Blueprint Labs

Discussion Paper #2025.04

College as a Marriage Market

Lars Kirkebøen
Edwin Leuven
Magne Mogstad
Jack Mountjoy

May 2025

MIT Department of Economics
77 Massachusetts Avenue, Bldg E52-300
Cambridge, MA 02139

National Bureau of Economic Research
1050 Massachusetts Avenue, 3rd Floor
Cambridge, MA 02138

Marginal Returns to Public Universities

Jack Mountjoy*

August 2025

Latest version available here.

Abstract

This paper studies the returns to American public universities by comparing the long-term outcomes of their barely admitted versus barely rejected applicants. I use administrative admission records spanning all 35 public universities in Texas, which collectively enroll 10 percent of all 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 gains an additional year of education in the four-year sector, becomes 12 percentage points more likely to ever earn a bachelor's degree, and eventually earns 8 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 26 percent for the marginal students themselves, 16 percent for society (which must pay for the additional education), and 6 percent for the government budget. Earnings gains are similar across admitting institutions of varying selectivity, but smaller for students from low-income families, who spend more time enrolled but complete fewer degrees and major in less lucrative fields. Finally, I develop a method to separately identify effects for students on the extensive margin of attending any university versus those on the margin of attending a more selective one, revealing larger effects on the extensive margin.

*University of Chicago Booth School of Business and NBER (email: jack.mountjoy@chicagobooth.edu). For helpful comments and conversations, I am grateful to Joe Altonji, Marianne Bertrand, Dan Black, Zach Bleemer, Chris Campos, Raj Chetty, David Deming, Jeff Denning, Michael Dinerstein, John Friedman, Michael Galperin, Luis Garicano, Andrew Garin, Owen Graham-O'Regan, Lancelot Henry de Frahan, Larry Katz, Pat Kline, Emily Leslie, Jonathan Meer, Lois Miller, Magne Mogstad, Dick Murnane, Rich Murphy, Derek Neal, Matt Notowidigdo, Amanda Pallais, Canice Prendergast, Michael Ricks, Evan Rose, Jon Roth, Jesse Shapiro, Doug Staiger, Alex Torgovitsky, Cody Tuttle, Chris Walters, Seth Zimmerman, four anonymous referees, and many seminar participants. I am also grateful for the expertise of the UT-Dallas Education Research Center staff, especially Holly Kosiewicz, Mark Lu, Trey Miller, and the tragically departed Rodney Andrews. The Robert H. Topel Faculty Research Fund at the University of Chicago Booth School of Business provided valuable research funding. The conclusions of this research do not necessarily reflect the opinions or official position of the Texas Education Research Center, the Texas Education Agency, the Texas Higher Education Coordinating Board, the Texas Workforce Commission, or the State of Texas.

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 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 all 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, 2024). 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.[2] 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

[1] See Barrow and Malamud (2015) for a review, along with recent contributions by Smith et al. (2025), Bleemer (2024), 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.

[2] Relatively simple admission procedures based in part on test score cutoffs are common at American public universities, in contrast to the "holistic" admission practices at highly selective private institutions that enroll far fewer students but receive far more media attention. "Public universities with huge applicant pools and large numbers of incoming students typically use an ax rather than the scalpel smaller private colleges employ" (Selingo, 2020).

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 answer the question of whether public universities should expand or contract along their 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 alternatives; on the other, they may incur substantial costs to themselves and to taxpayers 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 admission cutoffs at a handful of isolated institutions, this paper marshals hundreds of cutoffs spanning nearly the entire public university sector of the second largest U.S. state. 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 private and social returns to enrolling the typical marginal public university applicant, but also, as I detail below, separate estimates across universities of widely differing selectivity, across applicants from different demographic backgrounds, and across the economically distinct margins of attending any four-year institution versus attending a more selective one.

The paper proceeds as follows. After describing the data sources and linkages, Section 2 introduces the admission cutoffs 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 that university's SAT/ACT admission cutoff causes the probability of admission to jump abruptly by 27 percentage points, leading to a precisely estimated 15 percentage point first stage increase in the probability of enrolling at that university. The density of applicants and their pre-college characteristics are smooth through these cutoffs, justifying the use of local cutoff-crossing as an exogenous instrument for enrollment. Applicants can and do apply to multiple public universities, but virtually all applicants are marginal around at most one university's cutoff, which I refer to as that applicant's "target" university. The local average treatment effects identified by this fuzzy RD design therefore pertain to compliers who enroll in their target university if and only if they barely cross its admission cutoff. In terms of observable characteristics, the typical cutoff complier is significantly more disadvantaged than the average college applicant, but rather comparable to the average high school graduate.

The main causal estimates in Section 3 begin by documenting that marginal applicants to a given public university have multiple fallback options if they don't get in. Notably, half of cutoff compliers would fall back to another, typically less selective, four-year institution if rejected from their target university. The pooled results in Section 3 are therefore an equally weighted mix of both "intensive margin" effects of starting college at a more selective four-year institution and "extensive margin" effects of starting college at any four-year institution. Section 7 develops a

method to disentangle these economically distinct effects, while the results until then reflect the policy-relevant return to enrolling marginal applicants relative to their actual mix of next-best alternatives. Interestingly, only 6 percent of all cutoff compliers would 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. Nonetheless, due to substantial variation in resources and peers across different institutions, both within and across sectors, cutoff-crossing induces large changes in traditional quality measures of the institution that the marginal student first attends, including peer test scores, sticker-price tuition, educational spending per student, the institutional graduation rate, and peer earnings.

The remainder of Section 3 estimates long-run impacts on education and earnings. In terms of educational attainment, the average cutoff complier ultimately completes one full year's worth of additional credits in the four-year sector and becomes 12 percentage points more likely to ever earn a bachelor's degree from any institution. Only some of this gain 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. All of the gains in BA completion are in non-STEM majors, with no detectable increase in STEM degrees. Cutoff compliers become a bit more likely to attend graduate school and complete a graduate degree.

Turning to earnings impacts, I first show that marginal admission causes no change in the likelihood of appearing in the Texas earnings data, which assuages concerns about differential attrition (Foote and Stange, 2022). I then trace out the earnings trajectories of cutoff compliers using several different measures, including dollar earnings with and without imputed zeros, log earnings, and earnings rank within age cohort. All of these earnings measures deliver a consistent pattern of dynamic effects. Initially, in the first five years after application, admitted compliers earn less their rejected counterparts, as they are much more likely to be actively enrolled in a four-year program. Year six is the crossover age, at which point most of the admitted compliers have finished their college education, entered the workforce full-time, and just started to outearn their rejected counterparts. A statistically and economically substantive earnings premium of 8-9 percent starts to solidify around eight years out from application and persists thereafter. Since cutoff compliers complete almost one additional year of education in the four-year sector, these results imply a back-of-the-envelope gross return to a year of four-year college of 8-9 percent. In terms of ranks, the typical rejected complier ends up right around the 50th percentile of their cohort earnings distribution, with admission boosting that rank by 4-5 percentiles. 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, local polynomial functional forms, and alternative methods of inferring the relevant admission cutoffs.

In Section 4, I conduct a formal cost-benefit analysis of the private, social, and taxpayer returns to enrolling marginal public university students. I first use the main RD specification to estimate

the cost side of the ledger. I find that the average marginally admitted student actually pays no additional net tuition, with their $4,600 in additional gross tuition charges nearly fully offset by additional grant aid. Marginal students do end up taking out an additional $5,300 in student loans, likely to finance higher room and board charges and additional consumption at four-year colleges. From society's perspective, of course, the additional educational investments in marginal students are not free; I estimate that the average cutoff complier generates about $10,000 of additional educational expenditures at the institutions they attend.

I then combine these cost estimates with the estimated earnings trajectories to conduct a series of cost-benefit calculations. I first show that the undiscounted cumulative earnings benefits of enrolling marginal students eventually surpass the cumulative costs, but at different horizons for students, society, and taxpayers. For students, the initial earnings foregone by increased enrollment keep their benefits below their (roughly zero) net tuition costs until 8 years out from application, after which the net gain turns positive and grows steadily.[3] Society must wait longer for the earnings benefits to outweigh the resource costs of educating the marginal students, which occurs after 11 years.[4] For taxpayers, the additional tax revenue generated from the earnings gains surpasses the costs after 19 years. I then add a lifecycle horizon and discounting to the analysis to calculate net present values and internal rates of return. At a 3% discount rate, the lifetime net present value of enrolling the typical marginal applicant is about $80,000, with $70,000 pocketed by the student herself and taxpayers netting the remaining $10,000. The discount rates at which these NPVs decline to zero define the internal rate of return from each perspective: a substantial 26 percent for students, 16 percent for society, and 6 percent for the government budget. I conclude Section 4 by showing that these large net returns are robust across a range of alternative assumptions about future earnings growth and cost definitions.

In Sections 5, 6, and 7, I unpack the pooled RD estimates across three important dimensions of heterogeneity. First, Section 5 explores heterogeneity across the wide range of public universities in the sample. Cutoff compliers who are on the admission margin at more selective institutions experience substantially larger increases in peer quality compared to cutoff compliers who are on the admission margin at less selective institutions, but they are also less likely to be on the extensive margin of four-year college enrollment and experience smaller gains in BA attainment. These potentially offsetting factors lead to no systematic difference in the gross earnings gains reaped by cutoff compliers at more versus less selective institutions. There is also no systematic difference in the additional cumulative cost of educating cutoff compliers at more versus less selective institutions, while compliers themselves actually pay a bit less in additional net tuition when admitted at more selective institutions. When stratifying institutions by average educational expenditure per student rather than selectivity, there is some suggestive evidence that complier earnings gains are larger at

[3] I lack data on household formation and thus focus on the applicant's own earnings gain, which likely understates their household return inclusive of marriage market effects (Kirkebøen et al., 2025).

[4] I define the social gross benefit as simply the pre-tax earnings gain, which 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).

more resourced institutions. The earnings effects we would predict for cutoff compliers if we were to use the most common measure of a college's value-added, which is simply the mean earnings of its students relative to students at other institutions, would overpredict the value-added actually experienced by the average cutoff complier by a factor of two: they attend an institution with $6,700 higher average peer earnings as a result of admission, but gain only $3,300 themselves. Such an approach would also overpredict the relationship between selectivity and value-added by a factor of three: a 100 SAT point increase in the selectivity of a complier's target school predicts a $3,000 higher gain in peer earnings, but only a statistically insignificant $900 higher gain in the complier's own earnings.

Section 6 explores heterogeneity across students from different demographic backgrounds. With respect to gender, female and male cutoff compliers eventually reap similar gains in log earnings and earnings ranks, with somewhat larger dollar gains for men given their larger baseline levels. Women begin to reap their positive gains more quickly than men, however, likely explained by men taking longer to finish college. With respect to family income, as proxied by eligibility for free or reduced-price lunch in high school, compliers from low-income families experience significantly smaller earnings gains compared to compliers from higher-income families. The difference is not explained by differential gains in college quality induced by admission, but rather by low-income compliers being less likely to complete a degree, spending more time in college despite their fewer degrees, and majoring in less lucrative fields compared to their higher-income peers.[5] With respect to race, white and Asian compliers reap similar gains in earnings and degree completion compared to Black and Hispanic compliers, despite white and Asian students experiencing larger increases in college selectivity and spending per student as a result of admission.

Finally, in Section 7, I return the fact that half of cutoff compliers would initially fall out of the four-year sector if rejected, while the other half would fall back to another, typically less selective, four-year institution. Since these "extensive" and "intensive" treatment contrasts represent two economically distinct parameters of interest—the return to attending any four-year college versus the return to four-year selectivity—I develop a method to learn about their separate contributions to the pooled effects above. The method grapples with the challenge that these two types of compliers are not directly distinguishable in the data, since marginal applicants' fallback options are never explicitly measured in this setting. I do observe, however, each applicant's entire portfolio of admissions to all 35 Texas public universities, which 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. Unfortunately, this stratifier itself is somewhat endogenous to cutoff crossing: marginal applicants to a given university who are barely rejected are sometimes still able to secure admission to a fallback school within the same application year, thanks to rolling admissions and spring admission cycles at some schools. Nonetheless, I show how to make progress with this strong but endogenous stratifier towards the goal of isolating

[5] See Belley and Lochner (2007), Bailey and Dynarski (2011), Bleemer and Mehta (2024), and Bleemer and Quincy (2025), among others, for related results on family income gaps in degree completion, major choice, and earnings.

extensive versus intensive margin effects. I first separately identify treated and untreated mean potential outcome levels of different combinations of "response types" with respect to the effect of cutoff crossing on having any other admissions, and then impose an empirically-informed rank assumption across these response types that immediately delivers tightly informative upper and lower bounds on each of the extensive and intensive margin effects of interest. The results show that the pooled effects of enrolling marginal public university applicants are driven by larger effects on extensive margin compliers who would not initially enroll in any four-year college if rejected, with smaller contributions from the other half of compliers on the margin between a more-selective versus less-selective four-year school.[6]

1.1 Contributions and Comparisons to the Existing Literature

This paper advances the small but growing literature using exogenous admission variation to study causal impacts of American colleges on their students' outcomes.[7] 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 larger and more diverse public university sector, allowing for more statistical precision, student diversity, and institutional breadth in studying the returns to American public universities, both on average and across different types of institutions and students. I also develop methods to distinguish returns 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 (2024), 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 (2024) speak primarily to the intensive margin impacts of shifting high-GPA but low-SAT students across different institutions within the four-year sector.[8] This paper marshals broad ad-

[6]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 no 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.

[7]A more established literature has exploited admission cutoffs in other countries, many of which use centralized admissions systems that 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); Zimmerman (2019); Anelli (2020); 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). Also see Dynarski et al. (2023) for a recent review of the literature studying impacts of financial aid programs on college student outcomes, and Galperin (2023) and Londoño-Vélez et al. (2023) for recent contributions. See Miller (2023) for an admission discontinuity approach to identifying the impacts of transferring between postsecondary institutions.

[8]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,

missions variation across applicants, target institutions, and the intensive and extensive margins of four-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 both gross and net 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. (2025) 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.

Compared to the handful of prior papers that do estimate earnings effects of marginal admission into American public universities, this paper's estimate of 8% may initially appear small: Zimmerman (2014) reports a 22% earnings gain, 26% in Kozakowski (2023), 17% in Smith et al. (2025), 21% in Bleemer (2024), and 20% in Hoekstra (2009). Those larger prior estimates, however, are likely readily explained by the larger gains in educational attainment and access to college quality induced by their respective natural experiments. In Zimmerman (2014), enrollment compliers gain roughly three additional years of education in the four-year sector,[9] which is three times larger than the one-year gain experienced by the enrollment compliers in this paper. Similar calculations show that the estimates in Kozakowski (2023) and Smith et al. (2025) imply much larger gains in bachelor's degree completion per enrollment complier compared to this paper, likely explaining their larger earnings estimates.[10] In a similar fashion along on the college quality margin, the California public university cutoff compliers in Bleemer (2024) experience an order of magnitude larger increase in institutional spending per student (roughly $30,000) compared to the Texas pub-

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).

[9] In Table 4 of Zimmerman (2014), the reduced-form effect of admission cutoff-crossing is an additional .644 full-time-equivalent terms of enrollment in the Florida public university system, i.e. .322 years. The enrollment first stage of attending Florida International University (the single target university in the sample) is .104. Thus enrollment compliers gain $.322/.104 = 3.1$ additional years of four-year enrollment on average.

[10] In Kozakowski (2023), .151 additional BAs are created per marginal admission; scaling this by the effect of admission on enrollment in a state university of .58 implies a BA gain for enrollment compliers of $.151/.58 = 26$ percentage points, compared to 12 percentage points in this paper. Table 4 of Smith et al. (2025) reports a 37.2 percentage point gain in BA completion among enrollment compliers.

lic university compliers in this paper (roughly $3,000). This much steeper gradient in institutional resources across the hierarchy of California public universities compared to those in Texas may more generally explain why estimates of the return to college quality appear to be larger in California than Texas, e.g. Bleemer (2022, 2024), Bleemer and Quincy (2025), Chetty et al. (2023), and Mountjoy and Hickman (2021). Larger gains in institutional resources experienced by cutoff compliers could also explain the larger earnings effect estimate in Hoekstra (2009), though that paper only has access to enrollment data from the single unnamed target university and thus cannot measure the changes in institutional characteristics experienced by cutoff compliers.

Finally, 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 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 net present values and internal rates of return to educational investments.[11] Finally, my bounding approach to separately identify impacts for students on the intensive versus extensive margins of the four-year sector contributes to a growing literature that grapples with multiple treatment margins of educational choices and develops methods to disentangle their distinct causal contributions.[12]

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, 2025). Texas is the second largest U.S. state by land area and population (over 30 million people) and the fastest growing state in population level (nearly 500,000 net increase annually). As its own country, Texas would rank as the 8th largest economy in the world ($2.7 trillion GDP).

The analysis sample begins with the universe of students who graduate from a Texas public high school between 2004 and 2014. The 2004 cohort is the first to have SAT and ACT scores recorded in the college admission records, and I stop at the 2014 cohort to observe a balanced panel of ten 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

[11] 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).

[12] E.g., 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).

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 four-year universities, maintained by the Texas Higher Education Coordinating Board (THECB).[13] 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.[14]

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 2008-2014 cohorts, National Student Clearinghouse records are also 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, 2024) to construct student-year-level cost measures of gross tuition, grant aid, net tuition, loan accumulation, and colleges' per-student educational expenditures. The availability of the financial aid data lags one year behind the other outcome datasets, so the cost outcomes are observable for a balanced panel of nine years (instead of ten) for all sample cohorts.

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.[15] 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

[13]The THECB does not collect applications or admissions data from 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.

[14]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.

[15]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.

quartile of high school class GPA rank.[16,17] 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 cell as the combination of the "target" university targeted by the application,[18] the application year, an indicator for whether the applicant was in the top GPA quartile of their high school class,[19] and which test they submitted (SAT or ACT).[20] I exclude application 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 roughly 700 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

[16]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 such students from the RD analysis sample; Figure A.1 shows that Top Ten Percent applicants are relatively rare and distributed smoothly around the admission cutoffs used in this paper.

[17]For more recent cohorts of applicants well after those used in this paper, especially in the wake of the COVID-19 pandemic, several Texas public universities have gone "test optional" and now list separate guaranteed admission criteria for applicants with and without SAT/ACT scores.

[18]Applicants can apply to multiple universities and thus can appear in multiple application cells. However, only one percent of applicants in the RD sample are marginal to more than one university's cutoff, as measured by being within one ACT point or ten SAT points; dropping these applicants yields similar results. Generally, then, we can refer to a single "target school" that is relevant for each marginal applicant.

[19]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.

[20]The admission cutoffs studied in this paper do not appear to vary systematically by applicant race, as explored in Appendix B.

Figure 1: Distribution of SAT and ACT Admission Cutoffs, Complier-Weighted

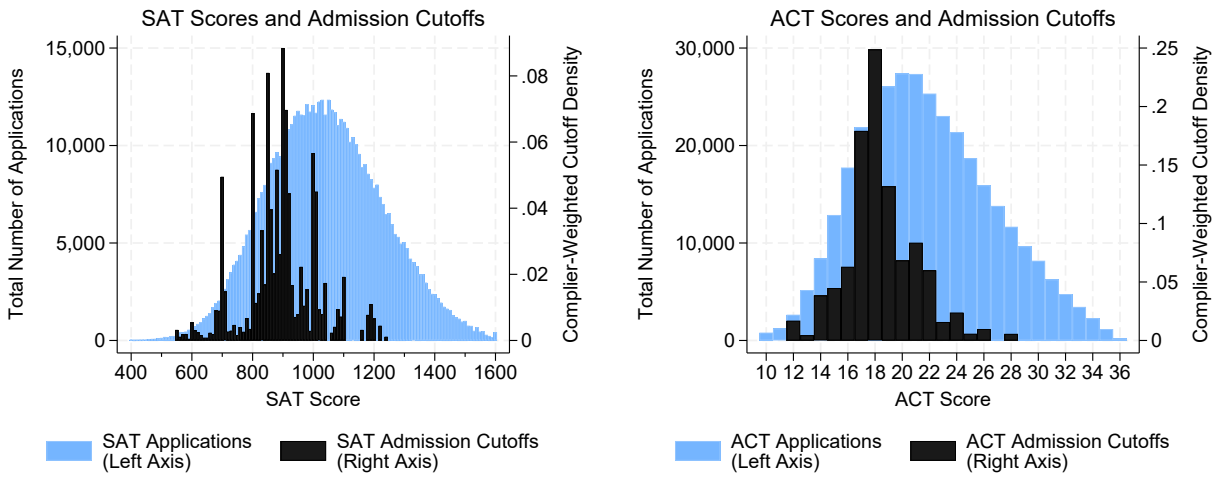

Notes: The lightly shaded histogram in the background shows the population distribution of test scores among all applicants. In the darker shaded foreground, each observation in the histogram is the location of the cutoff within a given application cell, weighted by its complier mass (the number of applications immediately around the cutoff multiplied by that cutoff's estimated first-stage discontinuity in enrollment).

the second-stage RD estimator unaffected by the fact that the cutoffs are estimated from the data rather than known ex ante. Roughly 20 percent of the initially inferred cutoffs exhibit statistically significant (at the 1% level) 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. For the main results, I disqualify these potentially publicly known cutoffs from consideration in the search algorithm. Since many individual application cells lack much statistical power to detect such manipulation, this exclusion does not mechanically ensure balance when pooling across all cutoffs below. Importantly, the robustness checks in Section 3.5 show that the main results are similar when including these cutoffs rather than disqualifying them. That section also shows robustness across additional alternative methods of inferring the relevant cutoffs, including requiring that each cutoff have a statistically significant first stage discontinuity and allowing each cell to contribute multiple cutoffs.

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 mass of compliers around that cutoff (the number of applicants 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 round numbers like 700, 800, 900, and 1000; these values are not prioritized in any way by the cutoff location algorithm, but rather reveal

Figure 2: Distribution of Universities Contributing Cutoffs, Complier-Weighted

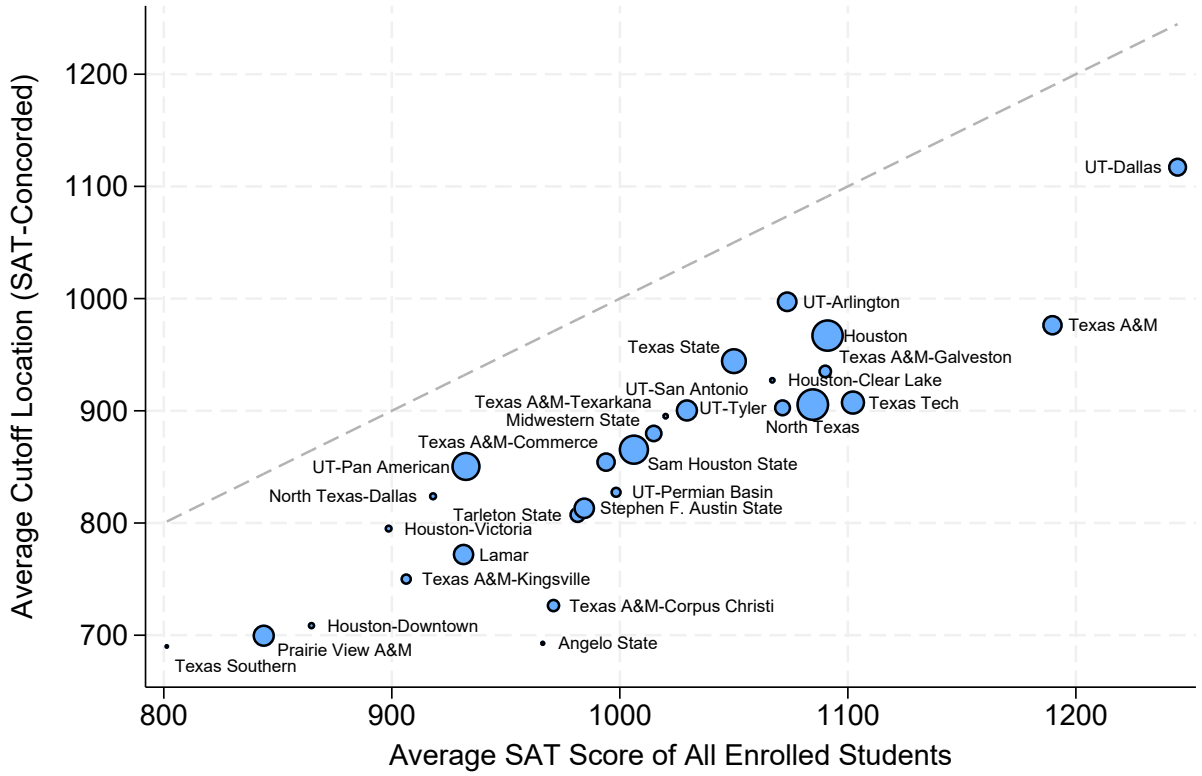

Notes: This figure plots the average SAT-concorded cutoff location within each public university in the analysis sample on the vertical axis against the average SAT score of all enrolled students on the horizontal axis, with the size of the circles reflecting the relative mass of cutoff compliers that each university contributes to the analysis sample.

the simple heuristics actually used by many admissions offices to ration admission offers.[21]

Figure 2 plots the average SAT-concorded cutoff location within each public university in the analysis sample, with the size of the circles reflecting the relative mass of cutoff compliers that each university contributes. On the horizontal axis is the university's average SAT score among enrolled students. The comparison to the dashed 45 degree line confirms that every school's marginal applicants score well below the average of their potential peers, typically around 100-200 SAT points lower. The figure also illustrates both the diversity of these institutions and their diffuse distribution, with the largest contributors of marginal applicants (University of North Texas and University of Houston) each contributing only 11 percent of the total mass of cutoff compliers. The paper's main pooled results in Sections 3 and 4 thus reflect a broad distribution of public

[21] The evidence suggests that these cutoffs arise from decentralized policy decisions at the level of the institution. First, most of the variation in cutoff locations in Figure 1 is across institutions, and within institutions across applicant GPA quartile, and thus not reflective of a uniform statewide policy, as in Smith et al. (2025) and Kozakowski (2023). Second, most institutions have relatively little variation in their cutoff locations across years; Figure 2 plots the average cutoff location of each school. Finally, the structure of the recovered cutoffs mirrors the structure of the decentralized "assured" admissions criteria that vary across university admissions websites, typically featuring distinct thresholds for applicants inside versus outside the top GPA quartile of their high school class, but with the recovered thresholds often appearing at lower SAT/ACT values to allow for some discretion relative to the assured criteria.

institutions; at the same time, the availability of cutoffs across such a wide diversity of schools motivates and enables Section 5's investigation of potential heterogeneity across them.

