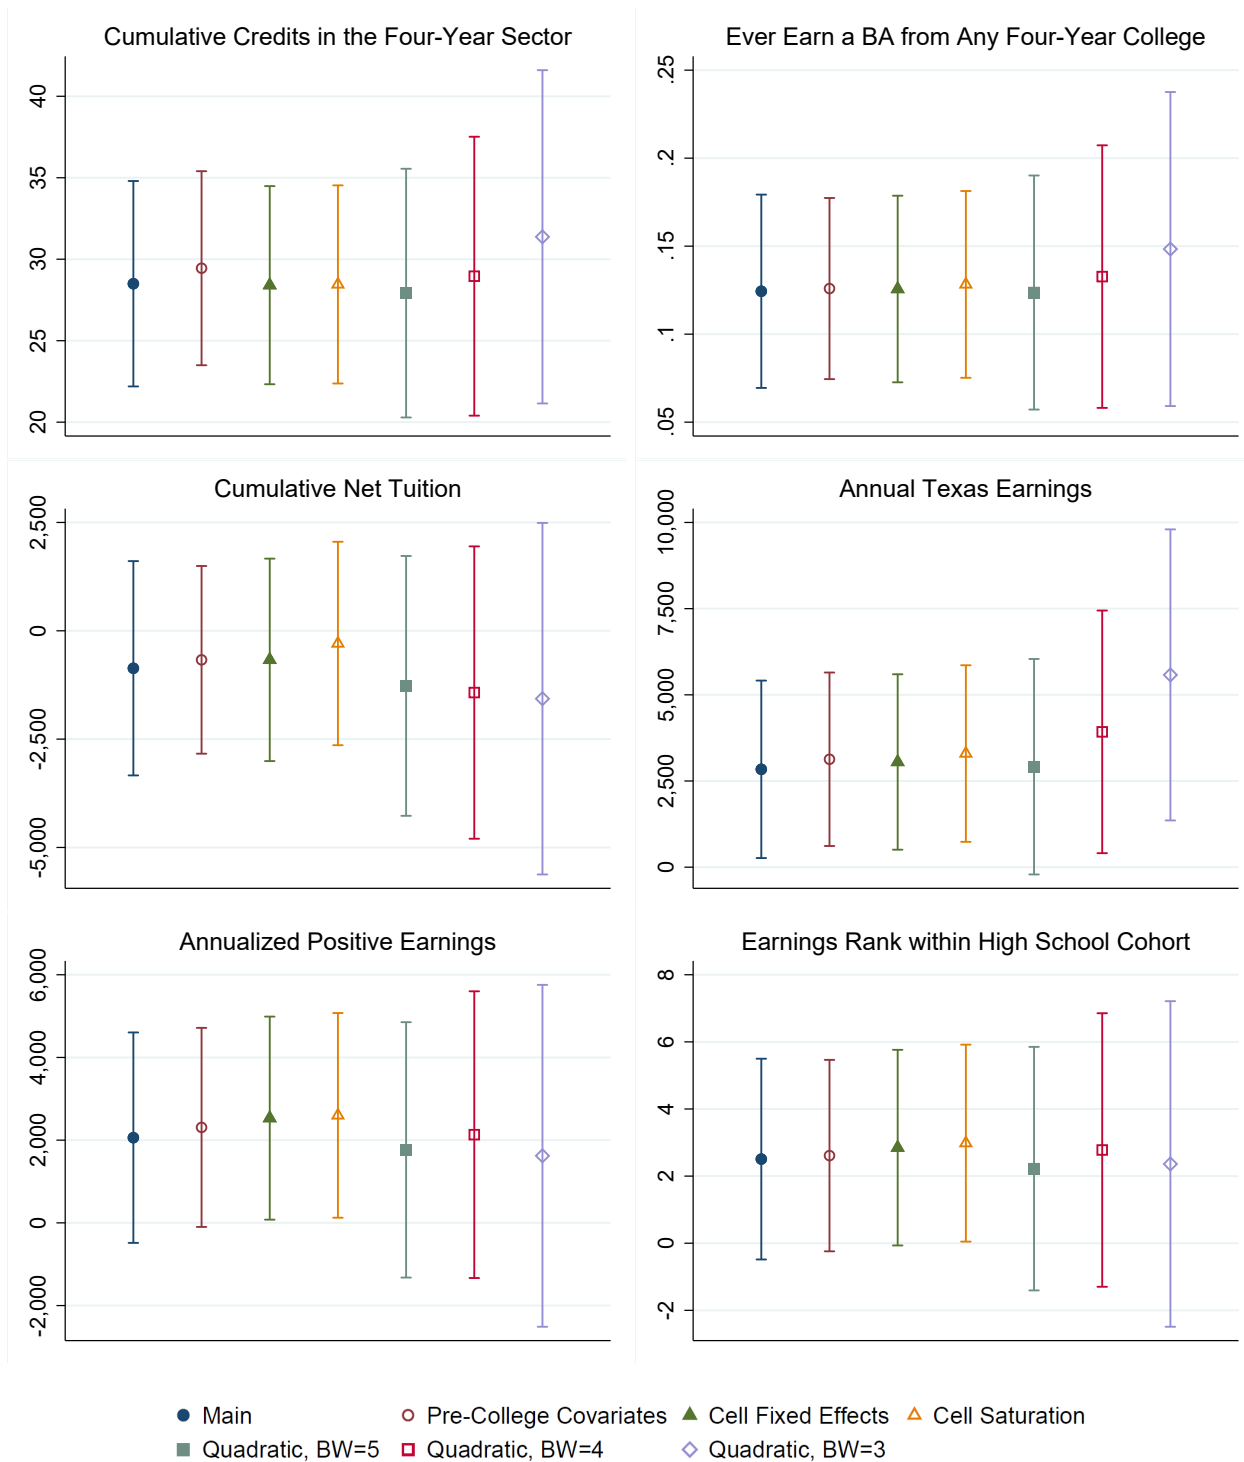
Notes: Each LATE point estimate and 95 percent confidence interval come from a different fuzzy RD specification. "Main" is the main specification described in Section 2.4, which solely controls for linear slopes in the running variable on either side of the cutoff. "Pre-College Covariates" adds all of the covariate controls described in footnote 14 to the main specification. "Cell Fixed Effects" adds indicators for each of the application cells, defined as the interaction between target school, application year, top quartile high school GPA status, and submitted test (SAT or ACT), to the main specification. "Cell Saturation" adds fully saturated interactions of each application cell with the linear running variable slopes. "Quadratic, BW=X" changes the local polynomial functional form from linear to quadratic in the running variable on either side of the cutoff and uses a bandwidth of $X \in \{3, 4, 5\}$ concorded ACT points.