2.3 RD Diagnostics

Figures 3, 4, and 5 conduct three important diagnostics on the suitability of these admission cutoffs as the basis of a fuzzy regression discontinuity research design. First, Figure 3 examines the first-stage relevance of cutoff-crossing for admission and enrollment at the school targeted by the application. The figure plots the nonparametric probabilities of admission and enrollment conditional on each unique value of the running variable, defined as the applicant's test score minus the admission cutoff she faces given her application cell (i.e., the school targeted by the application, the year she applied, her top quartile GPA status, and the test she submitted). Large discontinuities emerge clearly at the cutoffs for both SAT and ACT submitters. The bottom panels pool both types of submitters by concording SAT scores to ACT scores and use the main RD specification, detailed in the next subsection, to quantify the discontinuities. The bottom left panel shows that the typical marginal applicant's probability of admission into the target school jumps by a precisely estimated 27 percentage points across the cutoff. The bottom right panel shows the first stage in the fuzzy RD design: crossing the cutoff leads to a precisely estimated 15 percentage point increase in the probability of enrolling at the target institution, with a standard error of 0.3 percentage points and an F-statistic of 2,025.

Second, Figure 4 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 5 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 10 years on a suite of pre-college covariates.[22] 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 8-12 years out from application. Figure 5 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 with no discontinuities at the admission cutoff. 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

[22] 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 3: RD Diagnostics: First Stage Discontinuities in Admission and Enrollment

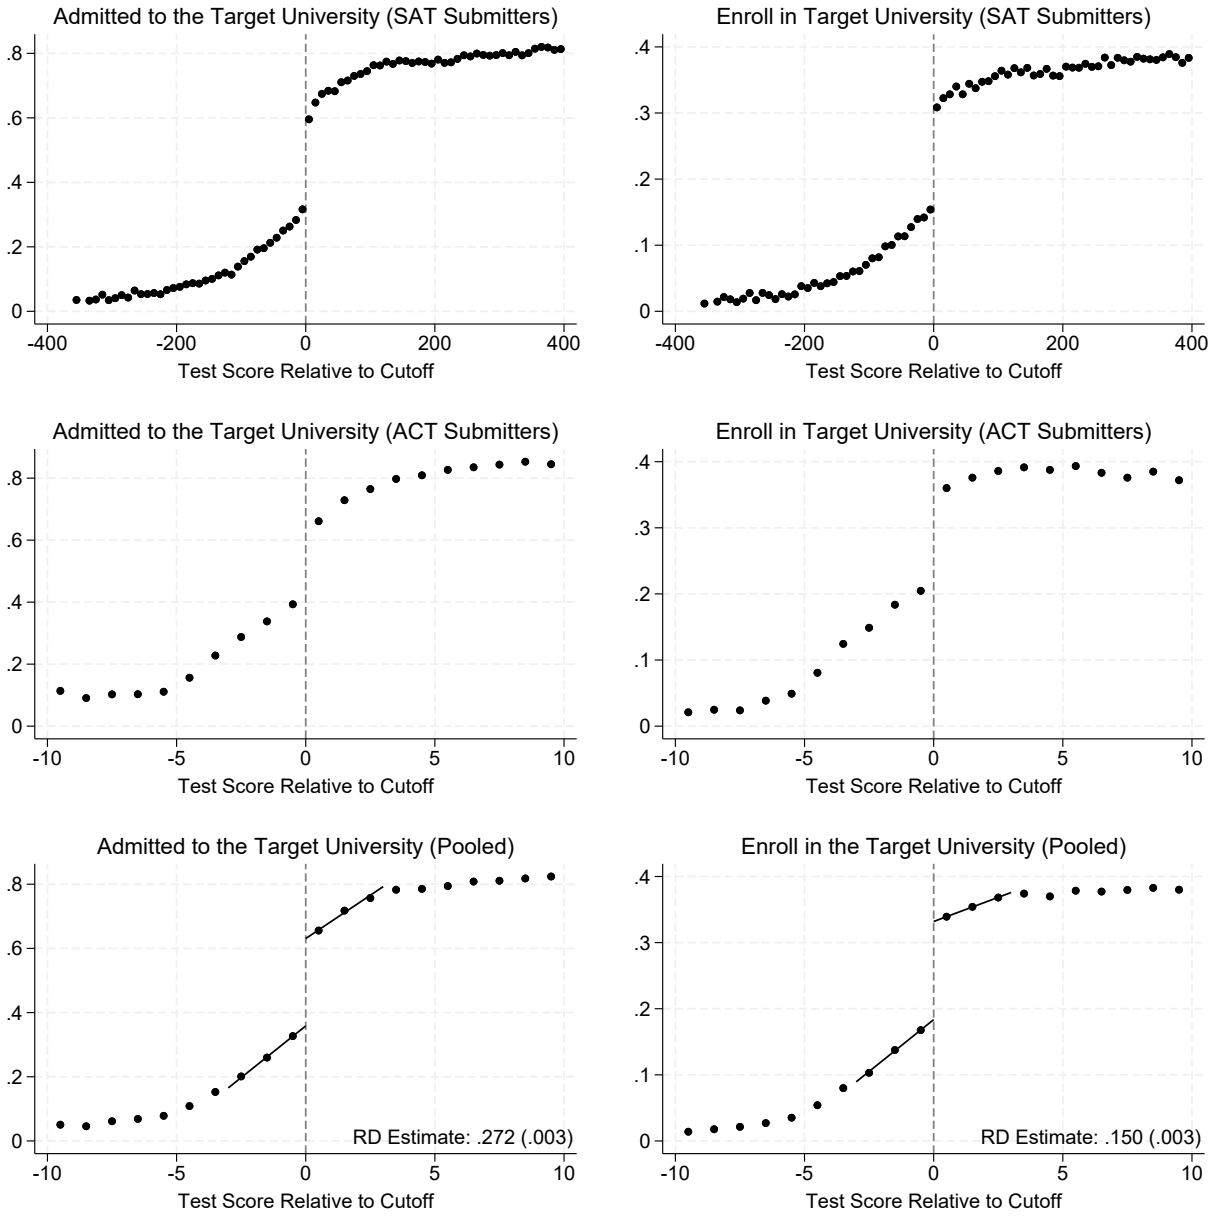

Notes: This figure plots the probability of admission and enrollment for each unique value of the running variable for each test, defined as the applicant's test score minus the admission cutoff she faces given her application cell. The bottom panel pools SAT and ACT submitters by dividing each SAT submitter's 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.

Figure 4: RD Diagnostics: Density of Applications

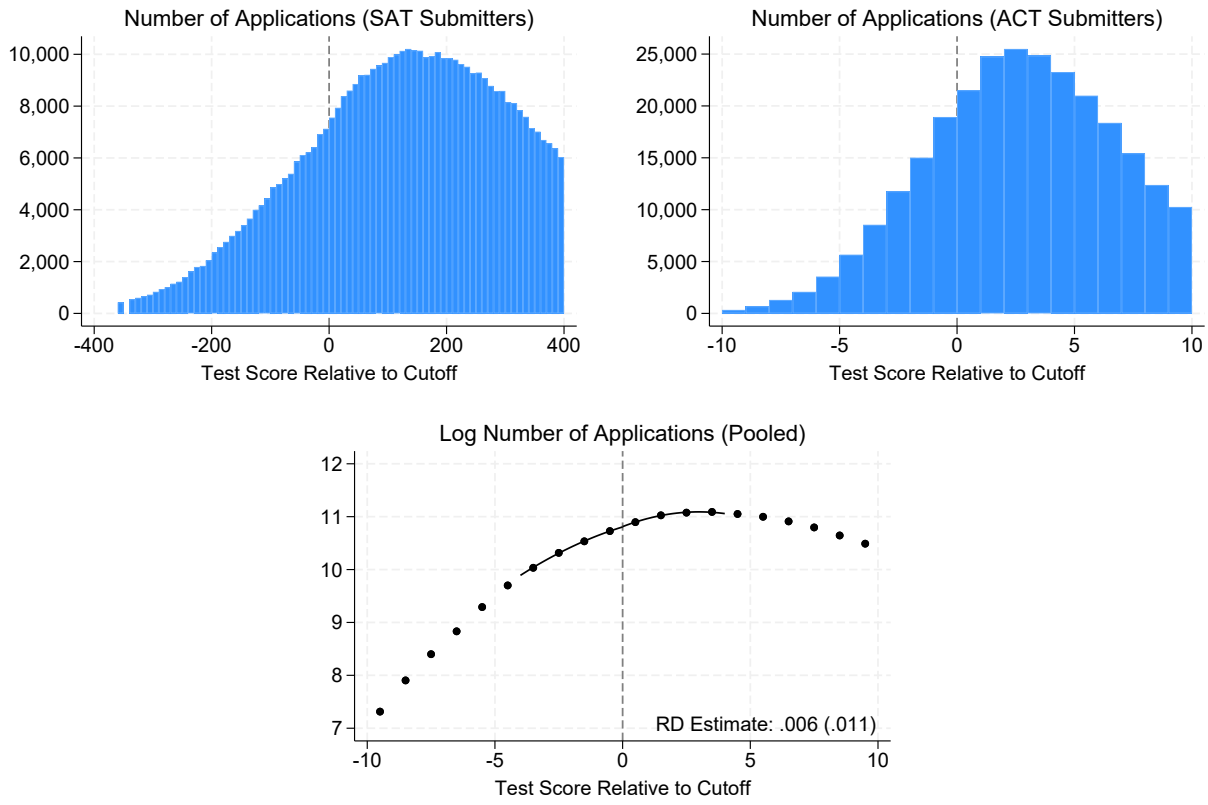

Notes: The top panel plots the nonparametric density of applications at each unique value of the running variable, separately for SAT and ACT submitters. The bottom panel conducts a McCrary (2008) test of manipulation by taking the log number of applications at each value of the concorded running variable 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 5: RD Diagnostics: Covariate Balance

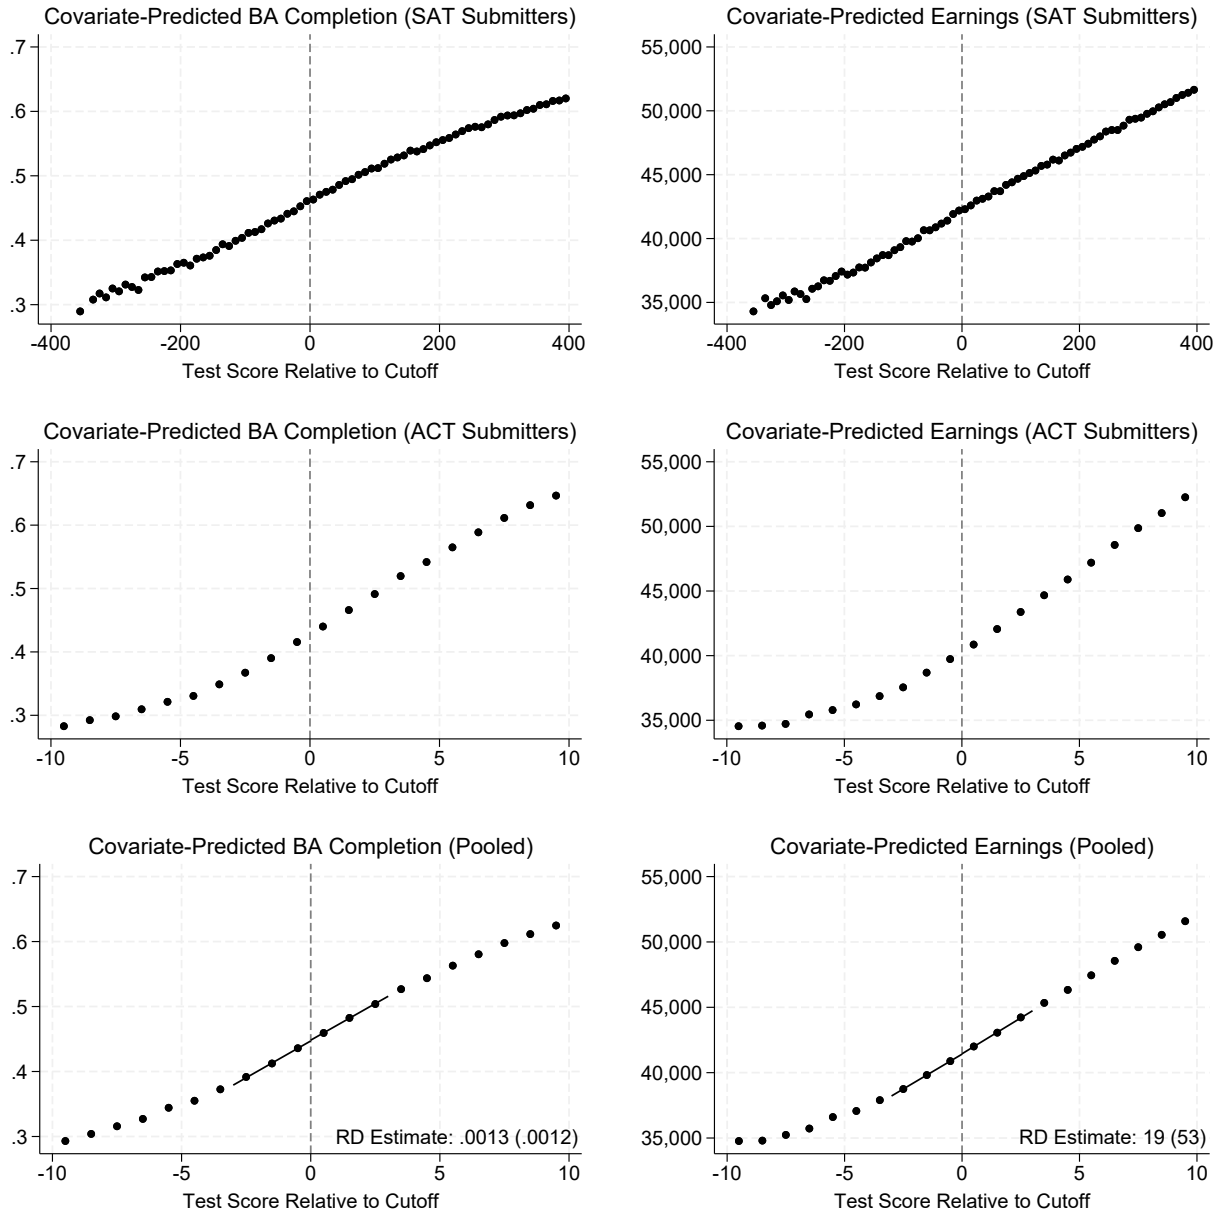

Notes: The left panels plot the covariate-predicted probability of bachelor's degree completion within 10 years of application, predicted via a logit regression using the covariates described in footnote 22. The right panels plot covariate-predicted earnings averaged over 8-12 years after application, 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.

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. $R = c$. 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 $\mathbb{E}[Y_1 - Y_0 | R = c, D_1(c) = 1, D_0(c) = 0]$, 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,

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

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 3 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 4 and 5. The institutional setting is also 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 summarize the returns to enrolling the typical marginal applicant to a public university, I first 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.[23] Cattaneo et al. (2016) show 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

[23] I define the cutoff as the midpoint between the two discrete running variable values that straddle the first stage discontinuity. For SAT submitters, I divide their SAT running variable by 40 to convert it to ACT units.

discontinuity). I then unpack these pooled estimates to explore heterogeneity across institutions (Section 5), across observable student backgrounds (Section 6), and across latent complier types on the intensive versus extensive margins of the four-year sector (Section 7).

I estimate (1) with local linear approximations of $\mathbb{E}[Y|R]$ and $\mathbb{E}[D|R]$, 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 very close to the bandwidth of three concorded ACT points that I use in the main specification, and the results are similar across a 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. Since application cell fixed effects and cell-specific running variable slopes absorb some residual outcome variation and slightly increase precision, I include these in the main specification, but also show robustness to their exclusion. 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 2.7 million graduates from Texas public high schools from 2004-2014, described in Column 1. Roughly one-third of those high school graduates apply to a Texas public university during their senior year of high school (Column 2). Roughly one-fourth 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 (thus pooling 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 bottom right panel of Figure 3, compliers comprise 15 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

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	.56	.55
FRPL	.48	.35	.45	.47	.47
White	.43	.46	.36	.32	.28
Hispanic	.39	.33	.36	.36	.32
Black	.14	.13	.22	.25	.33
Asian	.04	.07	.05	.05	.06
At-Risk	.55	.31	.48	.50	.51
Gifted	.12	.24	.09	.09	.09
HS Math (std.)	.11	.63	.09	.05	.00
HS English (std.)	.12	.57	.21	.18	.13
SAT Score (1600)	-	1039	916	906	892
ACT Score (36)	-	21.6	18.4	18.1	18.0
Applicants	2,721,970	885,070	234,271		

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 16). The means of marginal applicants and enrollment compliers in Columns 4 and 5 are estimated using the method of Abadie (2002) applied to the main RD specification described in Section 2.4.

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 6 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 6, the outcome is enrolling in the target university, so the reduced form discontinuity of 15 percentage points is the size of the enrollment first stage, and the 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 6 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 47 percent of compliers would fall back to a different Texas four-year college, including public and private schools; in the bottom left, 43 percent would fall back to a Texas two-year community college; and in the bottom right, the remaining 9 percent

Figure 6: Impacts on Enrolling in the Target University versus Next-Best Alternatives

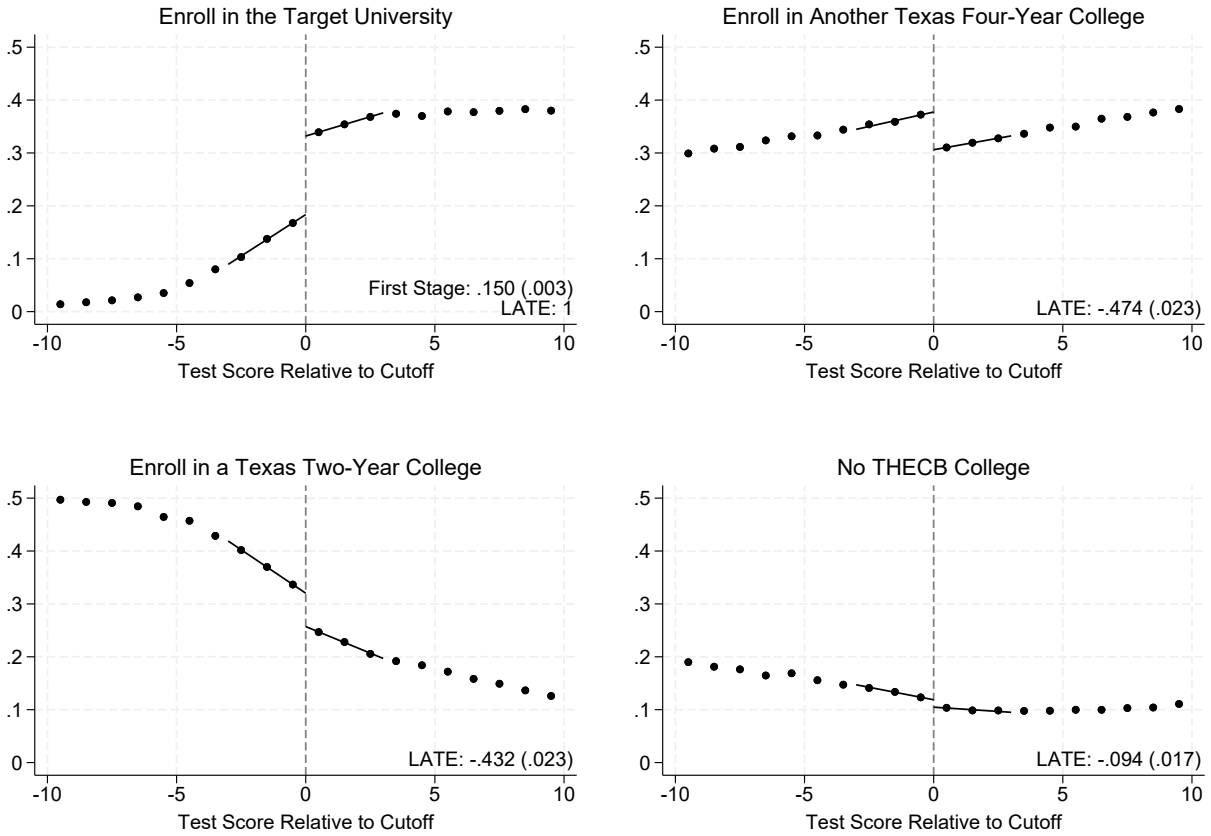

Notes: These sectors correspond to the applicant's first enrollment in the academic year after application. The LATE estimates come from the main fuzzy RD specification described in Section 2.4. The LATE of enrolling in the target university is 1 by construction, since that is the treatment variable. "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 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.

would not enroll in any Texas public or private college covered by the THECB data. 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, as only 4 percent of all untreated compliers attend a college outside of the THECB data universe but recorded in the National Student Clearinghouse. Figure A.3 also shows that only 5 out of the 47 percent of fallbacks to another Texas four-year college are to a private institution (which are included in the THECB data). 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 7 develops a method to identify separate effects for these distinct complier types.