Figure A.11: Crossing Point of Undiscounted Tax Revenue and Taxpayer Cost
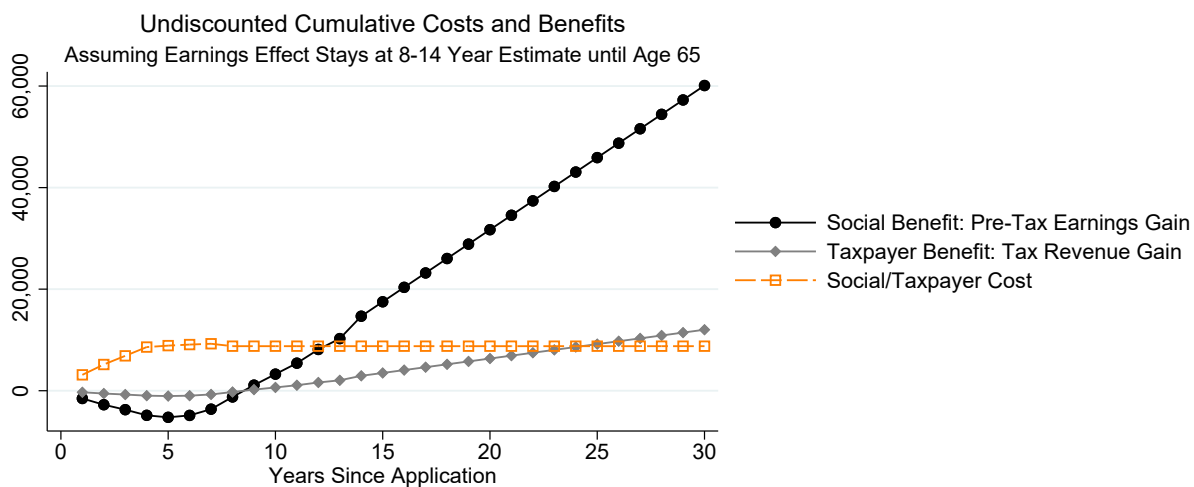
Notes: This figure extrapolates the cumulative tax revenue from the top panel of Figure 11 using the assumption from the middle panel of Figure 11, i.e. that the earnings effect stays at the pooled 8-14 year estimate from year onward.

Figure A.12: Cost-Benefit Calculations Using the Annualized Positive Earnings Measure
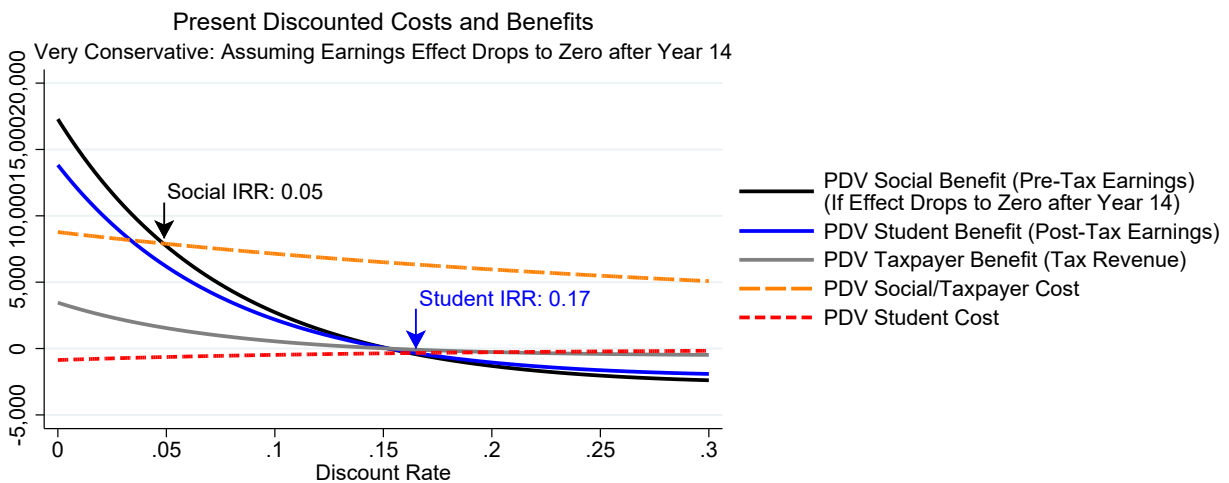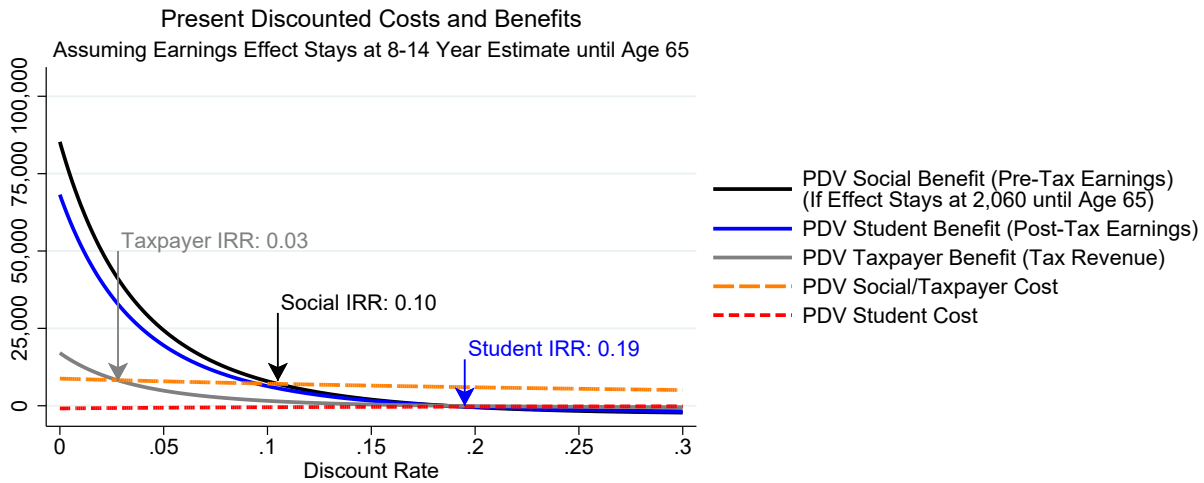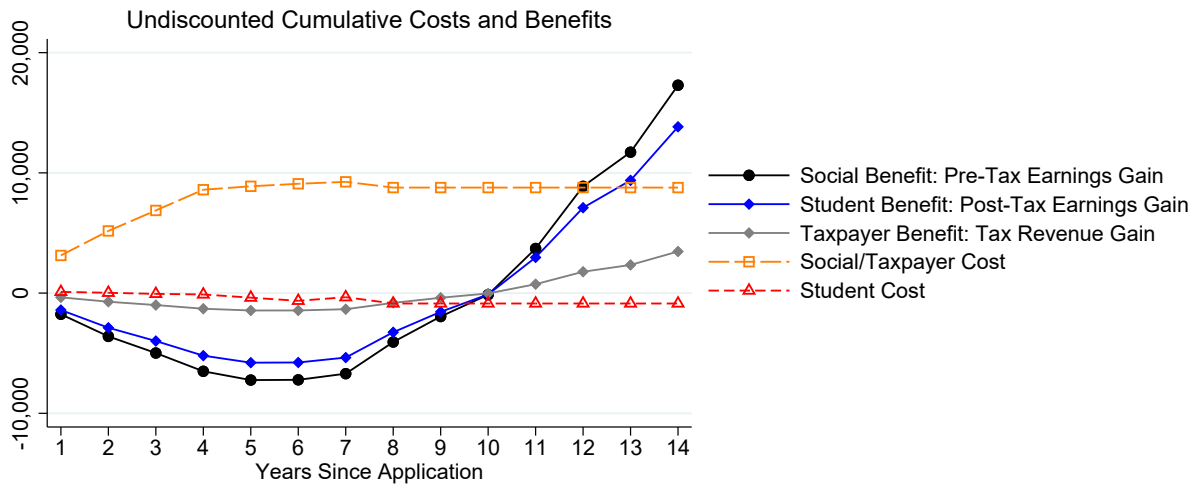
Notes: This figure is identical to Figure 11 except that it uses the annualized positive earnings measure from the bottom left plot of Figure 10 instead of the annual Texas earnings measure from the top right plot of Figure 10.