Figure 7 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

Figure 7: Impacts on Peer and Institutional Characteristics

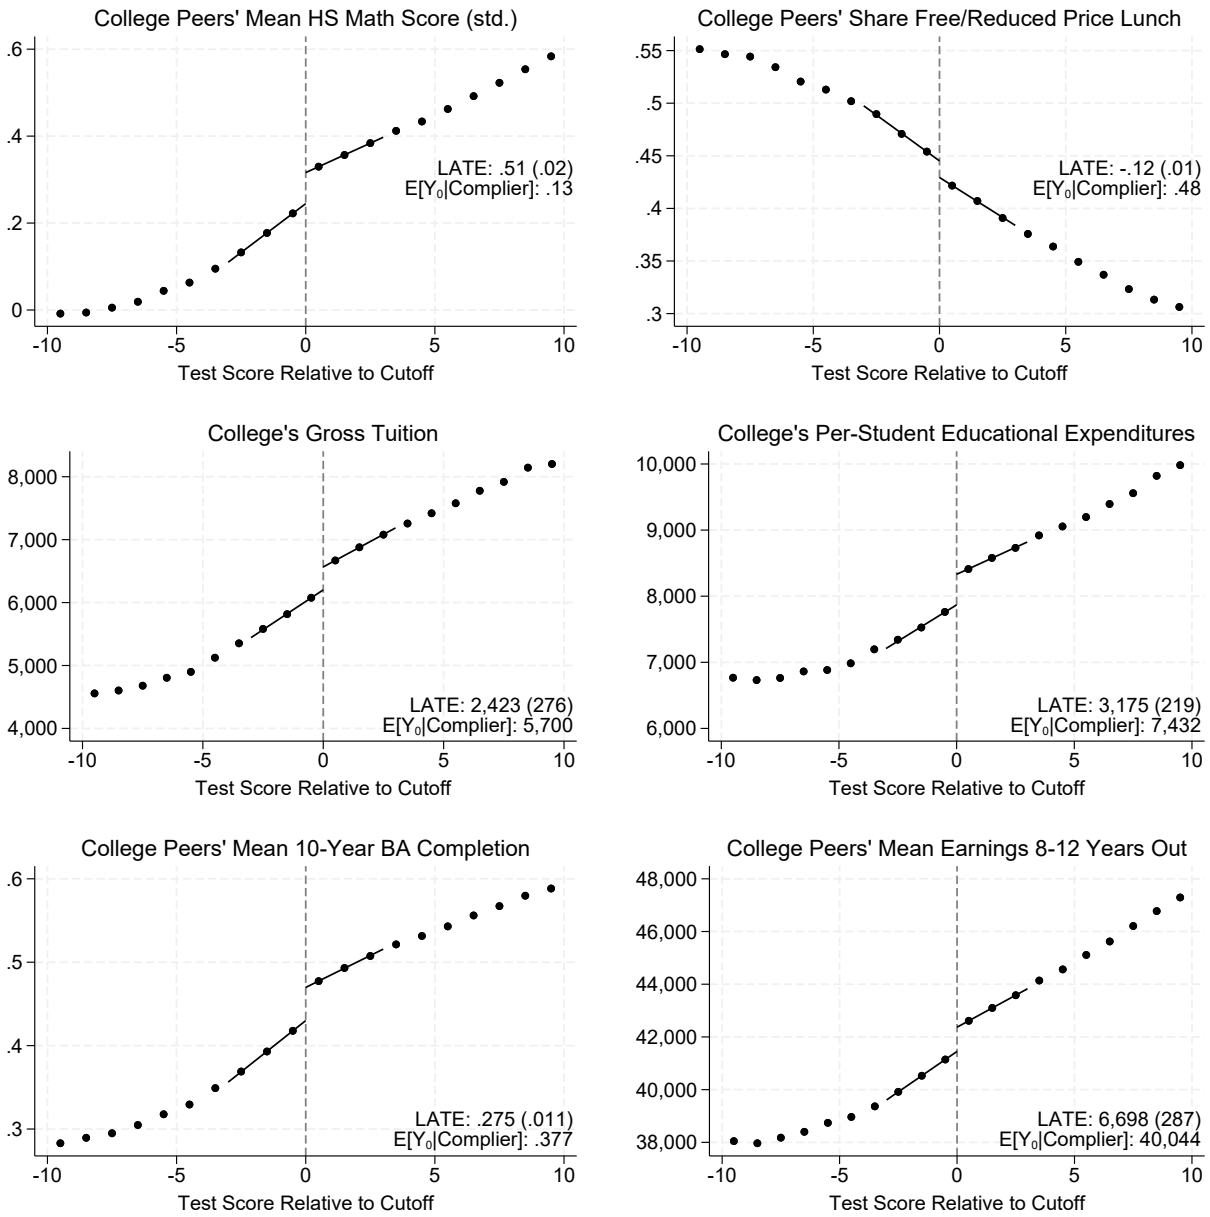

Notes: These college-level characteristics correspond to the applicant's first enrollment in the academic year after application. Applicants who end up enrolling nowhere are included in this analysis by assigning them the mean value of the dependent variable among Texas high school graduates who do not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

their next-best alternative.[24] The top panel of Figure 7 shows that the average high school math score of a complier's college peers increases by half 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 $2,400 more in gross tuition, which is a 42 percent increase relative to the untreated complier mean of $5,700. Those institutions also spend $3,200 more per year educating each student, a 43 percent increase over the untreated complier mean. The bottom panel of Figure 7 turns to average peer outcomes: cutoff compliers experience a dramatic 28 percentage point increase in their peers' 10-year BA completion rate and $6,700 higher peer mean earnings measured 8-12 years after college entry.

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 10 years out from application.

The plots in the left column of Figure 8 show that cutoff-crossing has a decisive influence on compliers' long-term engagement with the target university at which they are marginal.[25] Over the span of ten years from the initial application, the top left plot shows that only 11 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 educational attainment, and shows that cutoff-crossing causes compliers to eventually complete 72 more credits at the target university. This is roughly equivalent to 2.5 years of a four-year degree, which usually 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 8 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

[24] 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.

[25] See Figure A.4 for the reduced form RD plots corresponding to the outcomes in Figure 8.

Figure 8: Impacts on Long-Run Educational Attainment

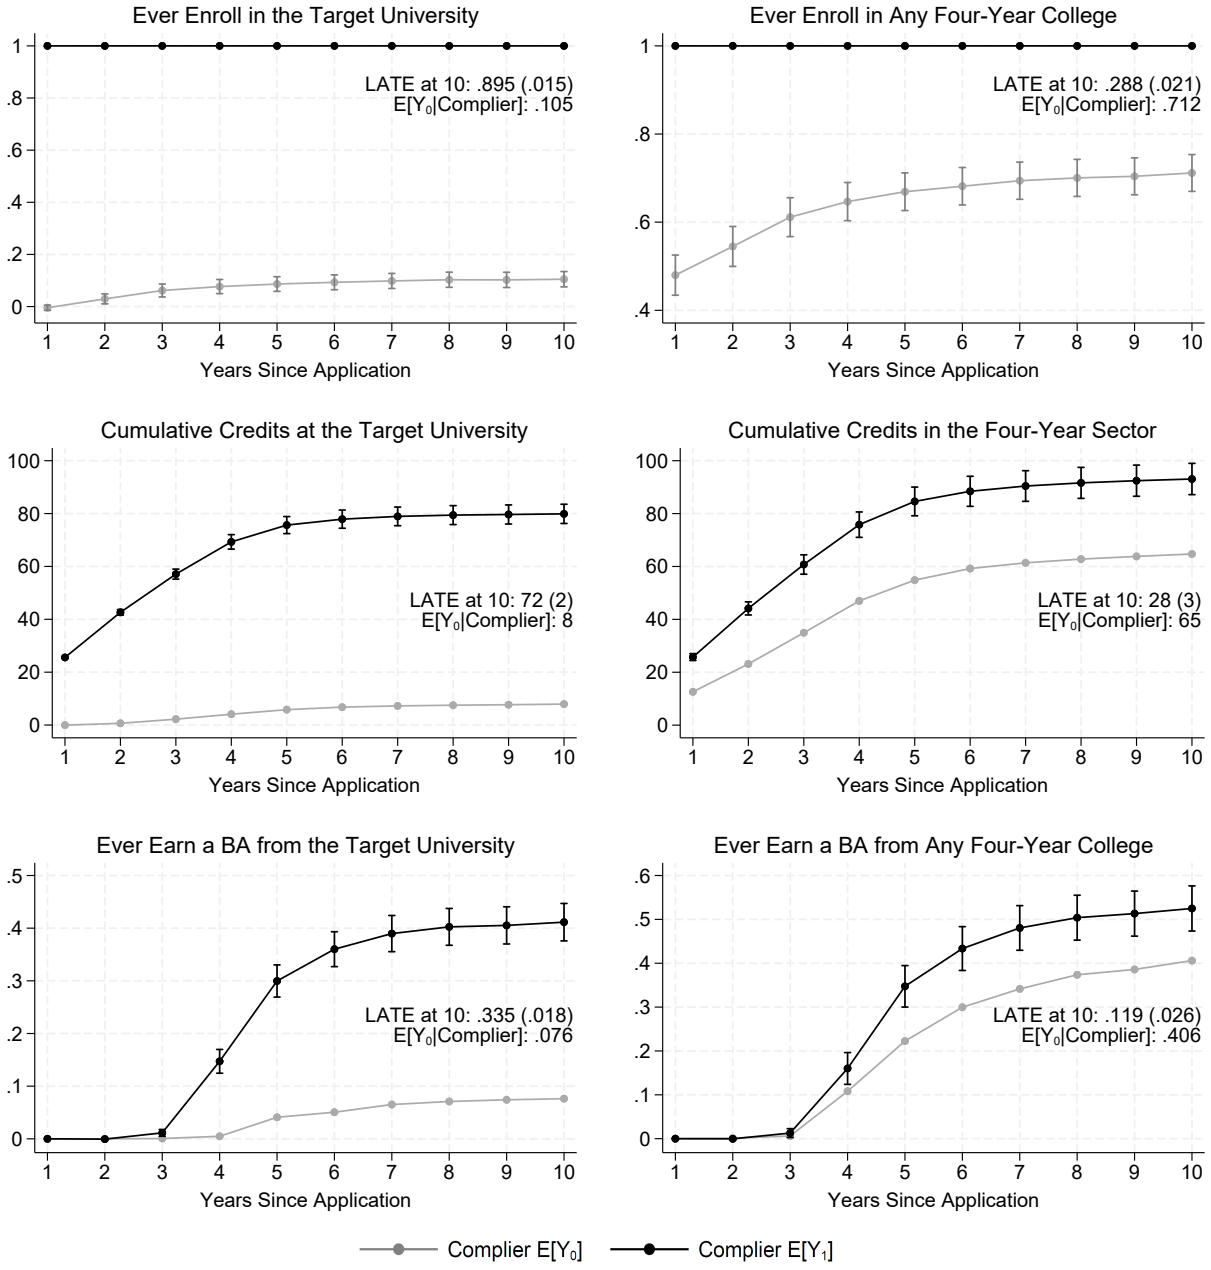

Notes: The gray dots in each plot show the mean outcome of compliers who fall just below the admission cutoff (untreated) at each year since application, and the black dots show the mean outcome of compliers who fall just above the cutoff (treated), which is the untreated mean plus the LATE. The estimates come from the main fuzzy RD specification described in Section 2.4, estimated separately for each year since application. See Figure A.4 for the reduced form RD plot corresponding to each outcome measured at 10 years.

plot shows that cutoff-crossing leads to a substantial 29 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 53 percent of compliers in Figure 6 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 8 shows that compliers become 12 percentage points more likely to ever complete a bachelor's degree from any four-year institution as a result of 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 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 at the target institution increasing that chance substantially but ultimately to a level just above 50 percent.

In terms of majors, the top left panel of Figure 9 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 8 is therefore driven entirely by additional degrees in non-STEM fields, as confirmed in the top right plot of Figure 9.

The previous results focused on educational trajectories in the four-year undergraduate sector; the remaining panels of Figure 9 explore substitution away from two-year community colleges and tests for downstream impacts on graduate education. The middle left panel of Figure 9 shows that cutoff-crossing causes compliers to complete 15 fewer credits in the two-year sector. Thus, roughly half of the additional credits completed in the four-year sector in Figure 8 are cannibalized from the two-year sector, with the other half comprising net gains in total postsecondary attainment. The middle right panel of Figure 9 shows that about 17 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 8 come at the expense of shorter degrees. In terms of graduate school, the bottom two plots of Figure 9 show weakly positive effects on graduate school enrollment and significantly positive effects on graduate degree completion, with compliers becoming about 3 percentage points more likely to hold a graduate degree by 10 years out from their initial undergraduate admission.

Figure 9: Impacts on Educational Attainment: Majors, Two-Year Credits, and Graduate School

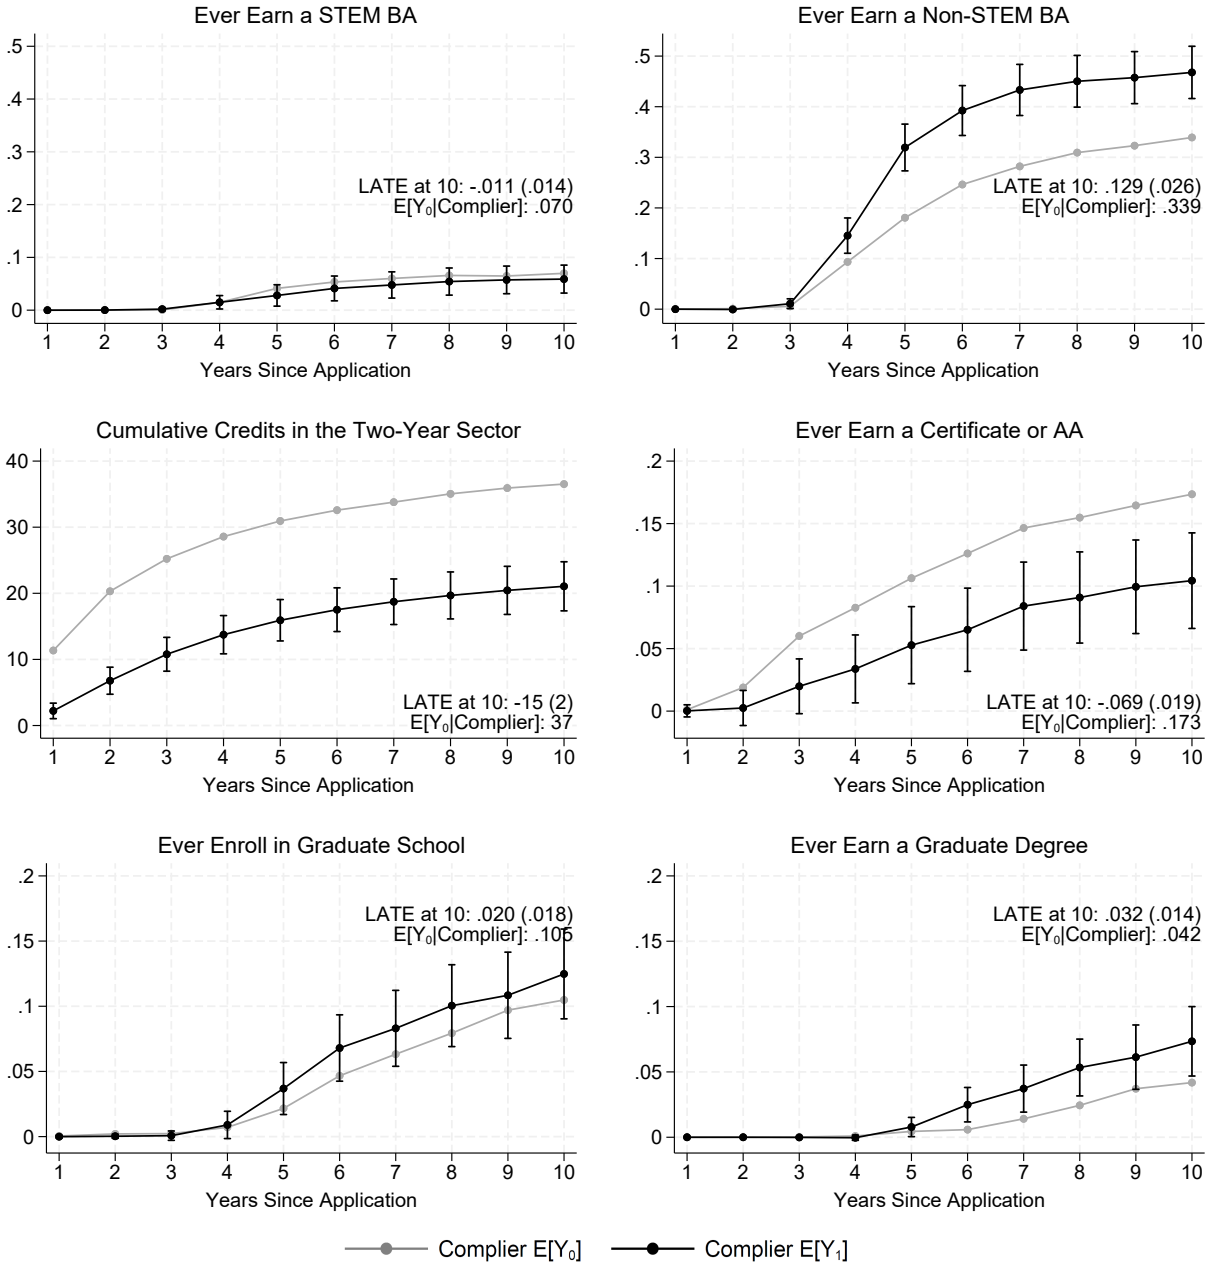

Notes: The gray dots in each plot show the mean outcome of compliers who fall just below the admission cutoff (untreated) at each year since application, and the black dots show the mean outcome of compliers who fall just above the cutoff (treated), which is the untreated mean plus the LATE. The estimates come from the main fuzzy RD specification described in Section 2.4, estimated separately for each year since application. STEM definition comes from U.S. Department of Homeland Security (2016). See Figure A.5 for the reduced form RD plot corresponding to each outcome measured at 10 years.

3.3 Earnings Trajectories

Do the gains in educational attainment generated by cutoff crossing ultimately generate earnings gains for compliers? 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 beyond the first year after application; roughly 75 percent of both appear in the earnings data in any given year after six years out. As is common in state earnings data, the individuals who do not appear in a given year could either be not working or working outside of Texas, making it difficult to determine which missing values are true zeros. Since cutoff crossing does not induce any change along the extensive margin of having positive Texas earnings, I proceed by conditioning on years with positive earnings and studying dynamic effects on three complementary measures: the dollar amount of earnings, log earnings, and an applicant's percentile rank of earnings among all Texas high school graduates who graduated in the same year. In all panels of Figure 10, I use the "sandwich" earnings measure common in analyses of quarterly state earnings data (e.g. Sorkin, 2018; Card et al., 2025), which keeps the quarters with positive earnings that are "sandwiched" between other quarters with positive earnings, thus ignoring the high-variance transition quarters between spells with earnings and spells without earnings. I annualize by multiplying the quarterly average each year by four. As alternatives to sandwich earnings, Figure A.6 shows very similar results when including the transition quarters and thus using all quarters with any positive earnings, as well as alternatively including the missing quarters and assigning them all the value of zero.

All 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 five years after application, as they are much more likely to be actively enrolled in a four-year program. Year six is the crossover age, at which point most of the admitted compliers have finished their college education, entered the workforce full-time, and just started to outearn their rejected counterparts. A statistically and economically substantive earnings premium starts to solidify around eight years out from application and persists thereafter: the middle left panel of Figure 10 shows gradually increasing earnings gains in dollar units, while the middle right and bottom panels show stabilizing relative gains in units of logs and ranks. The solid estimates in Figure 10 denote the balanced panel of outcomes up through year ten that are observed for the entire sample, while the dotted estimates are observed only for progressively older subsets of the sample but included to give a sense of the trajectories that lie beyond the currently available balanced panel.

When pooling across 8-12 years out from application, the LATE on log earnings in the middle right panel of Figure 10 implies that enrolling the typical marginal applicant to a public university yields a gross earnings return of 8.6 percent.[26] This lines up closely with the middle left panel on dollar earnings, as the 8-12 year LATE of \$3,341 divided by the untreated complier mean of \$40,815 yields a ratio of 8.2 percent. In terms of ranks, the bottom panels show that the

[26]Specifically, I run the main 2SLS specification on a long dataset of stacked observations within each individual across years 8-12, interact the usual cell fixed effect and slope controls with indicators for each year, and cluster at the individual level as usual. This stacked estimate is extremely similar to a precision-weighted average of the completely separate estimates for each year 8-12 plotted in Figure 10.

Figure 10: Impacts on Earnings Trajectories

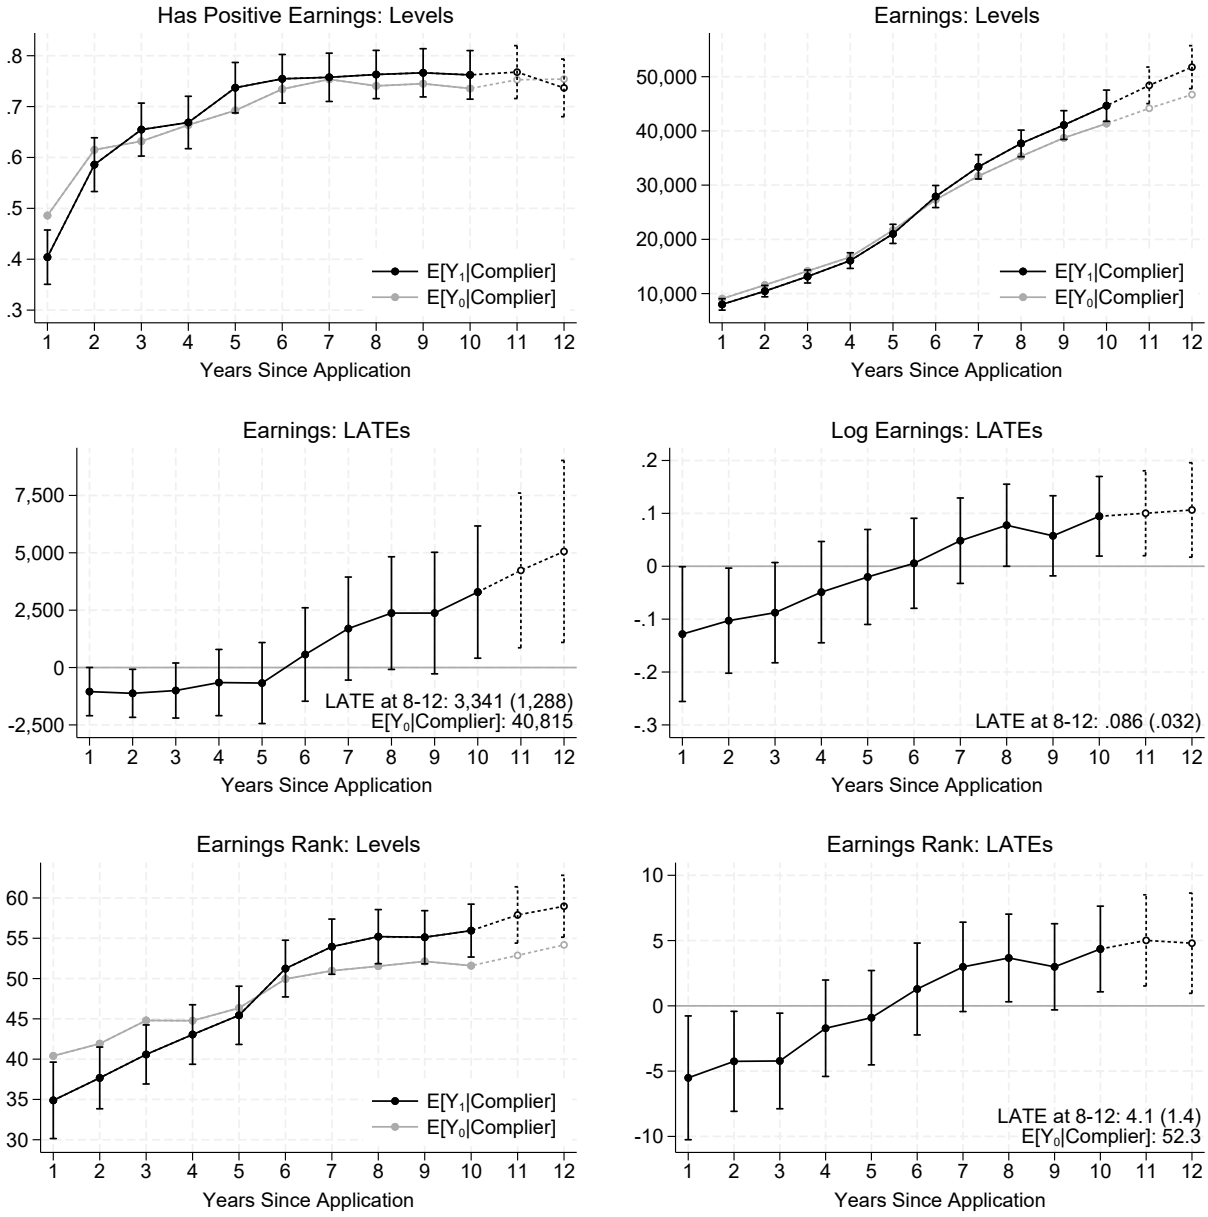

Notes: Each estimate comes from the main fuzzy RD specification described in Section 2.4 estimated separately for each year since application. In the "Levels" plots, the gray dots show the mean outcome of compliers who fall just below the admission cutoff (untreated), while the black dots show the mean outcome of compliers who fall just above the cutoff (treated), which is equal to the untreated mean plus the LATE. In the "LATEs" plots, the black dots show the LATE at each year. The solid dots denote the balanced panel of outcomes up through year 10 that are observed for the entire sample, while the hollow/dotted estimates come from progressively older subsets of the sample. The earnings measure used in these plots is annualized quarterly "sandwich" earnings, which averages the positive earnings quarters within each person-year that are surrounded by other positive earnings quarters and multiplies by four to annualize. See Figure A.6 for analogous results using alternative earnings measures. Earnings ranks correspond to each individual's percentile rank among their statewide high school graduating cohort. Earnings are winsorized at the 99th percentile and measured in real 2015 dollars. Log earnings are also winsorized at the 1st percentile. See Figure 11 for the reduced form RD plot corresponding to each outcome.

Figure 11: Impacts on Earnings: Reduced Form RD Plots

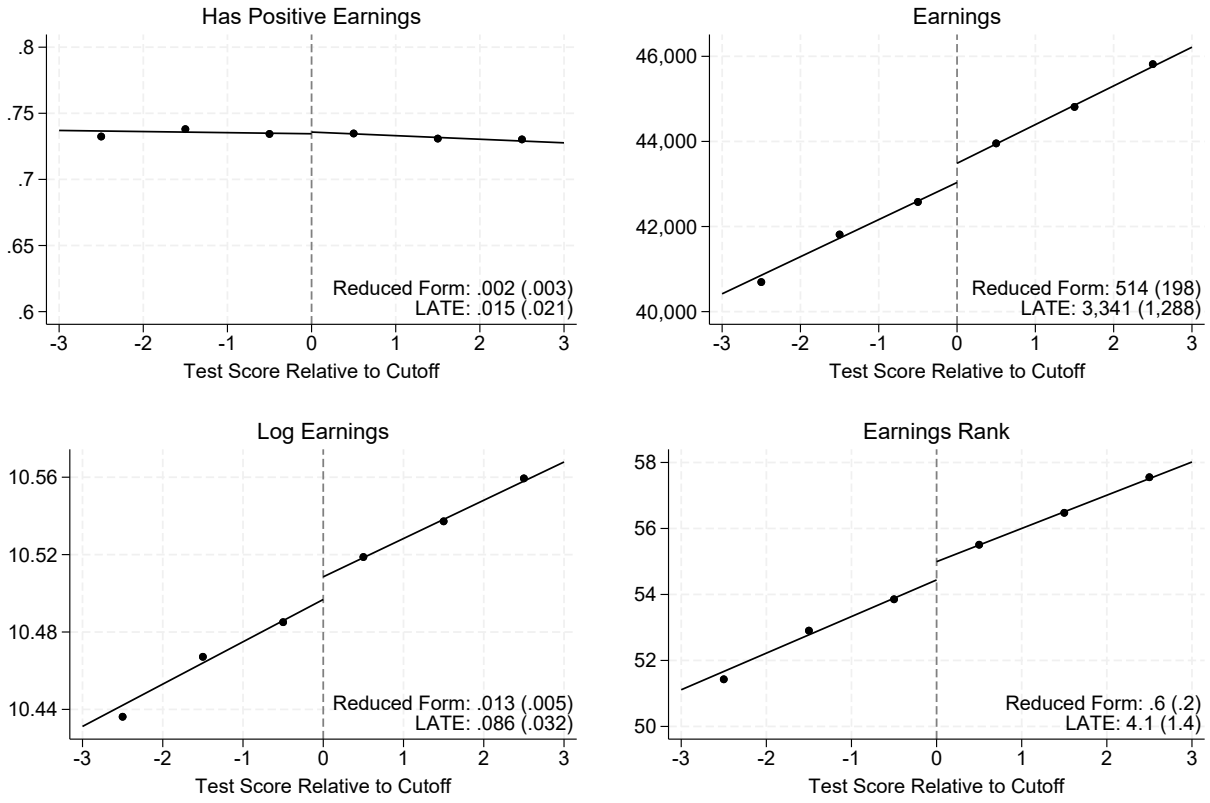

Notes: This figure plots the reduced-form discontinuities in the main earnings outcomes in Figure 10 measured 8-12 years out from application. The estimates come from the main fuzzy RD specification described in Section 2.4.

typical rejected complier ends up right around the 50th percentile of the earnings distribution, with admission boosting that rank by around 4-5 percentiles. Figure 11 shows the reduced form RD plots corresponding to each of these earnings outcomes. Figure A.6 shows very similar results across alternative earnings definitions. Section 3.5 below investigates robustness across a battery of alternative specifications and cutoff definitions.

3.4 Interpreting the Earnings Gains and Comparing to Observational Estimates

A few back-of-the-envelope calculations help to interpret these earnings effects vis-a-vis the gains in educational attainment and compare them to observational estimates. Table 2 collects several of the main RD estimates, many of which are used in this exercise.

First, recall that cutoff compliers earn an additional 28 credits in the four-year sector, which is equivalent to $28/30 = .93$ additional years of education in the four-year sector. If this gain in four-year attainment is the exclusive driver of the \$3,341/\$40,815=8.2% gain in earnings, that would imply a $8.2/.93 = 8.8\%$ gross return to a year of four-year college. Cutoff compliers do forgo half a year of credits (15) in the two-year sector, however, so the implied gross return also depends

Table 2: Causal Impacts of Enrolling Marginal Applicants to Public Universities

	LATE (SE)	$E[Y_0 \mid \text{Complier}]$
Characteristics of the Initial Institution		
Enroll in Any Four-Year College	.526 (.023)	.474
College Peers' Mean 10-Year BA Completion	.275 (.011)	.377
College Peers' Mean Earnings 8-12 Years Out	6,698 (287)	40,044
Long-Run Educational Attainment		
Cumulative Credits in the Four-Year Sector	28 (3)	65
Cumulative Credits in the Two-Year Sector	-15 (2)	37
Ever Earn a BA from Any Four-Year College	.119 (.026)	.406
Ever Earn a Certificate or AA	-.069 (.019)	.173
Earnings		
Has Positive Earnings	.015 (.021)	.745
Earnings	3,341 (1,288)	40,815
Log Earnings	.086 (.032)	1.455
Earnings Rank	4.1 (1.4)	52.3
Applicants for Education Outcomes	234,271	
Applicants for Earnings Outcomes	200,454	

Notes: Each estimate comes from the main fuzzy RD specification described in Section 2.4. Long-run educational attainment is measured 10 years out from the initial application. Earnings outcomes are pooled annual estimates over years 8-12; see Section 3.3 and Figure 10.

on the relative valuation of four-year versus two-year credits in the labor market. Table A.1 offers suggestive evidence from observational regressions of log earnings on credits that conditional on four-year credits, additional two-year credits do not increase earnings. The OLS coefficient on cumulative four-year credits with no controls, estimated in the entire population of high school graduates, suggests a log earnings return to a year of four-year education of 9.5%, and precisely zero return to two-year credits conditional on four-year credits. These estimates would line up fairly closely with the RD estimate's implied return of 8.8% calculated above. Adding detailed covariate controls, however, decreases the OLS return to a year of four-year college to just 5.1% (and still zero return to two-year credits), which is substantially below the RD-implied estimate. Limiting the OLS sample to the RD applicants produces a similarly small OLS estimate of 5.6% per year of four-year credits even with no controls, and now features a negative return to two-year credits conditional on four-year credits. Adding controls further reduces the four-year return. Thus, observational estimates of the gross return to a year of four-year college would roughly equal or substantially underpredict the RD-implied return, depending on the OLS sample and controls.

At the same time, the results on peer mean outcomes in Figure 7, reproduced in Table 2, imply that observational estimates of the return to college "quality" using institutional mean outcomes would substantially overpredict the educational and earnings gains actually reaped by marginally admitted students. Recall that cutoff compliers experience a 28 percentage point gain in the mean

BA completion rate of their initial institution, and a $6,700 gain in institutional (i.e., peer) mean earnings. These numbers correspond to the gains we would predict for compliers based on simple observational comparisons of mean outcomes across institutions, which are widespread in college guides, popular media, research reports, and government statistics. This approach would end up overpredicting the gains actually experienced by cutoff compliers by a factor of two: 28 vs. 12 percentage points for BA completion, and $6,700 vs. $3,300 for earnings. Altogether, these results suggest that observational estimates are an uneven guide to the causal estimates recovered by the RD design: OLS estimates of the return to a year of college can roughly equal or substantially underpredict the estimate implied by the RD results depending on the OLS sample and controls, while the returns to attending a higher quality college implied by simple comparisons of mean institutional outcomes substantially exceed those actually reaped by marginally admitted students.

3.5 Specification Checks

Figures A.8 and A.9 investigate the robustness of the RD estimates across a battery of alternative approaches. First, Figure A.8 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; that bandwidth is always very close to the main specification's bandwidth of three concorded ACT points. The point estimates are similar across the wide range of bandwidths, with expectedly less precision at smaller bandwidths and more precision at larger ones.

Next, Figure A.9 investigates several alternative ways of structuring the regression specification. Specification 1 reproduces the main estimates from the baseline specification described in Section 2.4. Specification 2 clusters the standard errors at the application cell level in addition to the applicant level, yielding very similar confidence intervals. Specification 3 adds the full suite of pre-college covariates described in footnote 22 to the main specification. Specification 4 removes the application cell-specific slopes from the main specification, controlling only for cell fixed effects and common slopes above and below the cutoff. Specification 5 further removes the application cell fixed effects, leaving only a common intercept and common slopes above and below the cutoff. Specification 6 changes the local polynomial functional form from linear to quadratic in the running variable on either side of the cutoff, first with a bandwidth of 6 concorded ACT points. Specification 7 narrows the quadratic specification bandwidth to 5 ACT points, and Specification 8 further narrows it to 4 ACT points.

Finally, the last set of estimates in Figure A.9 investigate several alternative ways of inferring the relevant cutoffs from the admission data. Specification 9 includes rather than disqualifies the roughly 20% of initially inferred cutoffs that exhibit statistically significant discontinuities in the log density of applicants or (more rarely) covariate-predicted BA completion or earnings. Specification 10 restricts the main set of cutoffs to those with statistically significant (at the 1% level) jumps in admission and enrollment, and Specification 11 does the same using a 5% significance threshold.

Finally, Specification 12 allows each application cell to potentially contribute multiple cutoffs, keeping all distinct cutoffs within a cell with statistically significant (at the 5% level) jumps in admission and enrollment. Altogether, Figure A.9 shows that the results remain similar across this entire battery of alternative approaches.

4 Cost-Benefit Analysis

The preceding results show that marginally admitted public university students eventually reap positive gross earnings returns. Those gross returns take many years to materialize, however, and do not account for the private and social costs of the additional education. In this section, I first use the main RD specification to estimate the cost side of the ledger. I then combine the dynamic effects on costs and earnings to conduct a formal cost-benefit analysis that quantifies the net returns to enrolling marginal students from the perspectives of the students themselves, taxpayers, and society.

4.1 Causal Impacts on Costs

Figure 12 presents estimates of the cumulative private and social costs generated by the average admission cutoff complier using the main fuzzy RD specification.[27] The top left panel begins by showing that cutoff compliers ultimately incur around $4,600 in additional gross tuition charges, ignoring financial aid, accumulated across all full and partial semesters enrolled at four-year and two-year colleges.[28] Many students receive grant aid to offset these gross tuition charges, however, and the top right panel shows that the additional accumulation of grants nearly 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 $5,300 in additional student loans, as shown in the middle right panel, likely to finance additional room and board charges and other consumption during college. Indeed, the bottom left panel shows that cutoff compliers would generate around $7,600 in additional room and board charges in the four-year sector if they lived on campus.[29]

[27] As noted in Section 2.1, the availability of the cost data lags one year behind the other outcome datasets, so the cost outcomes are observable for a balanced panel of nine years (instead of ten) for all cohorts.

[28] To construct this gross tuition measure, for each student enrolled at a given Texas institution in a given semester, I use the IPEDS panel data to first measure the gross tuition price of full-time enrollment at that institution in that semester, and then I prorate that measure by the individual student's credit enrollment intensity that semester relative to full-time status. Gross tuition includes mandatory fees and the college's estimated cost of required books and supplies, but excludes room and board. All students are assumed to pay in-state tuition rates, and community college students are assumed to pay in-district rates.

[29] Since I do not observe individual-level data on room and board charges, I use the IPEDS panel to measure what each student would pay each semester for room and board if they lived on campus, given the institution enrolled in and the year enrolled, prorated by the student's credit enrollment intensity relative to full-time. On one end, if all semesters enrolled at four-year institutions generate room and board charges, but no semesters enrolled at two-year institutions do (e.g. if all two-year students live at home with their parents with no opportunity cost), then the typical marginally admitted student would accumulate the additional room and board charges shown in the bottom

Figure 12: Impacts on Cumulative Costs

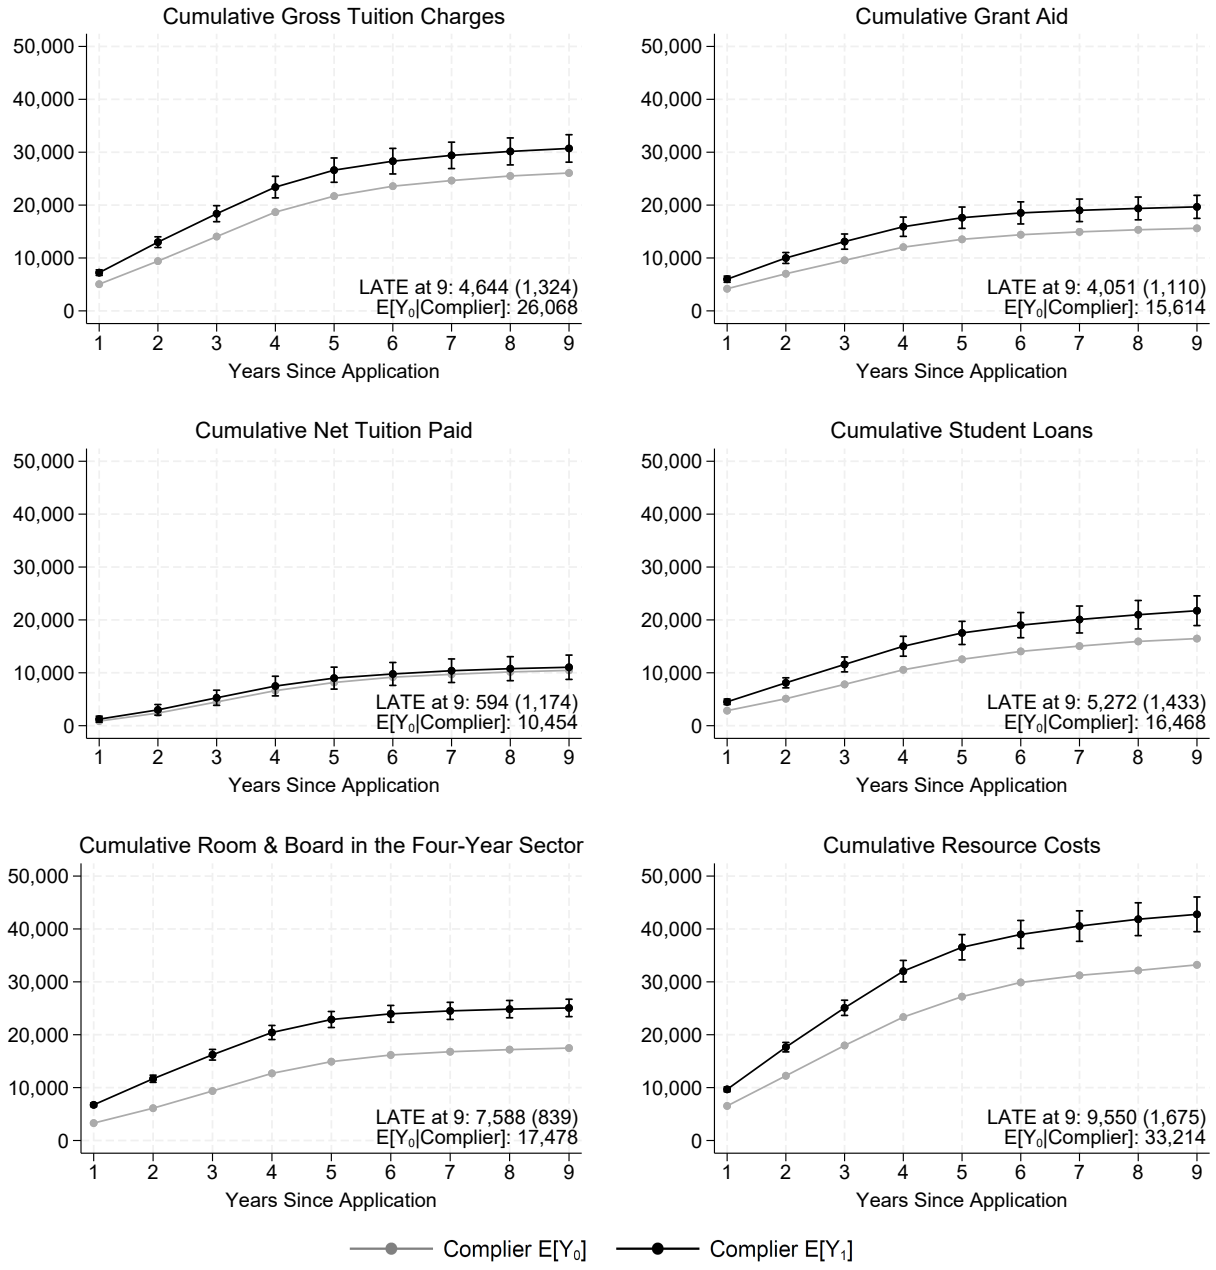

Notes: The gray dots in each plot show the mean outcome of compliers who fall just below the admission cutoff (untreated) at each year since application, and the black dots show the mean outcome of compliers who fall just above the cutoff (treated), which is the untreated mean plus the LATE. The estimates come from the main fuzzy RD specification described in Section 2.4, estimated separately for each year since application. All costs are deflated to real 2015 dollars. The gross tuition, room & board, 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 Figure A.10 for the reduced form RD plot corresponding to each outcome measured at 9 years.

From society's perspective, the additional resources used to educate these marginally admitted students could have been used for other purposes. To measure these social resource costs, I follow Hoxby (2019) and use the IPEDS panel data to measure the per-pupil "core" educational expenditures of each institution each year, which include spending on instruction, academic support, and student services. I then pro-rate that cost by the individual student's credit enrollment intensity each semester relative to full-time status. The bottom right panel of Figure 12 shows that the average cutoff complier ultimately adds around $9,600 of such resource costs to society's ledger.[30] Such 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 inherent difficulty of measuring marginal costs.

4.2 Cost-Benefit Calculations

I now combine the dynamic causal effects on costs and earnings to calculate the net returns to enrolling marginal public university students from the perspectives of the marginal students themselves, taxpayers, and society.

4.2.1 Undiscounted Cumulative Benefits Versus Costs

The left three panels of Figure 13 begin by plotting the dynamics of the undiscounted cumulative costs and benefits of enrolling the average cutoff complier to study the time horizon at which the investment eventually "pays off" from each perspective. The top left panel takes society's perspective: here, the cumulative gross benefit is simply the running sum of the annual pre-tax earnings effects in Figure 10.[31,32] The cumulative social cost is the accumulation of the resource costs of educating the marginal student documented in the bottom right panel of Figure 12.[33] Together, the dynamics of the undiscounted accumulations show that, from society's perspective, it takes 11 years for the benefits of enrolling the typical marginal student to surpass the costs.

The middle left panel of Figure 13 repeats this exercise from the perspective of the typical marginal student herself. From her private perspective, the gross benefit is her post-tax earnings gain, which is simply the pre-tax social earnings gain in the top left panel minus a 20 percent tax and transfer rate on the incremental earnings, following Hendren and Sprung-Keyser (2020)

left panel of Figure 12. On the other end, if I also include the room and board charges that would be incurred by all two-year students paying their college's official estimate of local living costs, marginally admitted students would accumulate around $5,000 in additional room and board charges on net, with the $7,600 gross increase in charges from four-year colleges partially offset by roughly $2,500 less in room and board charges from two-year colleges.

[30] This net increase is comprised of a roughly $12,000 gross increase in the four-year sector offset by a roughly $3,000 decrease in the two-year sector.

[31] The calculations in Figure 13 use the main "sandwich" earnings measure from Figure 10; Figure A.11 shows similar results when using the Texas earnings measure in the bottom panels of Figure A.6 that imputes all missing earnings quarters as zeros.

[32] 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).

[33] Informed by the dynamics in Figure 12, I assume that all of the cumulative cost effects stabilize after nine years.

and Angrist et al. (2022). Her cumulative private cost is the roughly zero additional net tuition estimated in Figure 12.[34,35] The dynamics show that for the typical marginal student, it takes around 9 years for her post-tax earnings benefit to surpass her (roughly zero) private cost.

The bottom left panel of Figure 13 repeats this exercise from the perspective of the taxpayer or government budget. From this perspective, the gross benefit is the 20 percent tax on the earnings gains. The taxpayer cost is almost the same as the social cost: since the student is paying almost none of the additional educational resource costs herself, taxpayers must subsidize those costs almost entirely. Some extrapolation of the earnings effects is required to estimate the eventual crossing point of the taxpayer's benefts and costs; the assumptions of the extrapolation are discussed in more detail below when calculating present discounted values. Under those assumptions, the bottom left panel of Figure 13 shows that it takes 19 years for taxpayers' benefits to exceed their costs.

4.2.2 Net Present Values and Internal Rates of Return

The right three panels of Figure 13 add a lifecycle horizon and 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, with the vertical distance between the benefit curve and cost curve defining the net present value (NPV) and their intersection point defining the internal rate of return (IRR) at which the discounted benefits just equal the discounted costs.

Calculating present values of the benefits stream requires assumptions about the future trajectory of the earnings effects beyond the observable range of 12 years out from application. My baseline approach follows Hendren and Sprung-Keyser (2020) and proceeds in four steps, visualized in Figure 14. First, I use the 2015 American Community Survey (ACS) to measure the earnings of the average American at each age from 18 to 65, which is the assumed retirement age for this analysis. Second, as in Hendren and Sprung-Keyser (2020), I convert this cross-sectional age-earnings profile into a plausible future lifecycle profile for my sample cohorts by assuming a constant earnings growth rate of 0.5% per year, multiplying the earnings at each age by 1.005^t, where $t = 1$ at age 18 and $t = 47$ at age 65. Third, I match the level of this ACS lifecycle profile to the observed early-career profile of my untreated compliers by taking the mean untreated complier earnings level at 12 years out (age 30), dividing by the mean earnings at age 30 in the ACS profile, and then multiplying the entire ACS lifecycle profile by this scaling factor. This yields the ex-

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

[35] The costs in Figure 13 exclude room and board charges, given that students 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. In Figure A.12, I show how the cost-benefit calculations would change with the inclusion of the additional room and board costs estimated in Figure 12, which are likely an upper bound on marginal students' actual increased room and board expenditures.

Figure 13: Cost-Benefit Calculations

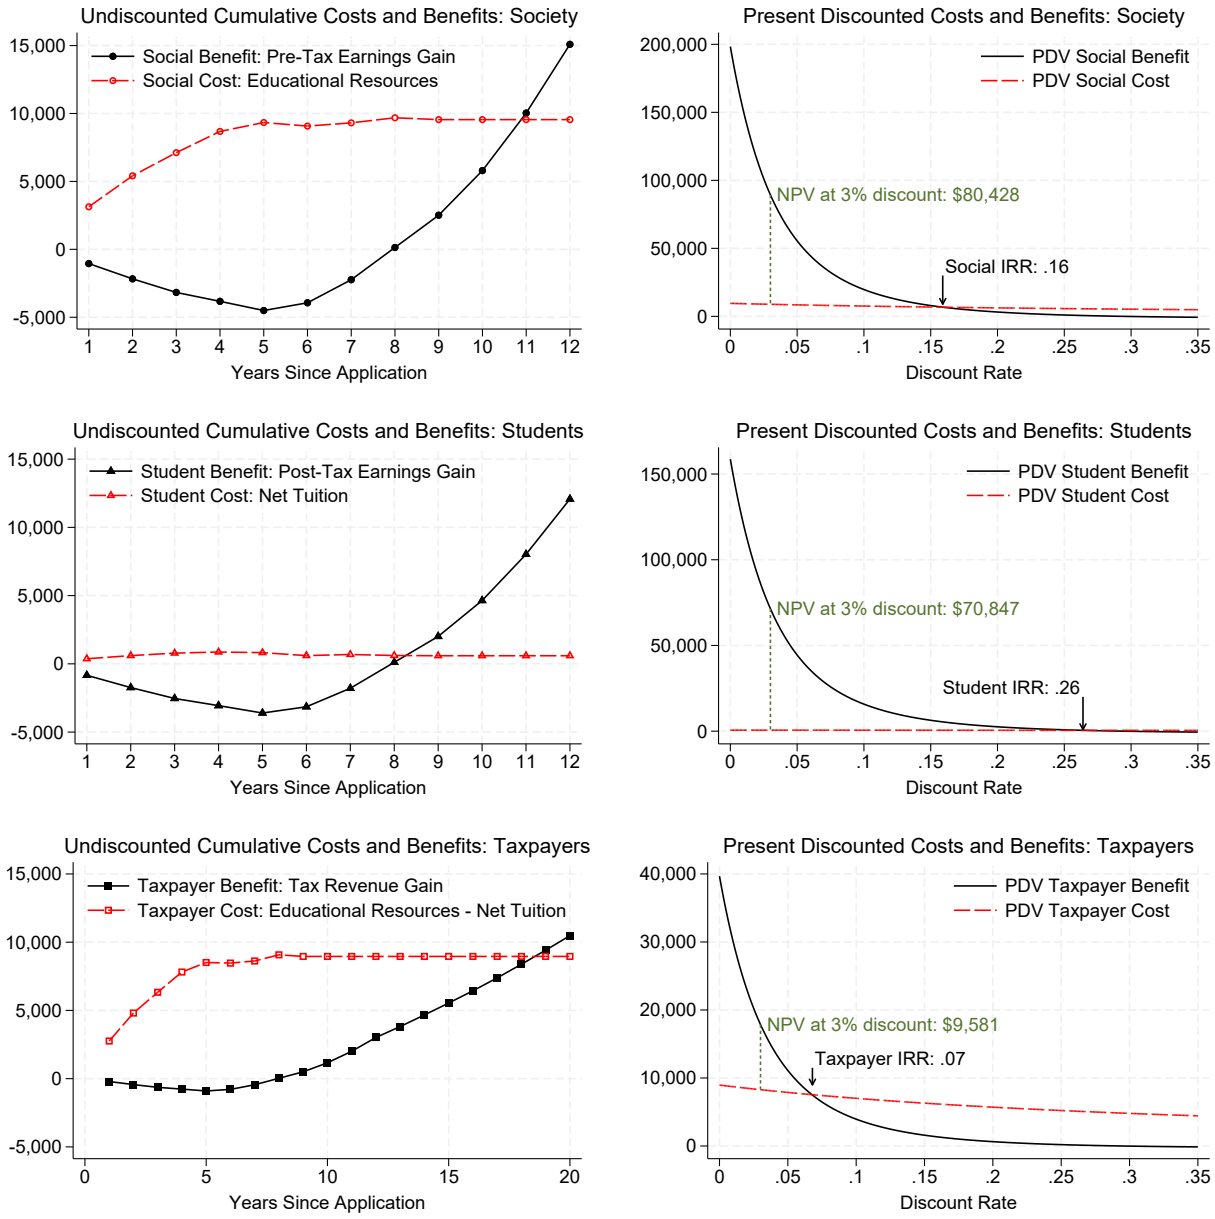

Notes: The left three panels plot the dynamics of the undiscounted cumulative costs and benefits of enrolling the average cutoff complier to study the time horizons at which the investment eventually "pays off" from each perspective; see Section 4.2.1 for details. The right three panels add a lifecycle horizon (see Figure 14) and discounting to the analysis to calculate net present values and internal rates of return; see Section 4.2.2 for details. The top two panels take the perspective of society; the middle two panels take the perspective of the marginally admitted students; and the bottom two panels take the perspective of taxpayers or the government budget. See Figures A.11, A.12, and A.13 for analogous exercises under alternative assumptions.