Figure A.13: The Other-Admissions Stratifer Is Endogenous to Cutoff-Crossing
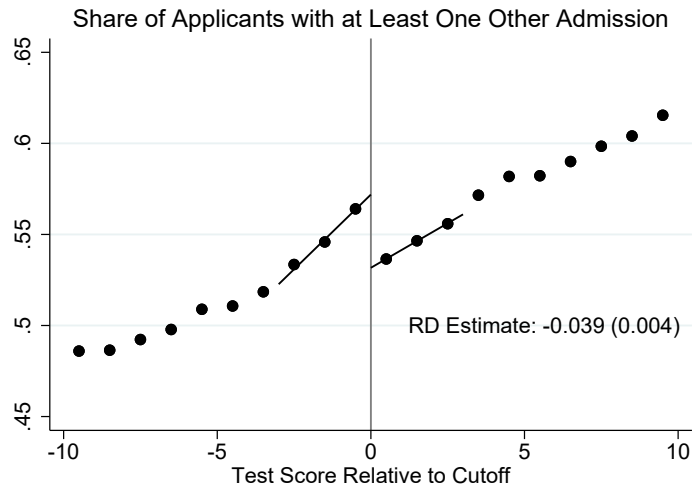
Notes: This figure plots the share of applicants at each of value of the concorded ACT running variable who have at least one other admission to a Texas public university in the year they applied to the target university. The reduced form RD estimate comes from the main RD specification described in Section 2.4.