Figure 14: Extrapolating the Earnings Effects Using the American Community Survey

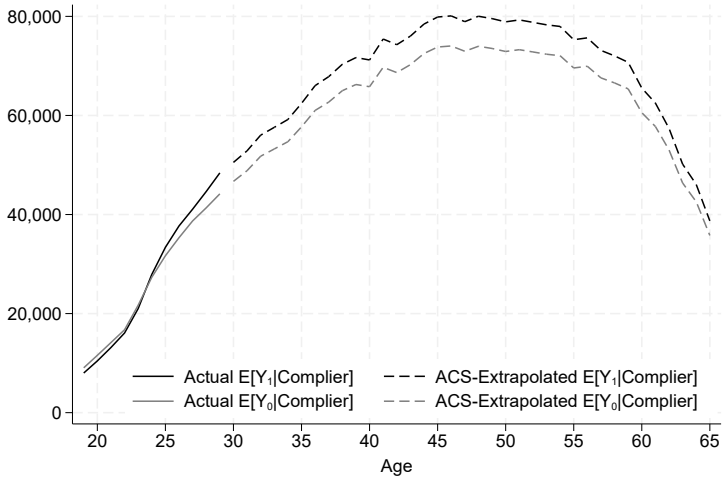

Notes: This plot visualizes the observed and extrapolated lifecycle earnings profiles of treated and untreated compliers using the 2015 American Community Survey. See Section 4.2.2 for the details of the extrapolation.

trapolated lifecycle earnings profile of untreated compliers, shown in the gray dotted line in Figure 14. Finally, I assume that the proportional earnings gain from treatment will remain at 8.2% (as implied by the dollar earnings estimate over years 8-12 in Figure 10) from age 30 until age 65, and thus multiply the untreated earnings profile by 1.082 to yield the lifecycle earnings profile of the treated compliers, shown in the black dotted line in Figure 14.

From society's perspective, the lifetime cumulative earnings benefit of enrolling the marginal applicant is the present discounted sum of all the vertical differences between the treated and untreated complier earnings profiles from Figure 14 over ages 18 to 65. The top right panel of Figure 13 shows that with no discounting, this lifetime gross pre-tax earnings benefit is about $200,000, and subtracting off the roughly $10,000 of social costs of educating the student yields a net present value (NPV) of about $190,000. Since the benefits take several years to materialize but the costs are upfront, the NPV declines quickly as the discount rate increases, e.g. to $80,000 at a 3% discount rate. At a discount rate of 16%, the present value of the benefits would just equal the costs to yield a zero NPV, which implies that society's internal rate of return (IRR) from investing in the typical marginally admitted public university student is 16%.

The middle right panel of Figure 13 repeats this exercise from the perspective of the marginal student herself, replacing the pre-tax social earnings benefit with her post-tax private earnings benefit and replacing the social cost of educating her with her private net tuition bill of roughly zero. With no discounting, the student's lifetime gross post-tax earnings benefit is about $160,000, which is also very close to her net benefit given her roughly zero cost. Her NPV declines to about $70,000 at a 3% discount rate. The NPV does not reach zero until a discount rate of 26%, implying a large private IRR to the student of 26%. Figure A.12 shows that this IRR would decline to a still substantial 15% if the $7,600 in additional room and board charges from Figure 12 were included

in the student cost, which (as discussed above) is likely an upper bound on the student's actual increase in room and board expenditures and may be partly consumption rather than investment.

Finally, the bottom right panel of Figure 13 takes the perspective of the taxpayer or government budget. With no discounting, the lifetime gross benefit of additional tax revenue is about $40,000, equal to the social pre-tax gross benefit of $200,000 minus the student's private post-tax benefit of $160,000. The lifetime net benefit, subtracting off the taxpayer's roughly $10,000 cost of subsidizing the student's education, is thus about $30,000 with no discounting. This NPV declines to about $10,000 at a 3% discount rate, and to zero at 6%. This positive IRR of 6% for taxpayers implies that for all discount rates below 6%, the marginal value of public funds (e.g., Mayshar, 1990; Hendren, 2016) is infinite: subsidizing the education of the marginal student pays for itself via increased future tax revenues.

Thus, even after accounting for the upfront costs and delayed benefits, enrolling marginal applicants to public universities generates substantial net returns for society, the marginal students themselves, and the government budget. These conclusions are robust across a range of alternative assumptions. First, Figure A.11 shows that the results are very similar (with slightly higher IRRs) if I replace the main "sandwich" earnings measure from Figure 10 with the Texas earnings measure from the bottom panels of Figure A.6 that imputes all missing earnings quarters as zero earnings. Second, as noted above, Figure A.12 shows that the student IRR would decline to a still substantial 15%, and the social IRR to 12%, if the upper bound estimate of $7,600 in additional room and board charges from Figure 12 were included in the student and social costs. Finally, Figure A.13 shows that the IRRs decrease only slightly if I dispense with the lifecycle extrapolation and more conservatively assume that the annual earnings gain stays constant at $3,341 (as estimated in Figure 10) from year 13 until retirement, rather than allowing the dollar amount of the earnings gain to rise proportionally over the lifecycle.

5 Heterogeneity Across Institutions

The next three sections unpack the impacts of public universities on their marginally admitted students across several important dimensions of heterogeneity. In this section, I explore heterogeneity across the wide diversity of admitting institutions in the sample, as visualized in Figure 2.

5.1 Across Institutional Selectivity

In the top panels of Figure 15, I begin by describing how the changes in peer quality experienced by cutoff compliers differ across the selectivity of the target university, i.e. the university at which the complier is marginally admitted. The darker "Pooled" estimate in the center of each plot reproduces the overall LATE, while the remaining lighter estimates are separate LATEs for each target university, estimated using the main RD specification in Section 2.4 separately for each target university in the sample. The LATEs are plotted against the selectivity of the institution on the horizontal axis, measured by the average SAT score of incoming students, and the diagonal

Figure 15: Across Institutional Selectivity: Effects on Peer Quality and Educational Attainment

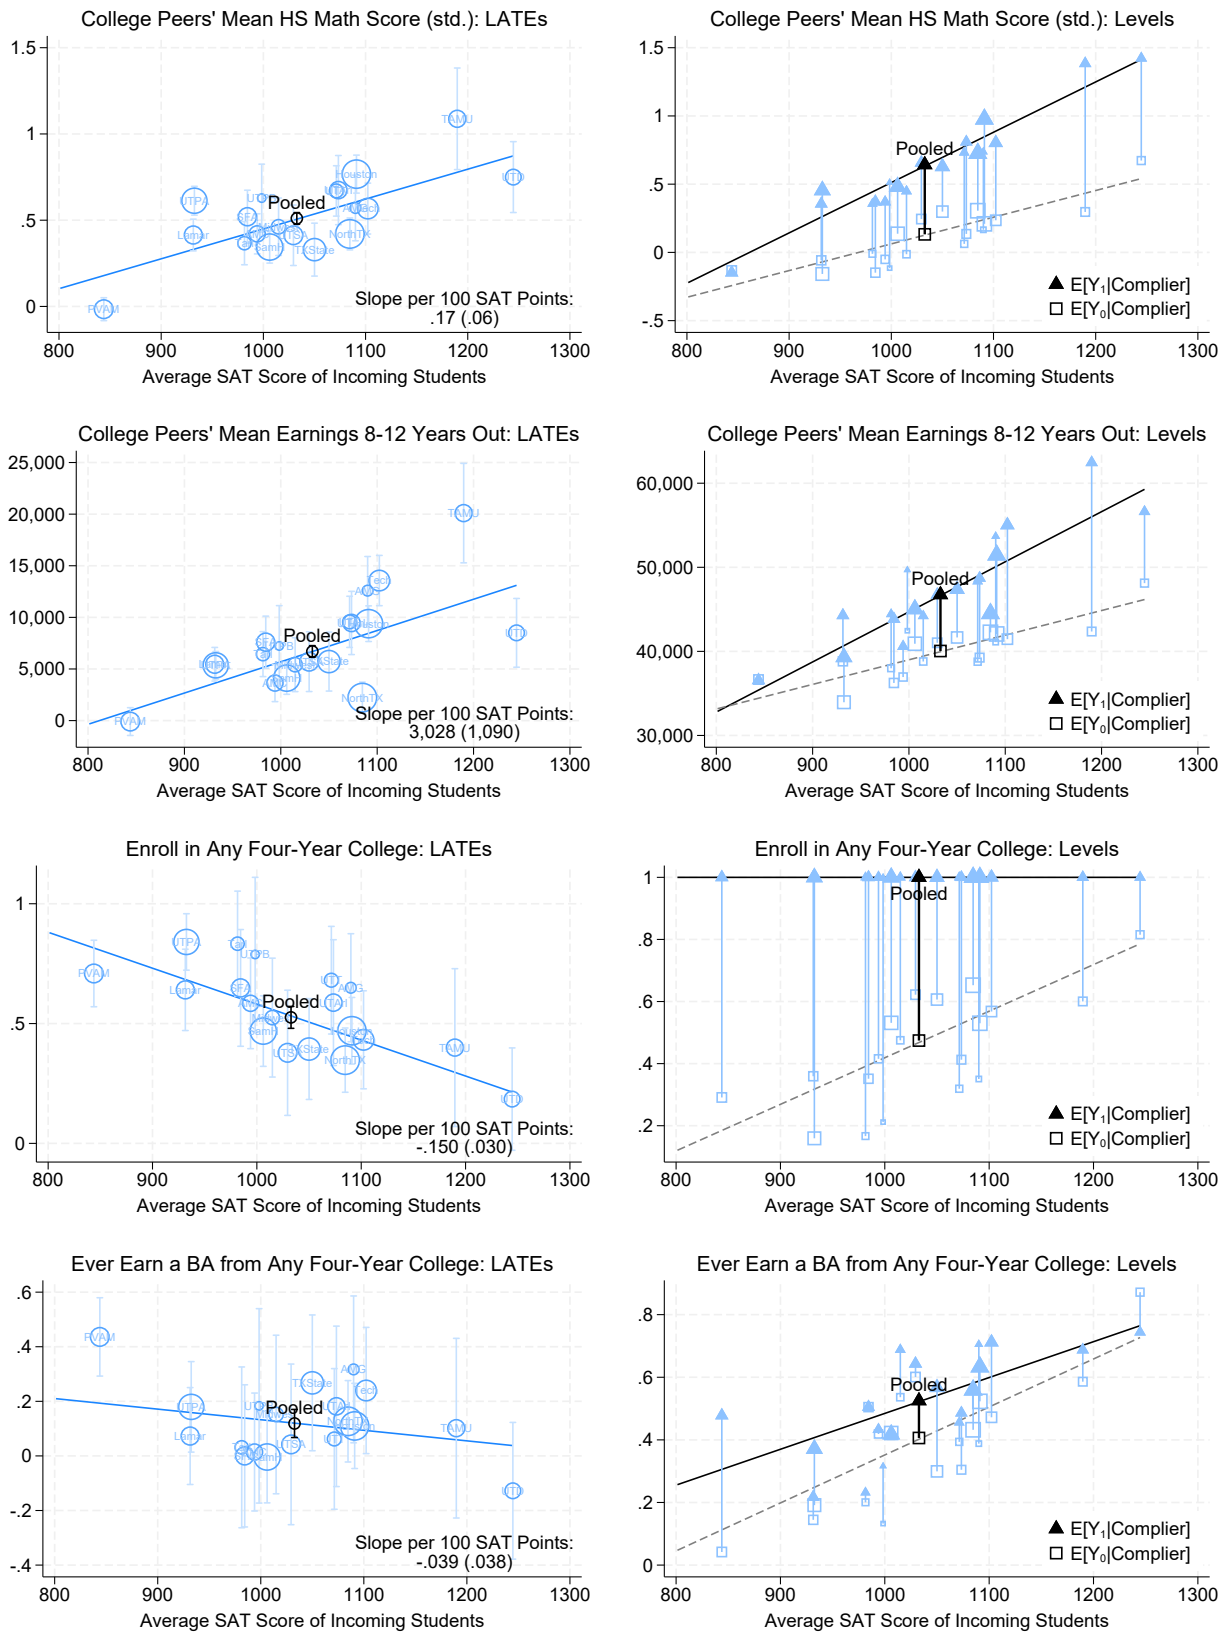

Notes: The darker "Pooled" estimate reproduces the pooled LATEs from Section 3, while the lighter estimates are LATEs for each target university estimated separately using the main specification in Section 2.4. The diagonal line is the estimated linear relationship between an institution's LATE (left plots) or potential outcome levels (right plots) and its selectivity. Estimates for the smallest institutions are not plotted due to their large CIs, but those institutions are included in the pooled estimates and slopes.

line is the estimated linear relationship between the target institution's LATE and its selectivity.[36]

The top left panel of Figure 15 shows that cutoff compliers who are on the admission margin at more selective institutions experience even larger increases in peer test scores than cutoff compliers who are on the admission margin at less selective institutions.[37] The panel below it shows a very similar pattern with respect to peers' eventual earnings, with the estimated slope coefficient predicting that the gain in peer earnings increases by about $3,000 for every 100 SAT point increase in the selectivity of the target institution. The right panels of Figure 15 decompose the LATEs into the treated and untreated complier potential outcome levels, showing that compliers targeting more selective institutions would have had higher peer quality even if rejected (the upward-sloping dashed line), but that peer quality increases even more steeply with selectivity when admitted (the steeper upward-sloping solid line).

At the same time, the bottom panels of Figure 15 show that cutoff compliers at more selective institutions are less likely to be on the extensive margin of the four-year sector, and they experience smaller (though not statistically different) gains in bachelor's degree completion, relative to compliers at less selective institutions. To take one contrasting pair of institutions, marginal admits into Texas A&M (TAMU) experience very large gains in peer quality (over 1 SD more in peer test scores and about $20,000 more in peer earnings), but 60% of them would have fallen back to another four-year college if rejected (compared to 50% across all compliers), and their estimated gain in BA completion is slightly smaller than the average complier (though not statistically different). On the other end, marginal admits into Prairie View A&M (PVAM), a historically Black university, experience no gain in peer test scores or earnings, but become 70 percentage points more likely to start at a four-year college and about 40 percentage points more likely to ever earn a BA.

These potentially offsetting effects of larger gains in peer quality but smaller gains in the quantity of four-year college engagement lead to a roughly flat relationship between a target institution's selectivity and its effect on complier earnings, as shown in Figure 16. The top panels show that cutoff compliers at more selective institutions experience weakly larger earnings gains in dollar levels, but the slope is not distinguishable from zero. Moreover, the middle and bottom panels show flat slopes when measuring earnings in logs and ranks, respectively, suggesting that the weakly increasing slope in dollar units results from constant proportional gains applied to the higher baseline earnings levels of compliers at more selective schools. Indeed, the potential outcome levels in the right panels of Figure 16 show that compilers at more selective schools do earn substantially more than compilers at less selective schools when admitted, but they also would have earned more in the rejected counterfactual.

These results highlight the fact that marginal applicants at more selective schools have underlying characteristics that lead to much higher earnings compared to marginal applicants at less

[36]The institution-specific estimates plotted in the figures in this section exclude the smallest target universities, as the lengths of their confidence intervals exceed the entire vertical range of the plots. But those institutions are included in each pooled estimate, and in the estimation of the linear slopes across institutional characteristics.

[37]Recall that applicants who do not end up enrolling in college are included in these results by assigning them the mean value of the dependent variable among Texas high school graduates who do not enroll in college.

Figure 16: Across Institutional Selectivity: Effects on Earnings

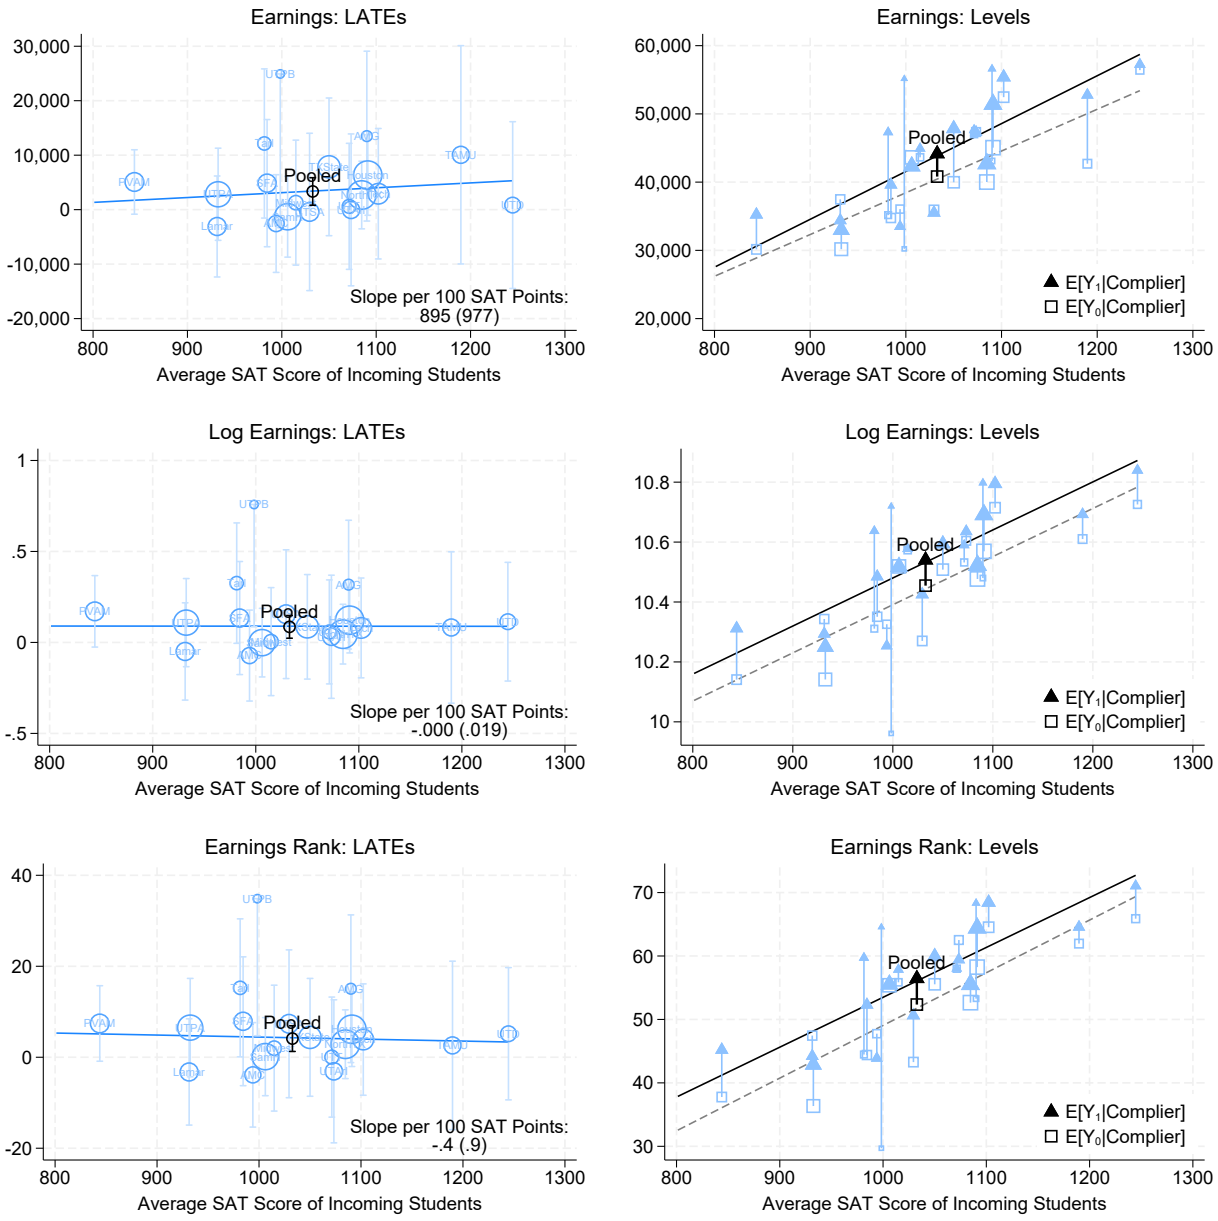

Notes: The darker "Pooled" estimate reproduces the pooled LATEs from Section 3, while the lighter estimates are LATEs for each target university estimated separately using the main specification in Section 2.4. The diagonal line is the estimated linear relationship between an institution's LATE (left plots) or potential outcome levels (right plots) and its selectivity. Estimates for the smallest institutions are not plotted due to their large CIs, but those institutions are included in the pooled estimates and slopes. The upper half of the confidence interval for UT-Permian Basin (UTPB) is omitted from the plot for readability, since it extends far beyond the current plot border.

selective schools, regardless of any causal impacts of the schools themselves. Relatedly, note that we can interpret the LATEs on average peer earnings in Figure 15 as the earnings effects we would predict for cutoff compliers if we were to use the most common measure of a college's value-added, which is simply the mean earnings of its students relative to students at other institutions. Section 3.4 discussed how such an approach would overpredict the value-added experienced by the average cutoff complier by a factor of two: they attend an institution with $6,700 higher average peer earnings as a result of admission, but gain only $3,300 themselves. The results in this section show that such an approach would also overpredict the relationship between selectivity and value-added by a factor of three: comparing the slopes in Figures 15 and 16 shows that a 100 SAT point increase in the selectivity of a complier's target school predicts a $3,000 higher gain in peer earnings, but only a statistically insignificant $900 higher gain in the complier's own earnings.

5.2 Across Institutional Spending

Selectivity, of course, is not the only important dimension of heterogeneity across postsecondary institutions. In Figure A.14, instead of ordering target universities by their average SAT score, I order them by their average educational spending per student. The resulting slopes offer suggestive, though not definitive, evidence that cutoff compliers at more "resourced" institutions experience larger earnings gains. The top panel of Figure A.14 shows a statistically significant positive slope when measuring complier earnings effects in dollars; the slopes are also positive, but not statistically significant, when measuring earnings gains in logs and ranks in the middle and bottom panels, respectively. Since each institution's LATE is a causal estimate but their collective slope through institutional spending is merely a correlation, these results cannot speak directly to the causal effect of additional educational spending per se on student earnings. The results simply illustrate that postsecondary institutions differ along multiple dimensions of "quality," and that some of these dimensions may have greater predictive power than others when predicting an institution's causal impacts on its students' outcomes.[38]

5.3 Costs

Finally, Figure A.15 explores how the private and social costs of educating marginally admitted students vary across institutions. The top row shows that cutoff compliers at more selective institutions tend to generate similar cumulative educational resource costs compared to cutoff compliers at less selective institutions, with a slope indistinguishable from zero. The second row, on the other hand, shows that cutoff compliers at more resourced institutions (those with greater educational expenditure per student) generate higher cumulative educational resource costs compared to cutoff compliers at less resourced institutions, as would be expected given the higher levels of per-student expenditure. From the marginal student's private perspective, the bottom two rows of Figure

[38] Using a different research design studying relative value-added across institutions within the four-year sector, Mountjoy and Hickman (2021) find a similar pattern of larger earnings effects of attending higher-resourced Texas public universities but not from attending more selective ones.

A.15 show that cutoff compliers at both more selective institutions and more resourced institutions actually pay weakly less in additional net tuition as a result of marginal admission compared to compliers at less selective and less resourced institutions.

6 Heterogeneity Across Student Backgrounds

This section explores heterogeneity in the effects of public universities on marginally admitted students from different demographic backgrounds. To do so, I run the main fuzzy RD specification from Section 2.4 separately across subgroups by gender, family income, and race.

6.1 Gender

First, with respect to gender, Figure 17 shows that male cutoff compliers reap somewhat larger gains in dollar earnings and BA completion compared to female compliers. These effects are not statistically distinguishable from each other, however, and the earnings gains are roughly similar in proportional terms given men's larger baseline earnings levels, as shown in both the top right panel of Figure 17 and in the similar log earnings estimates over years 8-12 plotted in Figure 18.

Interestingly, women begin to reap their earnings gains much more quickly than men, as shown in the left panel of Figure 18: the crossover point into positive earnings effects is around year 5 for women but not until year 8 for men.[39] The right panel of Figure 18 suggests that this pattern is driven by male compliers taking longer to leave college: cutoff crossing increases the probability that men are enrolled in any college in a given year up until 8 years out from application, compared to only 6 years out for women, mirroring the differential timing of their crossover points into positive earnings gains.

6.2 Family Income

Substantial heterogeneity in earnings gains emerges with respect to family income, as proxied by eligibility for free or reduced-price lunch (FRPL) in high school. The estimates in Figure 17 show that compliers from low-income families (FRPL) experience significantly smaller earnings gains, indistinguishable from zero, 8-12 years out from application compared to compliers from higher-income families (non-FRPL). In the top panels of Figure 19, I plot the dynamic paths of these earnings effects. The dynamics offer some suggestive evidence that low-income compliers may eventually reap meaningful earnings gains, thanks to an upward trend over years 8-12. What is clear, however, is that compliers from higher-income families start to enjoy significantly larger earnings gains much earlier than their peers from low-income families.

Digging into potential mechanisms, first note that these patterns are not readily explained by differential increases in college quality, as the bottom panels of Figure 17 show very similar increases in peer test scores and institutional spending for lower-income and higher-income compliers. Degree

[39] The earlier crossover point for women is also common in observational comparisons, e.g. Figure 3 of Mountjoy (2022).

Figure 17: Across Student Backgrounds: Effects on Earnings, Degrees, and Institutional Quality

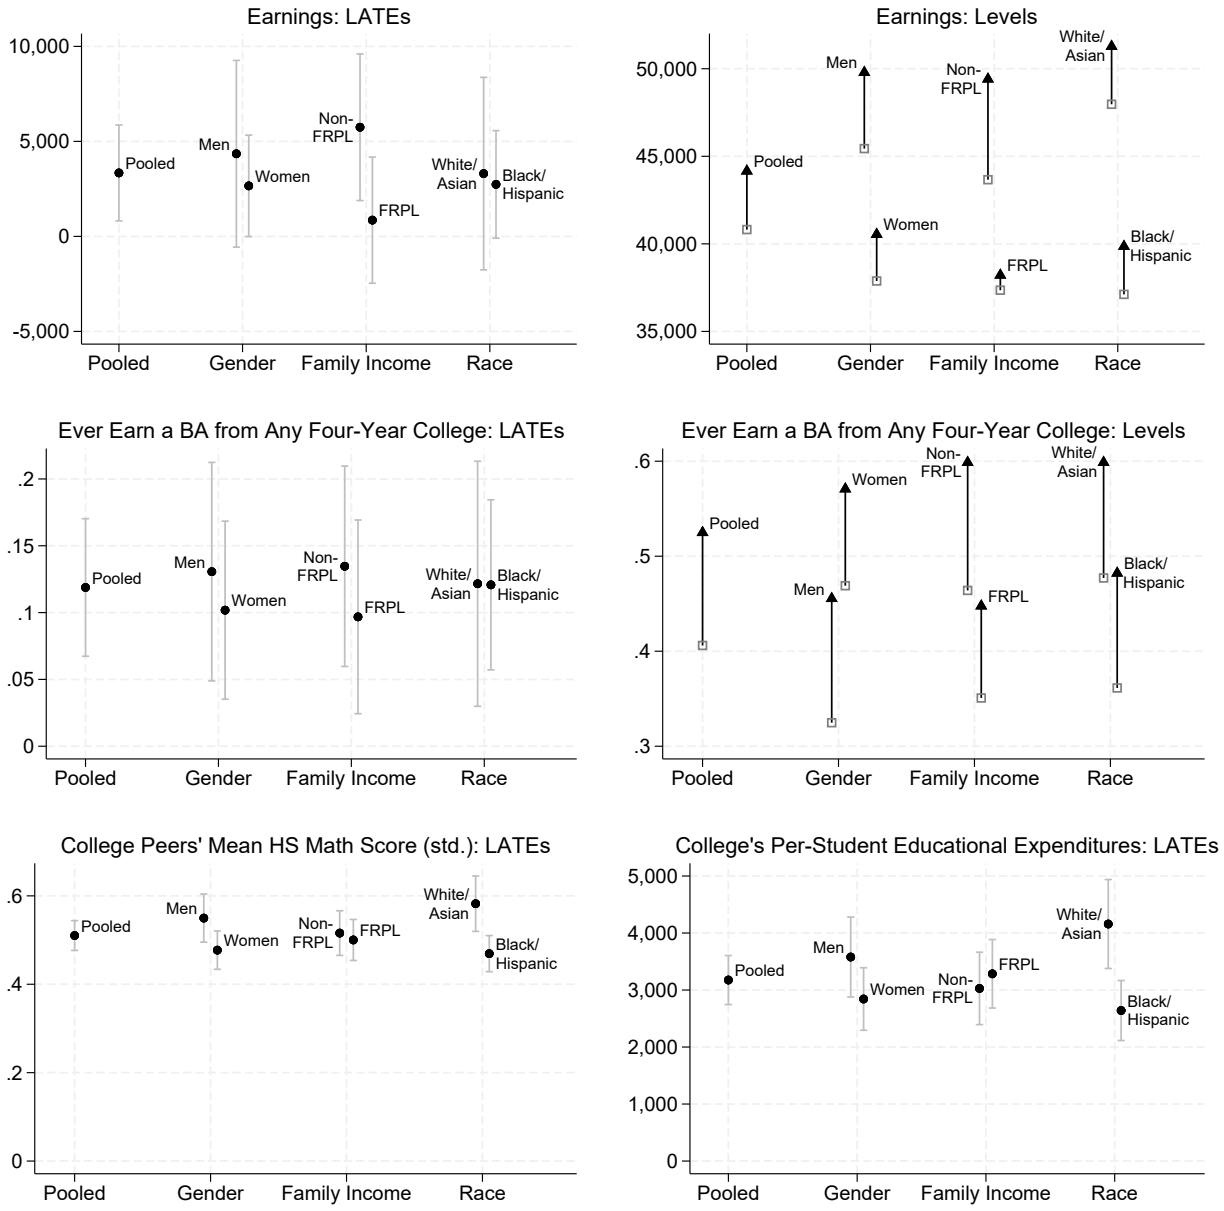

Notes: The "Pooled" estimate reproduces the pooled LATEs or potential outcomes from Section 3, while the other estimates come from running the main specification in Section 2.4 separately for each student subgroup. The "Levels" plots decompose the LATEs into the complier mean treated potential outcome level $\mathbb{E}[Y_1|\text{Complier}]$ in triangles and the complier mean untreated potential outcome level $\mathbb{E}[Y_0|\text{Complier}]$ in hollow squares. FRPL denotes eligibility for free or reduced-price lunch in high school, used as a proxy for low family income.

Figure 18: Across Gender: Dynamic Effects on Earnings and Enrollment

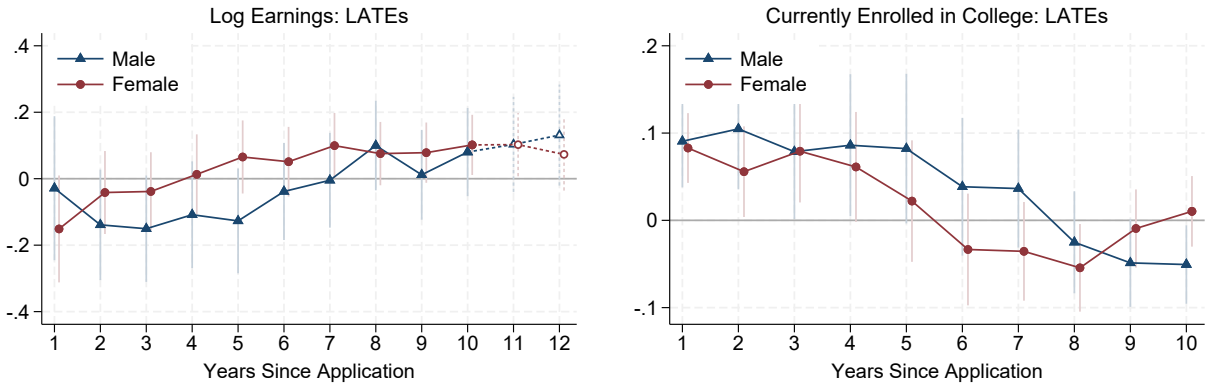

Notes: Each estimate comes from the main RD specification in Section 2.4 run separately by gender and year. In the left panel, the solid markers denote the balanced panel of outcomes up through year 10 that are observed for the entire sample, while the hollow/dotted markers come from progressively older subsets of the sample.

completion, on the other hand, may play an important role. Figure 17 shows that cutoff-crossing induces only a 9.7 percentage point gain in eventual BA completion for low-income compliers, compared to a 13.5 percentage point gain for higher-income compliers.[40] The middle left panel of Figure 19 shows that this degree gap opens up very early, with many higher-income compliers gaining an on-time BA at 4 years out while almost no low-income compliers do. Furthermore, the middle right panel of Figure 19 shows that low-income compliers, despite their lower BA completion gain, actually accumulate significantly more additional credits in the four-year sector than higher-income compliers, yielding an unfavorable combination of drawn-out college enrollment with fewer degrees to show for it.

The bottom panel of Figure 19 shows that differences in the types of majors completed by low-income versus higher-income compliers may also drive their differential earnings gains. The results on the left side of the plot replace each student's actual earnings outcome with the average earnings of all other applicants in the sample who completed the same majors at the same degree levels. The results show that cutoff-crossing induces lower-income compliers to end up with degrees in lower-earning majors compared to higher-income compliers. The results on the right side of the plot show that the gap is much smaller when defining premia at the institution-degree level, ignoring majors, which provides further evidence that these results are not readily explained by differences in institutional quality.

Thus, marginally admitted public university students from low-income and higher-income backgrounds both gain access to substantially higher-quality institutions, but low-income students struggle to generate early-career earnings gains for at least three potential reasons. First, they are less

[40] This family income gap in BA completion also shows up clearly in raw and controlled observational comparisons: low-income students in the broader applicant sample who enroll at a four-year Texas institution are 20 percentage points less likely to persist to a BA compared to their higher-income peers, with that gap remaining at 12 percentage points after controlling for rich pre-college covariates, and 10 percentage points after adding institution fixed effects.

Figure 19: Across Family Income: Dynamics and Potential Mechanisms

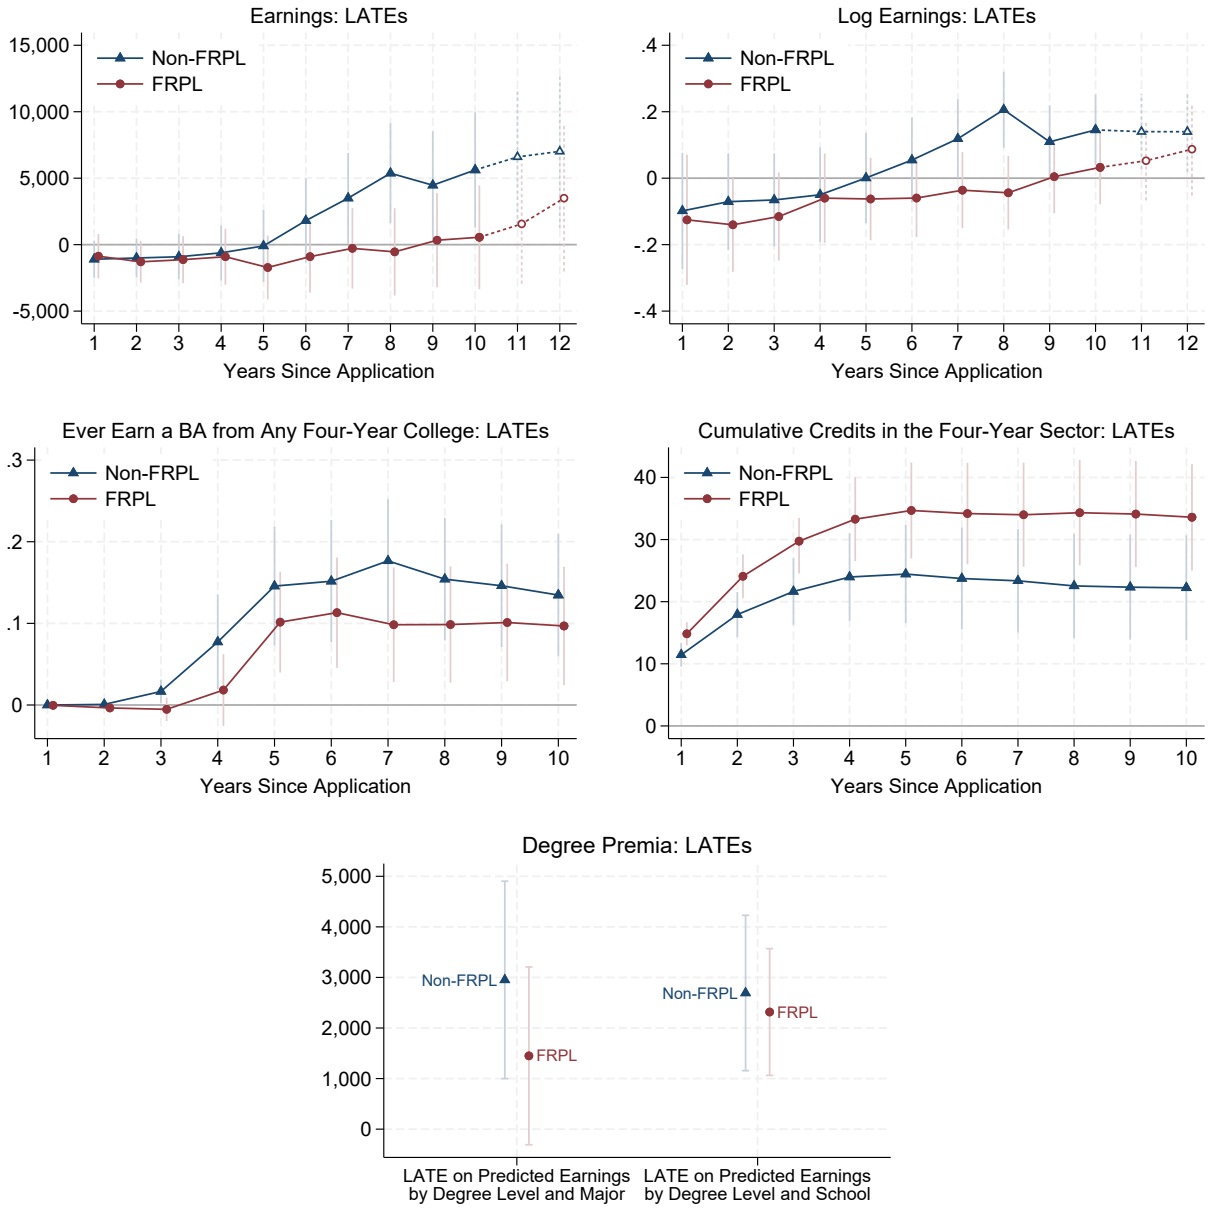

Notes: The estimates come from the main fuzzy RD specification described in Section 2.4 run separately for low-income (free or reduced-price lunch eligibility in high school, "FRPL") and higher-income (non-FRPL) applicants in each year. In the bottom plot, the results on the left side replace each student's actual earnings outcome with the average earnings of all other applicants in the sample who completed the same majors at the same degree levels, while the results on the right side replace each student's actual earnings outcome with the average earnings of all other applicants in the sample who completed the same degree levels (regardless of major) from the same institutions.

likely to persist to completing a degree. Second, despite completing fewer degrees, they actually spend more time enrolled, potentially giving up early-career earnings and experience. Finally, the degrees that they do complete are in less lucrative majors compared to those completed by their higher-income peers.

6.3 Race

With respect to applicant race, the rightmost estimates in each panel of Figure 17 show that white and Asian compliers reap similar gains in earnings and degree completion compared to Black and Hispanic compliers, despite white and Asian students experiencing larger increases in college selectivity and spending per student. The potential outcome levels reveal large baseline racial gaps in earnings and degree completion, with white and Asian students being over 10 percentage points more likely to have a BA and earning over $10,000 more than Black and Hispanic students with the same treatment status. Thus, marginally expanding admissions to public universities improves the outcomes of marginal students in a way that neither shrinks nor exacerbates large pre-existing racial disparities among them.

6.4 Costs

In Figure A.16, I document how the private and social costs generated by marginal students vary across these demographic groups. From society's perspective, the top row shows that compliers who are male, low-income, and Black or Hispanic generate relatively higher educational resource costs, but all of those subgroups also start from much lower untreated counterfactual levels and still fail to match the high treated levels of their counterpart subgroups (female, higher-income, and white or Asian compliers). In terms of net tuition paid by the student, the middle row shows that male and female compliers pay similar net tuition, both in LATEs and levels, while higher-income compliers pay dramatically more than low-income compliers, with treatment further increasing that gap. White and Asian students also pay much more than Black and Hispanic students, with treatment leaving that gap unchanged. With respect to student loans, the bottom row shows that marginal enrollment induces similar increases in loan amounts for male and female compliers, but women start from a much higher untreated baseline. Higher-income students generally take out more loans than lower-income students, but the LATE is larger for low-income compliers, so the enrollment treatment shrinks the gap. Finally, marginal Black and Hispanic students generally take out more loans than white and Asian students, and the enrollment treatment increases that gap.

7 Attending Any University vs. Attending a More Selective One

For the final dimension of heterogeneity, recall that roughly half of cutoff compliers would fall back to another, typically less selective, four-year institution if barely rejected from the target university, while the other half would initially fall out of the four-year sector, primarily to a two-year community college. In this section, I develop a method to disentangle the distinct contributions

of these "intensive" versus "extensive" margins of four-year enrollment. How much of the returns to enrolling marginal applicants are driven by extensive margin compliers induced to attend any four-year institution, versus intensive margin compliers induced to attend a more selective one?

7.1 Motivating and Summarizing the Separate Identification Method

Identifying separate treatment effects for these two complier types is hampered by the fact that they are not directly distinguishable in the data, as applicants' counterfactual fallback options are nowhere recorded and thus cannot be conditioned on. However, I do observe each applicant's portfolio of admissions to all 35 public universities in Texas, which enables a powerful stratification of marginal applicants into two observable groups: those who have at least one admission offer from *another* Texas public university besides the target school, and those who have none. This stratification does not perfectly separate intensive and extensive margin compliers, but it comes fairly close, and far closer 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 71 percent of compliers with another admission are on the intensive margin of the four-year sector, i.e. would fall back to that other available four-year option if rejected, while fully 93 percent of compliers with no other admissions are on the extensive margin, i.e. 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 Figure A.17, 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 admission to a fallback school within the same application year.[41]

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 summarize the method, developed in more detail in Appendix C, 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

[41] 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.

($Z = 0$) the target university's admission cutoff, respectively.[42] 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

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

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 ω is the share of cutoff compliers with $A_0 = 1$ instead of $A_0 = 0$.

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). However, in Appendix C I first show how to adapt the complier-describing logic of Abadie (2002) to identify several mean potential outcomes of the various "response types" of individuals with respect to the A stratifier: A-never-takers who never have another admission regardless of cutoff-crossing at the target school ($A_0 = 0, A_1 = 0$), A-always-takers who always have another admission ($A_0 = 1, A_1 = 1$), and A-compliers who secure another admission if and only if they fall below the cutoff at the target school ($A_0 = 1, A_1 = 0$). Estimating these mean potential outcomes amounts to running a series of fuzzy RD regressions that replace the outcome Y with interactions among Y, D, and A.

Figure 20 provides a visual summary of the method for the outcome of BA completion. The top left panel begins with the identified untreated (Y_0) mean potential outcomes among both subgroups of interest: the $A_0 = 0$ subgroup (A-never-takers), 93 percent of whom are on the extensive margin, and the $A_0 = 1$ subgroup (A-compliers and A-always-takers), 71 percent of whom are on the intensive margin. 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. To bound these quantities, I assume that A-always-takers, who always have at least one other admission, would tend to have weakly better treated outcomes than A-compliers, who scramble for another admission if and only 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—

[42] To reduce notation and simplify the exposition, here I suppress the conditioning on the running variable $R = r$ and define the binary instrument Z as falling "just above" versus "just below" the cutoff. The fuzzy RD structure remains in the background and in the estimation.

Figure 20: Disentangling the Intensive and Extensive Margins: BA Completion

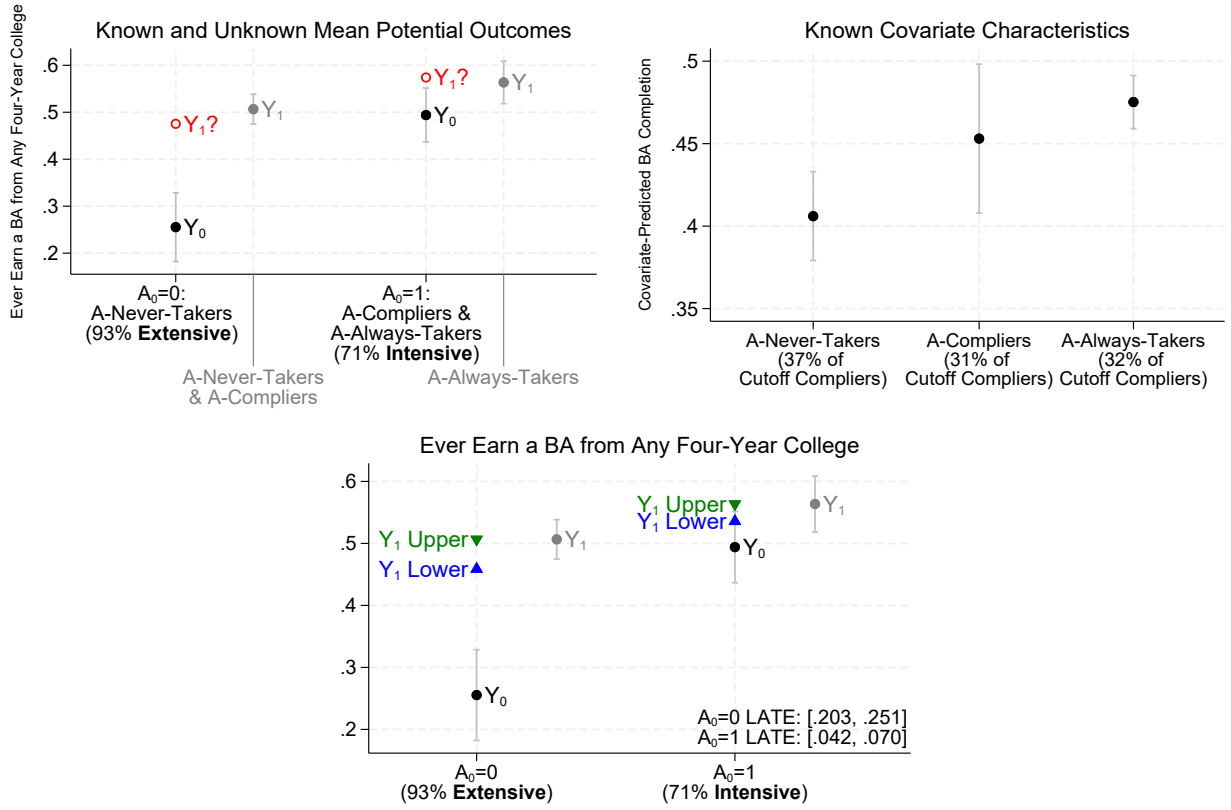

Notes: This figure visualizes the identification approach summarized in Section 7.1 and detailed in Appendix C. In the top left panel, the black mean Y_0's are identified under Assumptions 1 and 2; the hollow Y_1's are unknown but needed to identify the intensive vs. extensive margin LATEs of interest; and the gray mean Y_1's are not directly of interest but are identified under Assumptions 1 and 2 and will be used to learn about the unknown red Y_1's. The top right panel provides empirical support for Assumption 3; it plots the covariate characteristics (summarized as covariate-predicted BA completion) of the three response types. The bottom panel shows the upper and lower bounds on the mean Y_1's of interest implied by Assumption 3, and the resulting bounds on the separate LATEs for the almost entirely extensive margin $A_0 = 0$ group and the mostly intensive margin $A_0 = 1$ group.

an assumption reminiscent of Manski and Pepper (2000)'s monotone treatment selection, but here over the dimension of response types to the endogenous stratifier A rather than the endogenous treatment D.[43] The identified covariate characteristics of these three response types support such a ranking, as shown by the upward sloping pattern of covariate-predicted BA completion in the top right panel of Figure 20, as well as the upward sloping rankings across the identified Y_1's and across the identified Y_0's in the top left panel.