B   Inferring Admission Cutoffs

To infer the admission cutoffs, I first need to construct a measure of top quartile high school GPA status ("top 25") for each applicant. Many schools automatically admit applicants from the top GPA quartile of their high school class, and most more generally set distinct test score cutoffs for top quartile applicants versus applicants in other quartiles. A challenge is that the raw admissions data only selectively record an applicant as top 25 if she was positively admitted based on that status. However, I am able to infer top 25 status for most applicants by leveraging information about automatic top 25 admissions and information across multiple applications from the same applicant. I first identify all of the school-year cells that have automatic top 25 admission by plotting admission probabilities by test score among applicants to a given school who have a top 25 admission recorded at another school that year. I then assign positive top 25 status to applications to a school-year cell with an automatic top 25 admission policy that are recorded as admitted via top 25. I also assign positive top 25 status to an application to a given school if the applicant is recorded as admitted to another school that year via top 25. I assign negative top 25 status to applications to a school-year cell with an automatic top 25 policy that are not recorded as admitted via top 25. I then aggregate to the applicant level by assigning an applicant positive top 25 status if any of her applications were inferred as top 25; among the remaining applicants, I assign negative top 25 status if any applications were actively inferred as not top 25.

This procedure reveals 30 percent of applicants to be positively top 25 and 52 percent to be non-top 25, but leaves 18 percent unknown. To classify those remaining 18 percent, I use the following pre-college covariates to estimate a logit model of top 25 status among the 82 percent with known status from above. For indicators of female, black, Hispanic, white, graduating high school with minimum requirements, graduating with distinguished requirements, ever at risk of dropping out, ever eligible for free or reduced-price lunch, and gifted program participation, I include the student's indicator value, fixed effects for deciles of the student's high school-level mean of the indicator, and their interaction. For continuous measures of 10th grade standardized math and English scores, number of advanced high school courses taken, percent of days absent, and number of days of disciplinary suspension, I include a cubic in the student's percentile rank of that measure among her graduating high school class. I classify 45 percent of the previously unclassified applicants as positive top 25, and the remaining 55 percent as non-top25, by defining the cut point of the predicted values as the point at which the densities of the predicted values in the training sample among top 25 and non-top 25 applicants cross.

Next, with well-defined application cells in hand, I infer an admission cutoff for each cell using a procedure similar to Hoekstra (2009), Andrews et al. (2017), Carneiro et al. (2019), Altmejd et al. (2021), Brunner et al. (2023), and Miller (2023), among others. I exclude cells in which virtually all applicants are admitted, as these cells do not have meaningful cutoffs to find, as well as a small number of cells with incomplete admissions data. A few cells, for example, only record the applications of students who end up enrolling at the university. A few others record only a trivial number of applicants relative to adjacent years. I also exclude all cells at UT-Austin, as the highly selective flagship's "holistic" admission process does not employ simple cutoffs in SAT/ACT scores in any of my sample years. Within each of the 828 remaining application cells, I estimate a series of local linear RDs centered at each distinct test score value up to the 75th percentile within that cell, using the same specification as the main RD implementation described in Section 2.4: local linear slopes differing arbitrarily on either side of the cutoff within a bandwidth of three concorded ACT points (120 SAT points) and weighted with a triangular kernel. I then define the cutoff for each cell as the test score value with the largest discontinuity in a composite of admission and enrollment, defined as zero if the applicant is not admitted, one if the applicant is admitted but does not enroll,