The bottom panel of Figure 20 shows that this ranking assumption 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, the originally identified mean Y_1 for A-always-takers is an upper bound for the mean Y_1 of the composite of

[43] See Lee and Salanié (2023) and Galperin (2023) for different but related assumptions.

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.

7.2 Decomposition Results

The bounds that result from this method end up being quite informative about the separate contributions of intensive and extensive margin compliers. For the outcome of BA completion, the bottom panel of Figure 20 shows that the LATE among the $A_0 = 0$ subgroup, who are almost entirely extensive margin compliers, is bounded between 20-25 percentage points,[44] which is roughly twice the 12 percentage point effect in the pooled complier population. Meanwhile, for the other subgroup with $A_0 = 1$, who are mostly (71%) intensive margin compliers and thus would tend to fall back to another less-selective four-year college if rejected, the LATE of enrolling at the target institution is just a 4-7 percentage point gain in BA completion. Furthermore, this small LATE for the $A_0 = 1$ group almost surely overstates the true intensive margin effect, since the stratification by A_0 does not perfectly isolate intensive margin compliers: 29% of the $A_0 = 1$ group are actually on the extensive margin, i.e. they have another four-year admission but do not excercise the option if rejected from the target school, falling out of the four-year sector instead. If these extensive margin applicants mixed into the $A_0 = 1$ group reap large treatment effects similar to the 93% extensive margin compliers with $A_0 = 0$, then the pure intensive margin treatment effect would be even smaller than the 4-7 percentage point effect for the $A_0 = 1$ group, since that effect is inflated by the 29% extensive margin compliers. Thus, the vast majority of the gain in BA completion estimated in Section 3 is driven by a large gain for extensive margin compliers who would not enroll in any four-year college if rejected, with a much smaller contribution from the intensive margin compliers who upgrade the selectivity of their four-year institution. Figure A.18 shows that these effects are statistically distinct via bootstrapped confidence intervals for the midpoints of each bound.

Figure 21 applies the bounding method to all the main outcomes of interest, with Figure A.18 providing bootstrapped confidence intervals for the midpoint of each bound. The top row shows larger increases in peer quality on the extensive margin, though note that intensive margin compliers also experience substantial increases in these measures of college selectivity: peer test scores increase by .40-.45 standard deviations in the mostly intensive $A_0 = 1$ group, and peer earnings increase by \$5,200-\$6,100. The second row shows dramatically larger extensive margin effects on cumulative credits in the four-year sector, with extensive margin compliers gaining almost two years of additional four-year education. That row also reproduces the earlier BA results for completeness. The next three panels decompose effects on the three different earnings measures. While the bounds for dollar earnings heavily overlap across the two margins, the bounds for log earnings and earnings rank do not, with substantively larger relative earnings gains for extensive versus intensive margin compliers.[45] This conclusion is further strengthened when noting, as above, that 29% of the

[44] Using a distance-based research design applied to older cohorts of Texas students, Mountjoy (2022) estimates a similar 20 percentage point increase in BA completion when shifting students from two-year to four-year colleges.

[45] This pattern across the different earnings measures could be driven by the intensive margin effects of college

Figure 21: Disentangling the Intensive and Extensive Margins

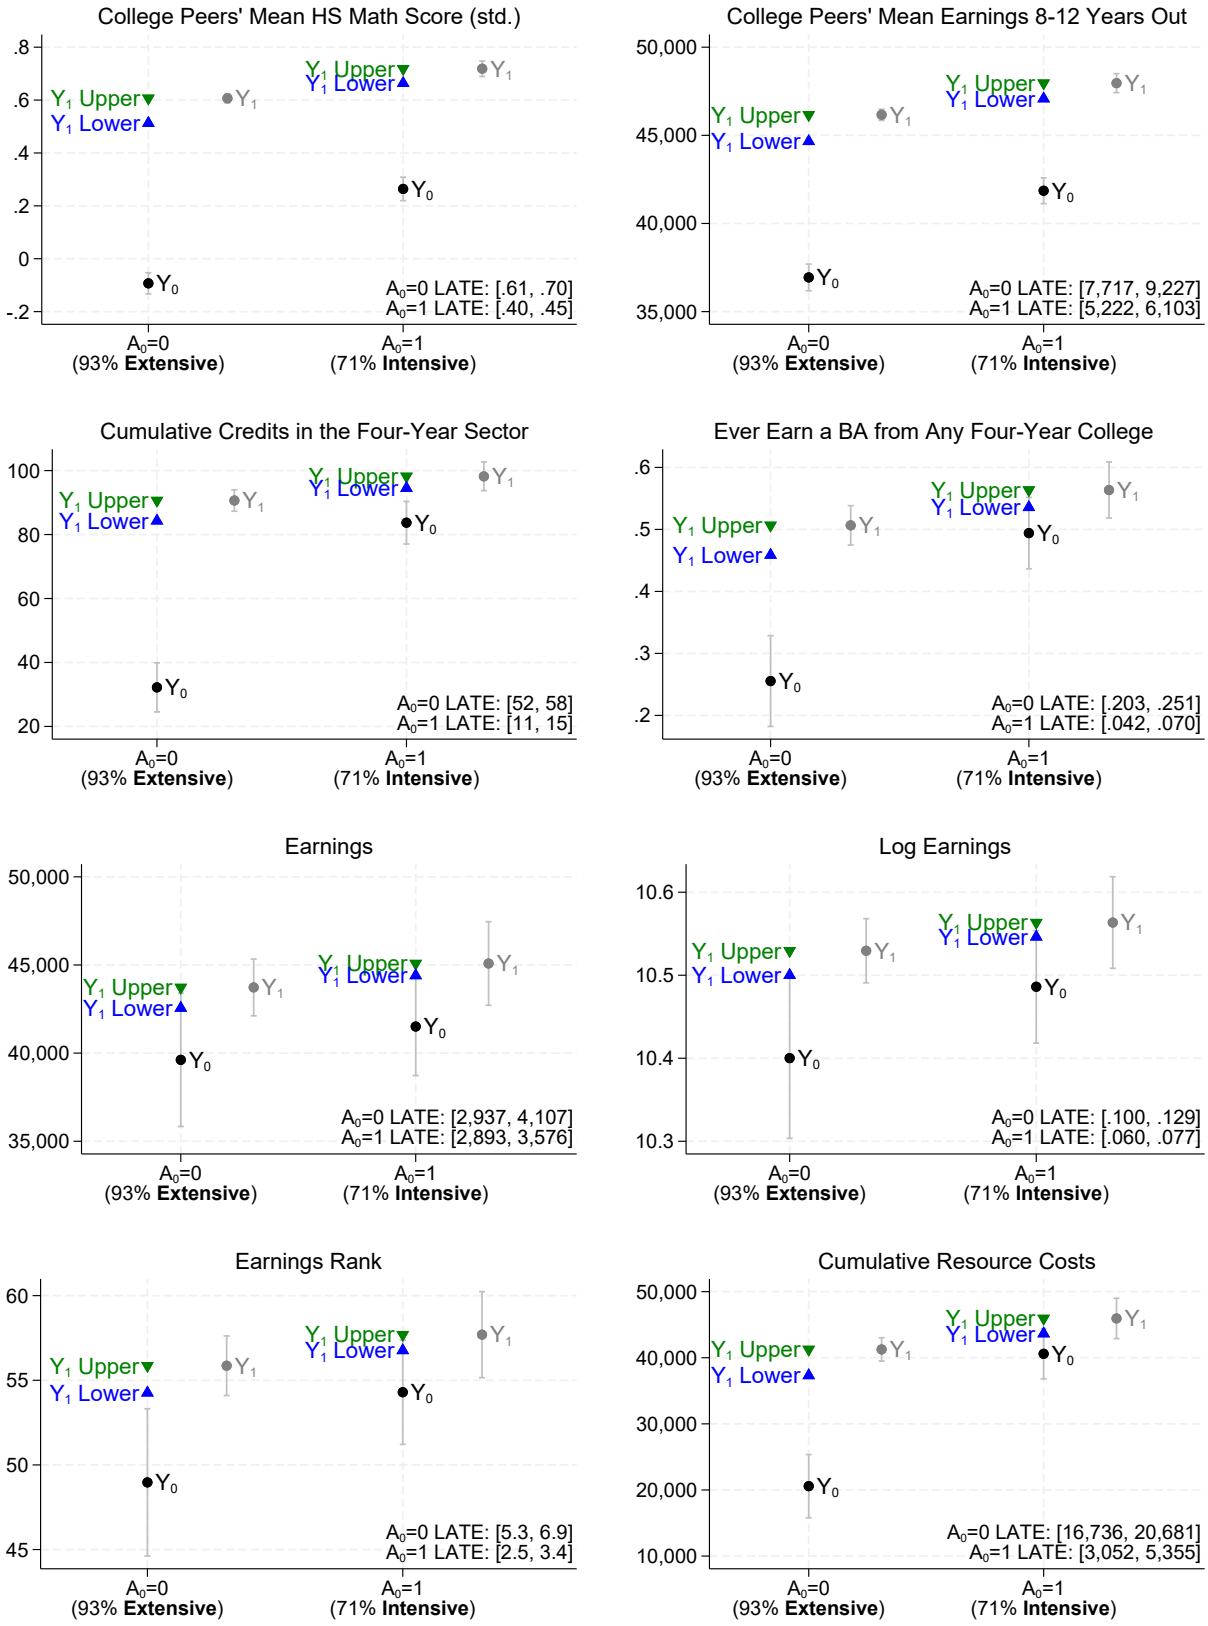

Notes: This figure plots the estimates that result from the identification approach summarized in Section 7.1 and detailed in Appendix C. See the notes to Figure 20. Figure A.18 provides bootstrapped confidence intervals for the midpoint of each bound.

$A_0 = 1$ group are actually on the extensive margin, which likely inflates the $A_0 = 1$ LATE relative to an effect for purely intensive margin compliers. Thus, the gross gains from enrolling marginal applicants are driven to a larger extent by extensive margin compliers induced into the four-year sector from outside it, with smaller gains for intensive margin compliers induced to attend a more selective school within it. The final panel of Figure 21, however, shows that the additional resource costs of educating intensive margin compliers are also small, allowing their small gross gains to still generate positive net returns.

8 Conclusion

Is the marginal American public university student a good investment? This paper used hundreds of decentralized admission discontinuities linked to administrative data spanning the second largest U.S. state to estimate the private and social returns to enrolling marginal applicants to public universities. Marginally admitted students, compared to their marginally rejected but otherwise identical counterparts, complete an additional year of four-year education, are 12 percentage points more likely to earn a bachelor's degree, and eventually earn 8 percent more. Marginally admitted students pay none of the additional tuition costs of these investments thanks to offsetting grant aid; formal cost-benefit calculations imply internal rates of return of 26 percent for the marginal students themselves, 16 percent for society, and 6 percent for the government budget.

Thus, despite the fact that marginally admitted students tend to arrive on campus with substantially weaker academic preparation and end up with below-average degree attainment and earnings relative to their university peers, the analysis in this paper shows that these outcomes are significant improvements over the typical trajectories these marginal students would have experienced had they been rejected instead. Moreover, since the benefits of enrolling marginal students surpass the costs, the results in this paper also suggest that marginally expanding admissions slots at public universities would tend to generate positive net returns, both for the newly admitted students themselves and for the taxpayers subsidizing the investment. Of course, when venturing beyond marginal expansions, these positive net returns could start to decrease if either the marginal benefits started to decrease, e.g. if students with even weaker academic preparation struggle to gain as much from the opportunity, or if the marginal costs of educating them started to increase, perhaps due to capacity constraints or more costly remedial education.

The analysis also uncovered notable dimensions of heterogeneity. Across institutions, the selectivity of a public university is not a strong predictor of its impact on the earnings of its marginally admitted students, while institutional spending per student may be a more useful predictor. In either case, students actually tend to face lower net tuition when marginally admitted into institutions that score higher on these quality metrics. Across student demographics, women begin to reap their positive earnings gains much more quickly than men, likely explained by men taking longer to finish college. Students from low-income families reap much smaller earnings gains compared

selectivity being more concentrated at higher incomes, where smaller changes in logs and ranks can still translate into large changes in dollars.

to their peers from higher-income backgrounds, likely due to low-income students spending more time enrolled but completing fewer degrees, as well as majoring in less lucrative fields. Finally, I developed a method to derive tight bounds on separate effects for students on the extensive margin of attending any four-year college versus those who would fall back to a less-selective one if rejected, revealing a larger role for the extensive margin in driving the overall gains.

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.

Bailey, M. J. and S. M. Dynarski (2011): "Gains and Gaps: Changing Inequality in U.S. College Entry and Completion," in *Inequality in Postsecondary Education*, ed. by G. Duncan and R. Murnane, New York: Russell Sage.

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*.

Belley, P. and L. Lochner (2007): "The Changing Role of Family Income and Ability in Determining Educational Achievement," *Journal of Human Capital*, 1, 37–89.

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

——— (2024): "Top Percent Policies and the Return to Postsecondary Selectivity," Working paper.

Bleemer, Z. and A. Mehta (2024): "College Major Restrictions and Student Stratification," National Bureau of Economic Research Working Paper No. 33269.

Bleemer, Z. and S. Quincy (2025): "Changes in the College Mobility Pipeline since 1900," National Bureau of Economic Research Working Paper No. 33797.

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*, 69, 1127–1160.

Card, D., J. Rothstein, and M. Yi (2025): "Location, Location, Location," *American Economic Journal: Applied Economics*, 17, 297–336.

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," National Bureau of Economic Research 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. 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," *The 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.

Kirkebøen, L., E. Leuven, M. Mogstad, and J. Mountjoy (2025): "College as a Marriage Market," NBER Working Paper No. 28688.

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.

Manski, C. F. and J. V. Pepper (2000): "Monotone Instrumental Variables: With an Application to the Returns to Schooling," *Econometrica*, 68, 997–1010.

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," National Bureau of Economic Research Working Paper No. 29276.

National Center for Education Statistics (2024): "Integrated Postsecondary Education Data System," .

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.

Selingo, J. (2020): *Who Gets in and Why: A Year inside College Admissions*, New York: Simon and Schuster.

Smith, J., J. Goodman, and M. Hurwitz (2025): "The Economic Impact of Access to Public Four-Year Colleges," *Journal of Human Resources*, Forthcoming.

Sorkin, I. (2018): "Ranking Firms Using Revealed Preference," *The Quarterly Journal of Economics*, 133, 1331–1393.

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," .

UTD-ERC (2025): "UT Dallas Education Research Center Data Holdings," .

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 Tables and Figures

Table A.1: Observational Estimates of the Return to a Year of College in Each Sector

	Effect on Log Earnings			
	(1)	(2)	(3)	(4)
One Year of Credits in the Four-Year Sector	.0948	.0505	.0558	.0394
(SE)	(.0002)	(.0003)	(.0007)	(.0007)
One Year of Credits in the Two-Year Sector	-.0001	.0004	-.0195	-.0156
(SE)	(.0003)	(.0003)	(.0012)	(.0011)
Sample: All High School Graduates	✓	✓		
Sample: RD Applicants			✓	✓
Covariate Controls		✓		✓
Number of Students	2,171,461	2,171,125	200,527	200,527

Notes: Each estimate comes from a multivariate OLS regression of log earnings (averaged over years 8-12 out from high school graduation) on cumulative years of credits (measured 10 years out from high school graduation) in the four-year sector and in the two-year sector. Covariate controls 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 A.1: Top Ten Percent Applicants Are Smooth Through the Cutoff

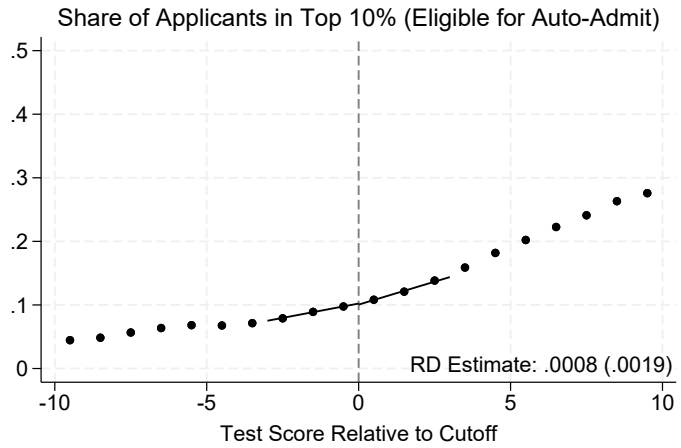

Notes: This figure plots the share of applicants at each of value of the concorded ACT running variable who are in the top GPA decile of their high school graduating class, making them eligible for automatic admission at Texas public universities (outside of UT-Austin, which has a stricter threshold)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

Figure A.2: Balance across Individual Pre-College Covariates

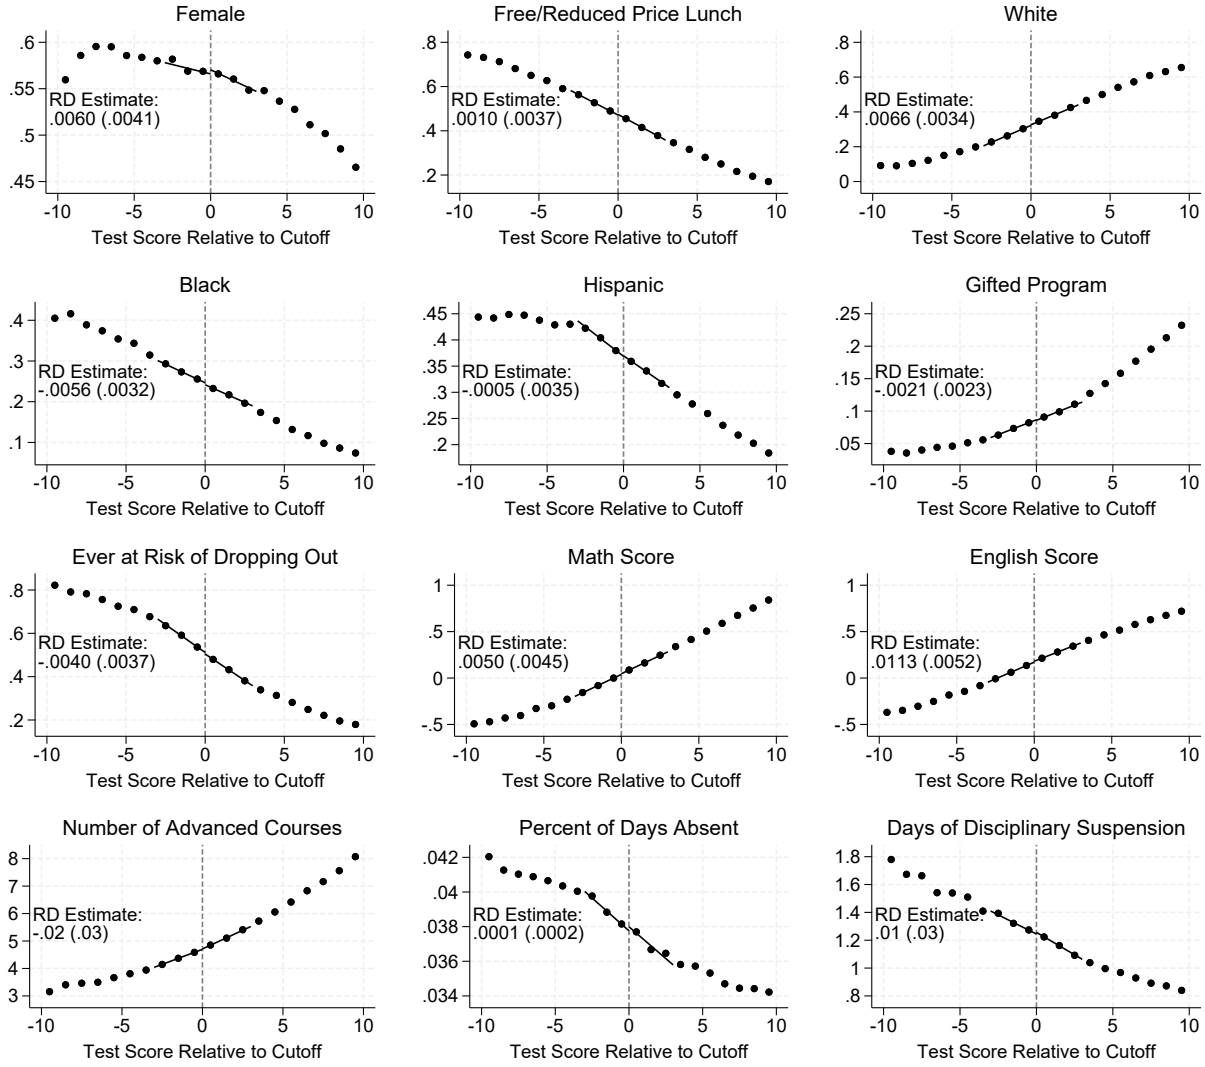

Notes: These plots show balance across the individual pre-college covariates used to predict BA completion and earnings in Figure 5. All covariates are measured during an applicant's observable high school years. Math and English scores are standardized to have mean zero and standard deviation one across all Texas public high school students taking the same test in the same year. The RD estimates come from the main specification described in Section 2.4.

Figure A.3: Private and Out-of-State College Enrollment

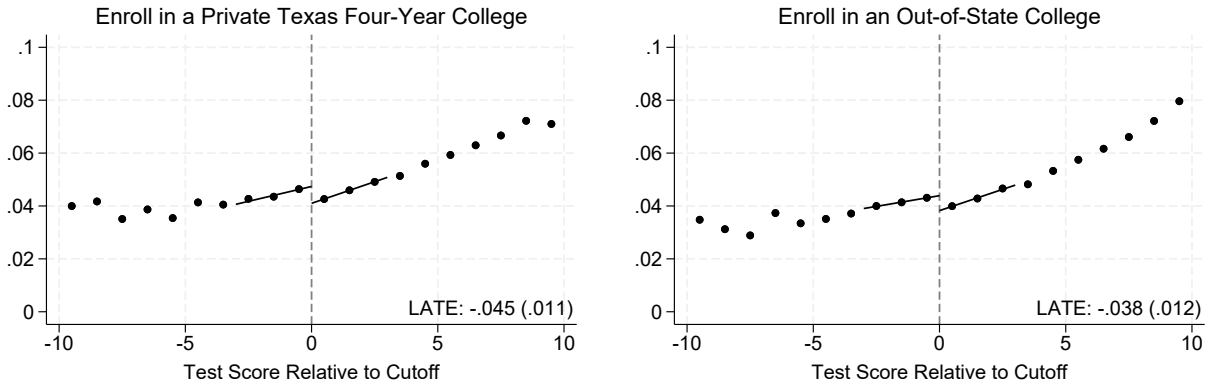

Notes: Enrolling in a private Texas four-year college is observed in the THECB data for all sample cohorts, and included in the outcome "Enroll in Another Texas Four-Year College" in Figure 6. Enrolling in an out-of-state college in the National Student Clearinghouse data is not observed for the 2004-2007 cohorts but observed for the 2008-2014 cohorts.