and two if the applicant is admitted and enrolls. Roughly one-fourth of the candidate cutoffs exhibit statistically significant ($t > 1.96$) discontinuities in the log density of applications or (more rarely) the covariate-predicted BA completion or earnings measures described in Section 2.3, likely because these cutoffs actually corresponded to a university's publicly posted criteria in a given year. I disqualify these potentially publicly known cutoffs from consideration in the search algorithm, but since many individual application cells lack much statistical power to detect manipulation, this exclusion does not mechanically ensure balance when pooling across all cutoffs.

C   Proofs for Section 5

First, I show that

$$\frac{E[YDA|Z=1] - E[YDA|Z=0]}{E[DA|Z=1] - E[DA|Z=0]} = E[Y_1|D\text{-complier and } A\text{-always-taker}].$$

The first term in the numerator is

$$
\begin{aligned}
E[YDA|Z=1] &= E[Y_1|D_1 = 1, A_1 = 1]Pr[D_1 = 1, A_1 = 1] \\
&= E[Y_1|D_1 = 1, D_0 = 0, A_1 = 1, A_0 = 1]Pr[D_1 = 1, D_0 = 0, A_1 = 1, A_0 = 1] \\
&+ E[Y_1|D_1 = 1, D_0 = 1, A_1 = 1, A_0 = 1]Pr[D_1 = 1, D_0 = 1, A_1 = 1, A_0 = 1]
\end{aligned}
$$

thanks to instrument validity and Assumption 1 (no A-defiers). Likewise, thanks to Assumption 2 (all A-compliers are D-compliers),

$$
\begin{aligned}
E[YDA|Z=0] &= E[Y_1|D_0 = 1, A_0 = 1]Pr[D_0 = 1, A_0 = 1] \\
&= E[Y_1|D_1 = 1, D_0 = 1, A_1 = 1, A_0 = 1]Pr[D_1 = 1, D_0 = 1, A_1 = 1, A_0 = 1],
\end{aligned}
$$

so the numerator is

$$
\begin{aligned}
E[YDA&|Z=1] - E[YDA|Z=0] \\
&= E[Y_1|D_1 = 1, D_0 = 0, A_1 = 1, A_0 = 1]Pr[D_1 = 1, D_0 = 0, A_1 = 1, A_0 = 1].
\end{aligned}
$$

Similarly, the terms in the denominator are

$$
\begin{aligned}
E[DA|Z=1] &= Pr[D_1 = 1, A_1 = 1] \\
&= Pr[D_1 = 1, D_0 = 1, A_1 = 1, A_0 = 1] \\
&+ Pr[D_1 = 1, D_0 = 0, A_1 = 1, A_0 = 1],
\end{aligned}
$$

$$
\begin{aligned}
E[DA|Z=0] &= Pr[D_0 = 1, A_0 = 1] \\
&= Pr[D_1 = 1, D_0 = 1, A_1 = 1, A_0 = 1],
\end{aligned}
$$

which implies

$$
\begin{aligned}
E[DA&|Z=1] - E[DA|Z=0] \\
&= Pr[D_1 = 1, D_0 = 0, A_1 = 1, A_0 = 1],
\end{aligned}
$$

which together with the numerator yields the desired result.

Second, I show that

$$\frac{E[Y(1-D)(1-A)|Z=1] - E[Y(1-D)(1-A)|Z=0]}{E[(1-D)(1-A)|Z=1] - E[(1-D)(1-A)|Z=0]} = E[Y_0|D\text{-complier and } A\text{-never-taker}].$$

The first term in the numerator is

$$\begin{aligned}
E[Y(1-D)(1-A)|Z=1] &= E[Y_0|D_1=0, A_1=0]Pr[D_1=0, A_1=0] \\
&= E[Y_0|D_1=0, D_0=0, A_1=0, A_0=0]Pr[D_1=0, D_0=0, A_1=0, A_0=0],
\end{aligned}$$

and the second term is

$$\begin{aligned}
E[Y(1-D)(1-A)|Z=0] &= E[Y_0|D_0=0, A_0=0]Pr[D_0=0, A_0=0] \\
&= E[Y_0|D_1=1, D_0=0, A_1=0, A_0=0]Pr[D_1=1, D_0=0, A_1=0, A_0=0] \\
&+ E[Y_0|D_1=0, D_0=0, A_1=0, A_0=0]Pr[D_1=0, D_0=0, A_1=0, A_0=0],
\end{aligned}$$

so the numerator is

$$\begin{aligned}
E[Y(1-D)(1-A)|Z=1] &- E[Y(1-D)(1-A)|Z=0] \\
&= -E[Y_0|D_1=1, D_0=0, A_1=0, A_0=0]Pr[D_1=1, D_0=0, A_1=0, A_0=0].
\end{aligned}$$

Similarly, the terms in the denominator are

$$\begin{aligned}
E[(1-D)(1-A)|Z=1] &= Pr[D_1=0, A_1=0] \\
&= Pr[D_1=0, D_0=0, A_1=0, A_0=0]
\end{aligned}$$

and

$$\begin{aligned}
E[(1-D)(1-A)|Z=0] &= Pr[D_0=0, A_0=0] \\
&= Pr[D_1=1, D_0=0, A_1=0, A_0=0] \\
&+ Pr[D_1=0, D_0=0, A_1=0, A_0=0],
\end{aligned}$$

which implies

$$\begin{aligned}
E[(1-D)(1-A)|Z=1] &- E[(1-D)(1-A)|Z=0] \\
&= -Pr[D_1=1, D_0=0, A_1=0, A_0=0],
\end{aligned}$$

which together with the numerator yields the desired result.

Third, I show that

$$\frac{E[YD(1-A)|Z=1] - E[YD(1-A)|Z=0]}{E[D(1-A)|Z=1] - E[(D(1-A)|Z=0]} = E[Y_1|D\text{-complier and }(A\text{-never-taker or }A\text{-complier})].$$

The first term in the numerator is

$$
\begin{aligned}
E[YD(1-A)|Z=1] &= E[Y_1|D_1=1, A_1=0]Pr[D_1=1, A_1=0] \\
&= E[Y_1|D_1=1, D_0=0, A_1=0, A_0=1]Pr[D_1=1, D_0=0, A_1=0, A_0=1] \\
&+ E[Y_1|D_1=1, D_0=0, A_1=0, A_0=0]Pr[D_1=1, D_0=0, A_1=0, A_0=0] \\
&+ E[Y_1|D_1=1, D_0=1, A_1=0, A_0=0]Pr[D_1=1, D_0=1, A_1=0, A_0=0],
\end{aligned}
$$

and the second term is

$$
\begin{aligned}
E[YD(1-A)|Z=0] &= E[Y_1|D_0=1, A_0=0]Pr[D_0=1, A_0=0] \\
&= E[Y_1|D_1=1, D_0=1, A_1=0, A_0=0]Pr[D_1=1, D_0=1, A_1=0, A_0=0],
\end{aligned}
$$

so the numerator is

$$
\begin{aligned}
E[YD(1-A)|Z=1] &- E[YD(1-A)|Z=0] \\
&= E[Y_1|D_1=1, D_0=0, A_1=0, A_0=1]Pr[D_1=1, D_0=0, A_1=0, A_0=1] \\
&+ E[Y_1|D_1=1, D_0=0, A_1=0, A_0=0]Pr[D_1=1, D_0=0, A_1=0, A_0=0].
\end{aligned}
$$

Similarly, the terms in the denominator are

$$
\begin{aligned}
E[D(1-A)|Z=1] &= Pr[D_1=1, A_1=0] \\
&= Pr[D_1=1, D_0=1, A_1=0, A_0=0] \\
&+ Pr[D_1=1, D_0=0, A_1=0, A_0=0] \\
&+ Pr[D_1=1, D_0=0, A_1=0, A_0=1]
\end{aligned}
$$

and

$$
\begin{aligned}
E[D(1-A)|Z=0] &= Pr[D_0=1, A_0=0] \\
&= Pr[D_1=1, D_0=1, A_1=0, A_0=0],
\end{aligned}
$$

which implies

$$
\begin{aligned}
E[D(1-A)|Z=1] &- E[(D(1-A)|Z=0] \\
&= Pr[D_1=1, D_0=0, A_1=0, A_0=1] \\
&+ Pr[D_1=1, D_0=0, A_1=0, A_0=0],
\end{aligned}
$$

which together with the numerator yields the desired result.

Fourth, I show that

$$\frac{E[Y(1-D)A|Z=1] - E[Y(1-D)A|Z=0]}{E[(1-D)A|Z=1] - E[(1-D)A|Z=0]} = E[Y_0|D\text{-complier and }(A\text{-always-taker or }A\text{-complier})].$$

The first term in the numerator is

$$\begin{aligned}
E[Y(1-D)A|Z=1] &= E[Y_0|D_1=0, A_1=1]Pr[D_1=0, A_1=1]\\
&= E[Y_0|D_1=0, D_0=0, A_1=1, A_0=1]Pr[D_1=0, D_0=0, A_1=1, A_0=1],
\end{aligned}$$

and the second term is

$$\begin{aligned}
E[Y(1-D)A|Z=0] &= E[Y_0|D_0=0, A_0=1]Pr[D_0=0, A_0=1]\\
&= E[Y_0|D_1=1, D_0=0, A_1=1, A_0=1]Pr[D_1=1, D_0=0, A_1=1, A_0=1]\\
&\quad + E[Y_0|D_1=1, D_0=0, A_1=0, A_0=1]Pr[D_1=1, D_0=0, A_1=0, A_0=1]\\
&\quad + E[Y_0|D_1=0, D_0=0, A_1=1, A_0=1]Pr[D_1=0, D_0=0, A_1=1, A_0=1],
\end{aligned}$$

so the numerator is

$$\begin{aligned}
E[Y(1&-D)A|Z=1] - E[Y(1-D)A|Z=0]\\
&= - E[Y_0|D_1=1, D_0=0, A_1=1, A_0=1]Pr[D_1=1, D_0=0, A_1=1, A_0=1]\\
&\quad - E[Y_0|D_1=1, D_0=0, A_1=0, A_0=1]Pr[D_1=1, D_0=0, A_1=0, A_0=1].
\end{aligned}$$

Similarly, the terms in the denominator are

$$\begin{aligned}
E[(1-D)A|Z=1] &= Pr[D_1=0, A_1=1]\\
&= Pr[D_1=0, D_0=0, A_1=1, A_0=1]
\end{aligned}$$

and

$$\begin{aligned}
E[(1-D)A|Z=0] &= Pr[D_0=0, A_0=1]\\
&= Pr[D_1=0, D_0=0, A_1=1, A_0=1]\\
&\quad + Pr[D_1=1, D_0=0, A_1=1, A_0=1]\\
&\quad + Pr[D_1=1, D_0=0, A_1=0, A_0=1]
\end{aligned}$$

which implies

$$\begin{aligned}
E[(1-D)A|Z=1] &- E[(1-D)A|Z=0]\\
&= - Pr[D_1=1, D_0=0, A_1=1, A_0=1]\\
&\quad - Pr[D_1=1, D_0=0, A_1=0, A_0=1],
\end{aligned}$$

which together with the numerator yields the desired result.