Figure A.4: Impacts on Long-Run Educational Attainment: RD Plots

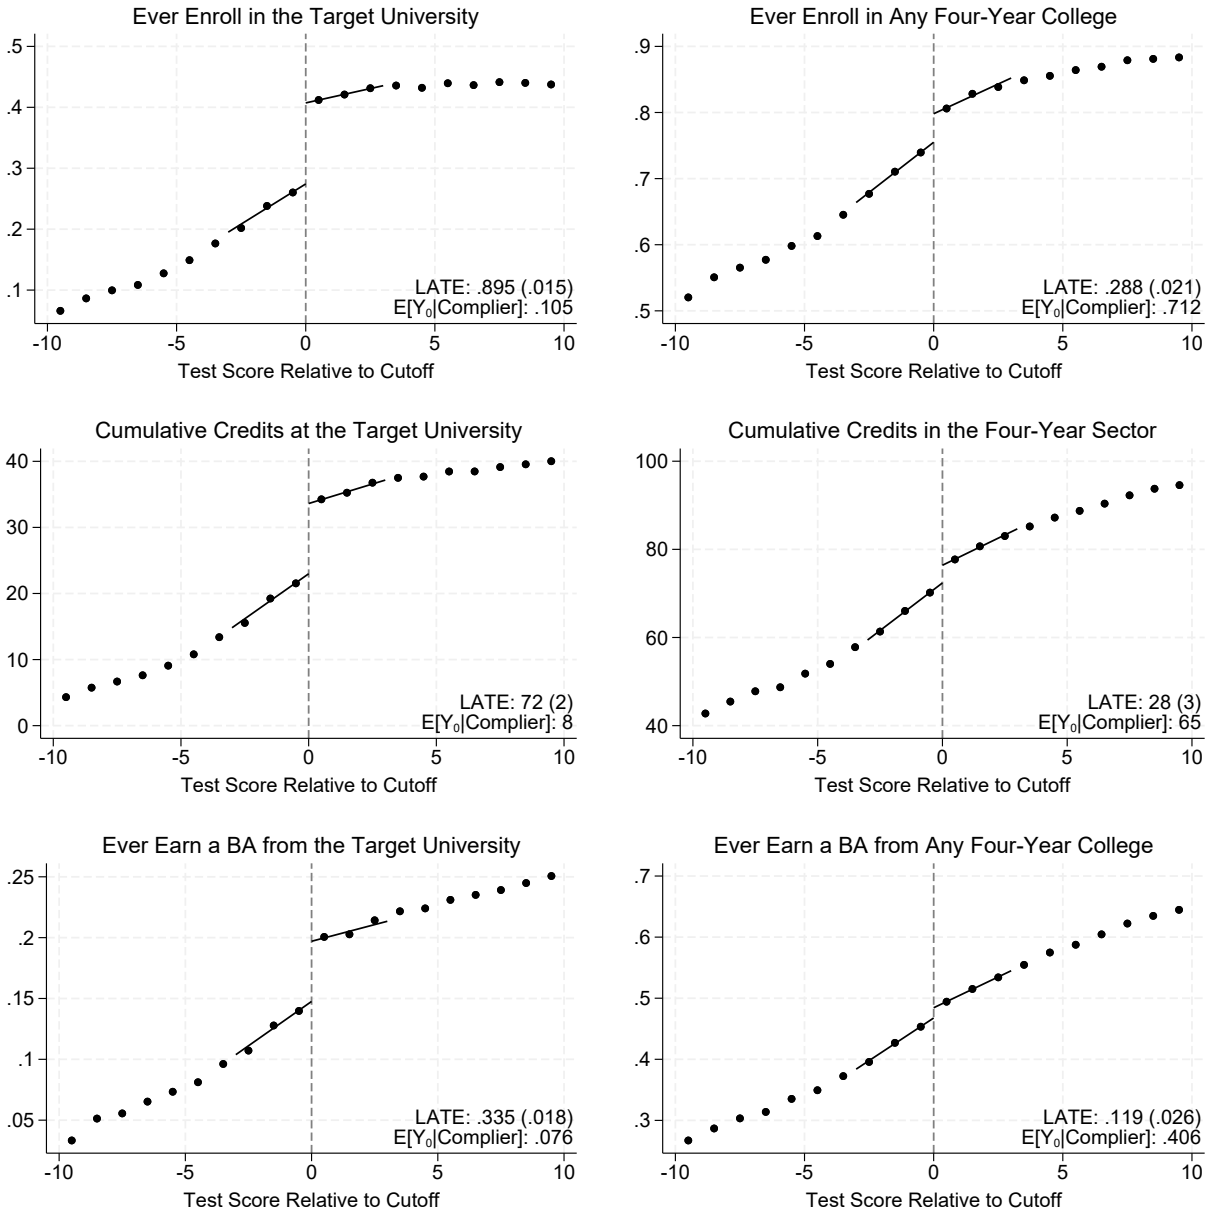

Notes: These RD plots correspond to the educational outcomes in Figure 8 measured 10 years out from application. The LATE estimates come from the main fuzzy RD specification described in Section 2.4.

Figure A.5: Impacts on Majors, Two-Year Credits, and Graduate School: RD Plots

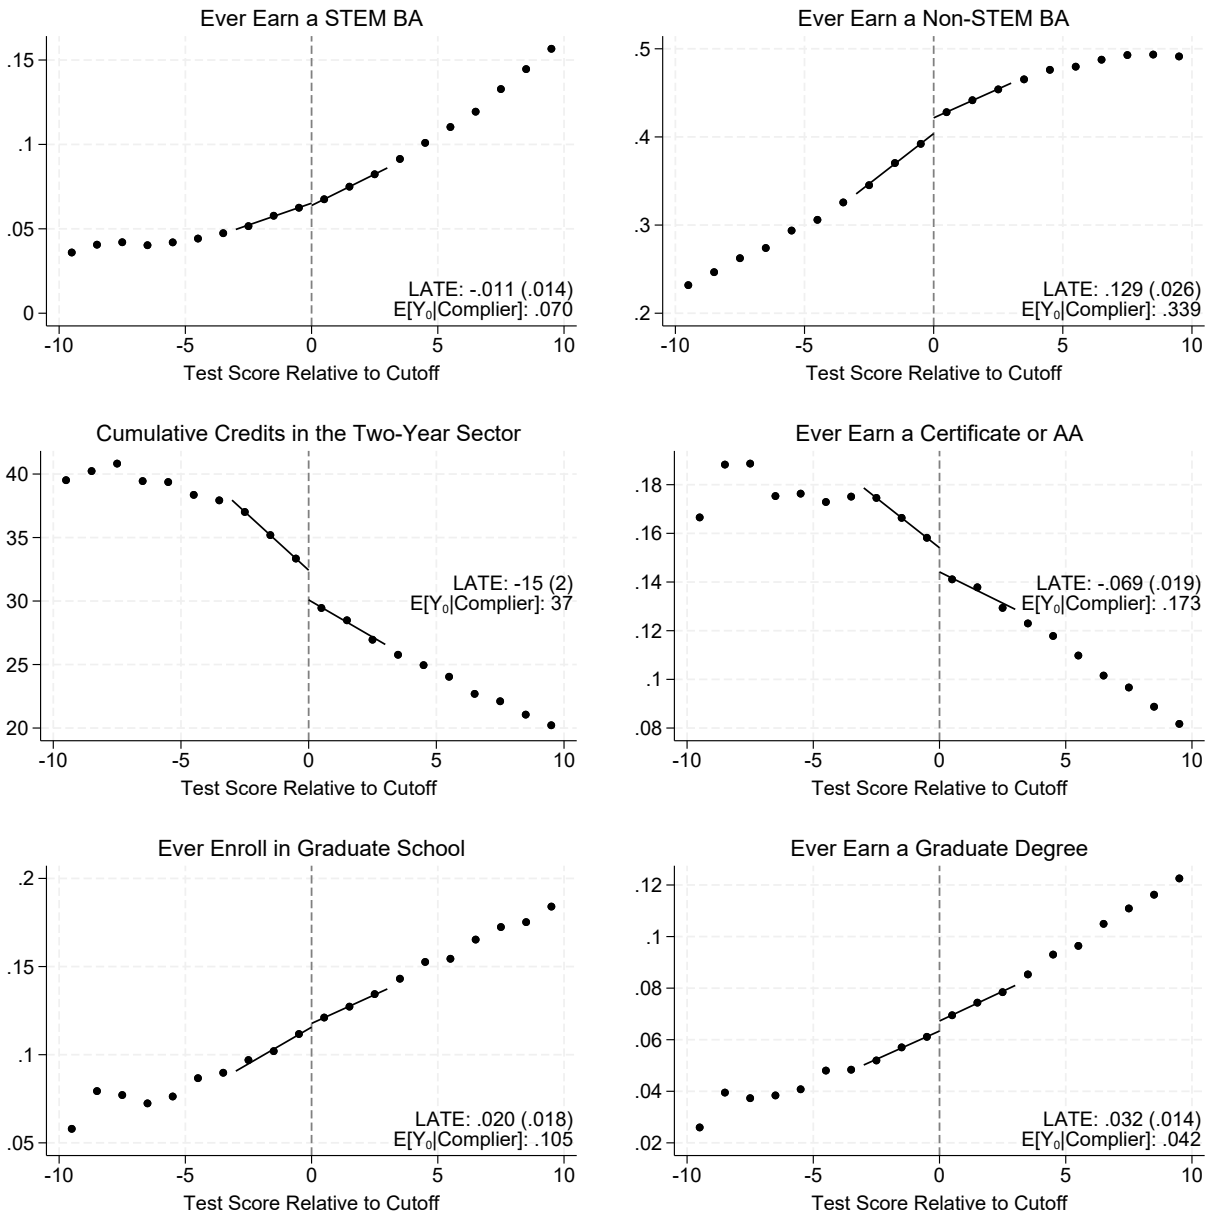

Notes: These RD plots correspond to the educational outcomes in Figure 9 measured 10 years out from application. The LATE estimates come from the main fuzzy RD specification described in Section 2.4.

Figure A.6: Robustness to Alternative Earnings Measures

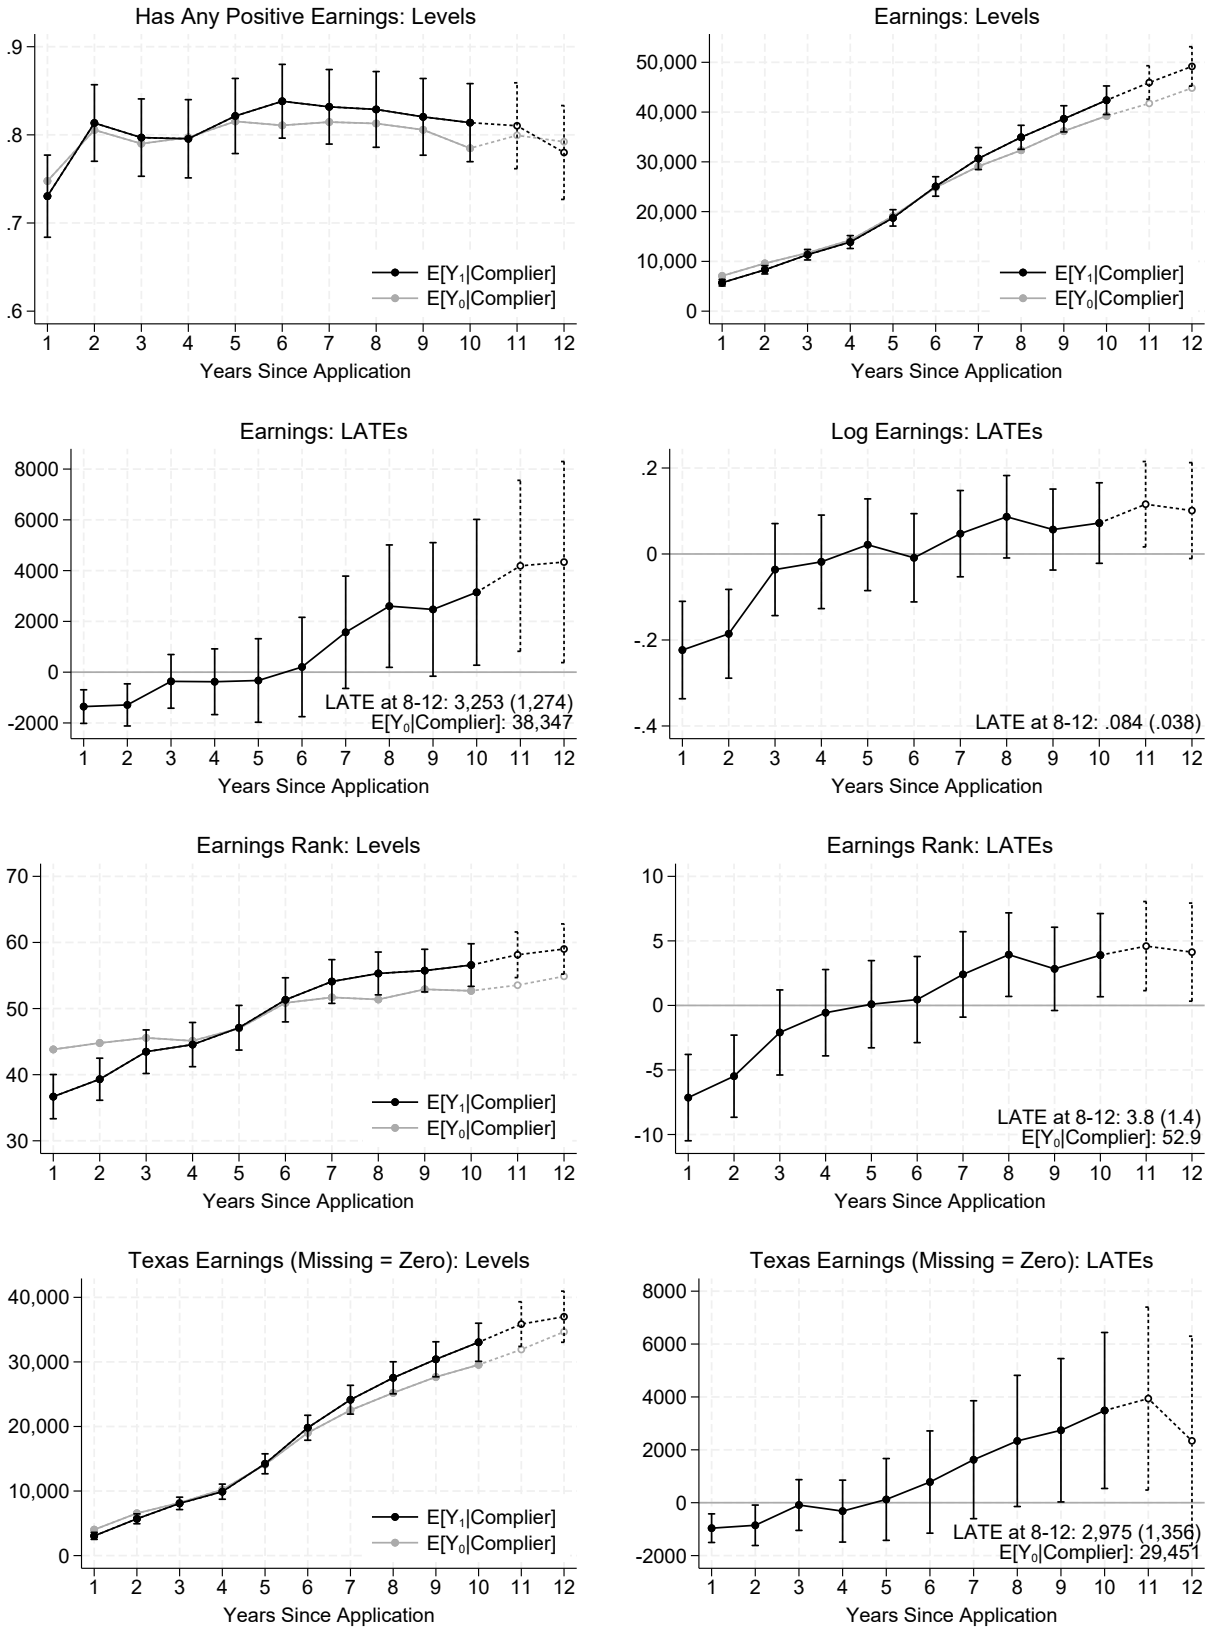

Notes: This figure presents an alternative to Figure 10 to explore robustness across alternative earnings measures. The top six panels use all positive earnings quarters within a person-year, keeping the transition quarters dropped in the sandwich earnings measure used in Figure 10. The bottom panels further include all missing quarters, imputing them as zero earnings, thus summing all earnings observed in Texas for each person-year. See Figure A.7 for the reduced form RD plot corresponding to each outcome.

Figure A.7: Robustness to Alternative Earnings Measures: Reduced Form RD Plots

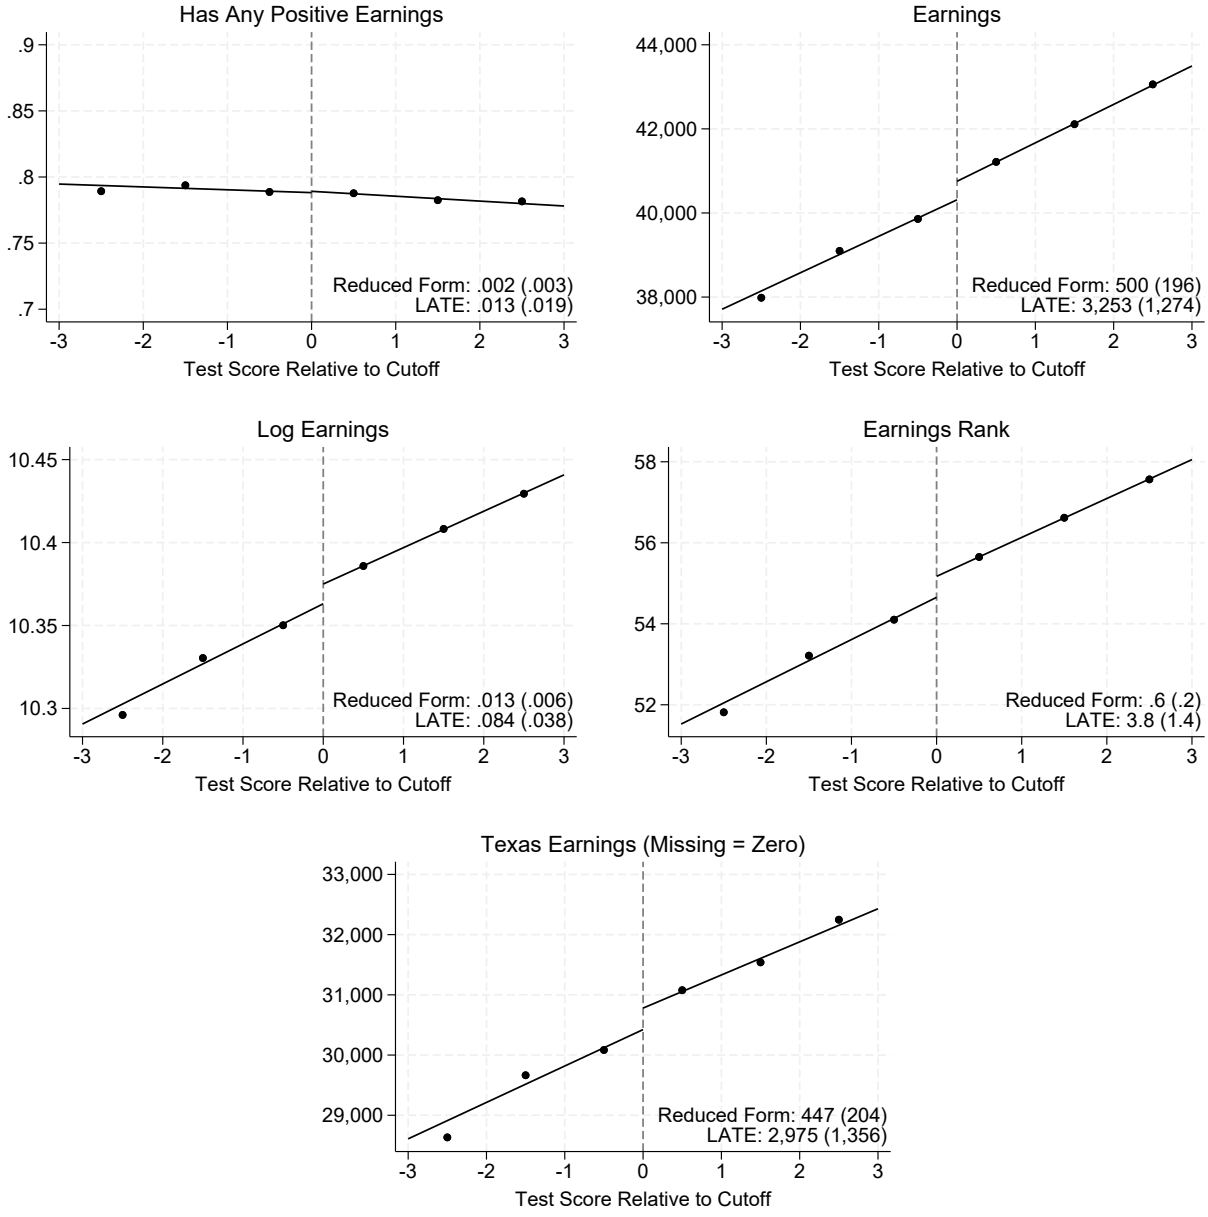

Notes: This figure plots the reduced-form discontinuities in the alternative earnings outcomes in Figure A.6 measured 8-12 years out from application. The estimates come from the main fuzzy RD specification described in Section 2.4.

Figure A.8: Robustness to Alternative Bandwidths

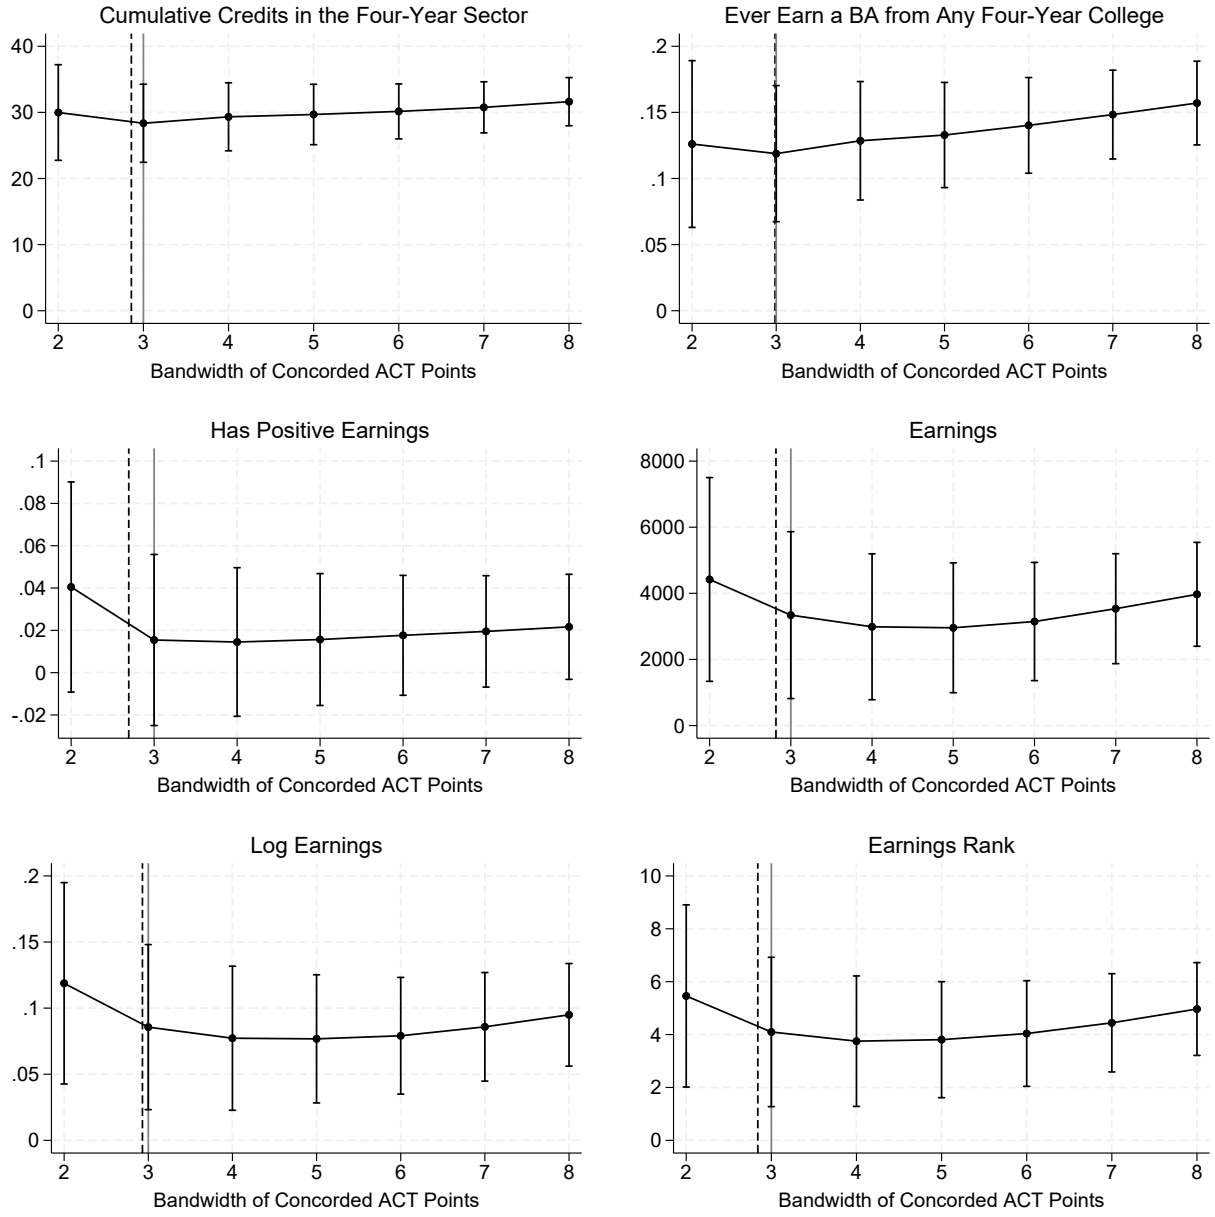

Notes: Each LATE point estimate and 95 percent confidence interval come from the main fuzzy RD specification described in Section 2.4, estimated separately using the different bandwidths of concorded ACT points indicated on the horizontal axis. The solid gray vertical line marks the bandwidth of 3 used 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.9: Robustness to Alternative Specifications

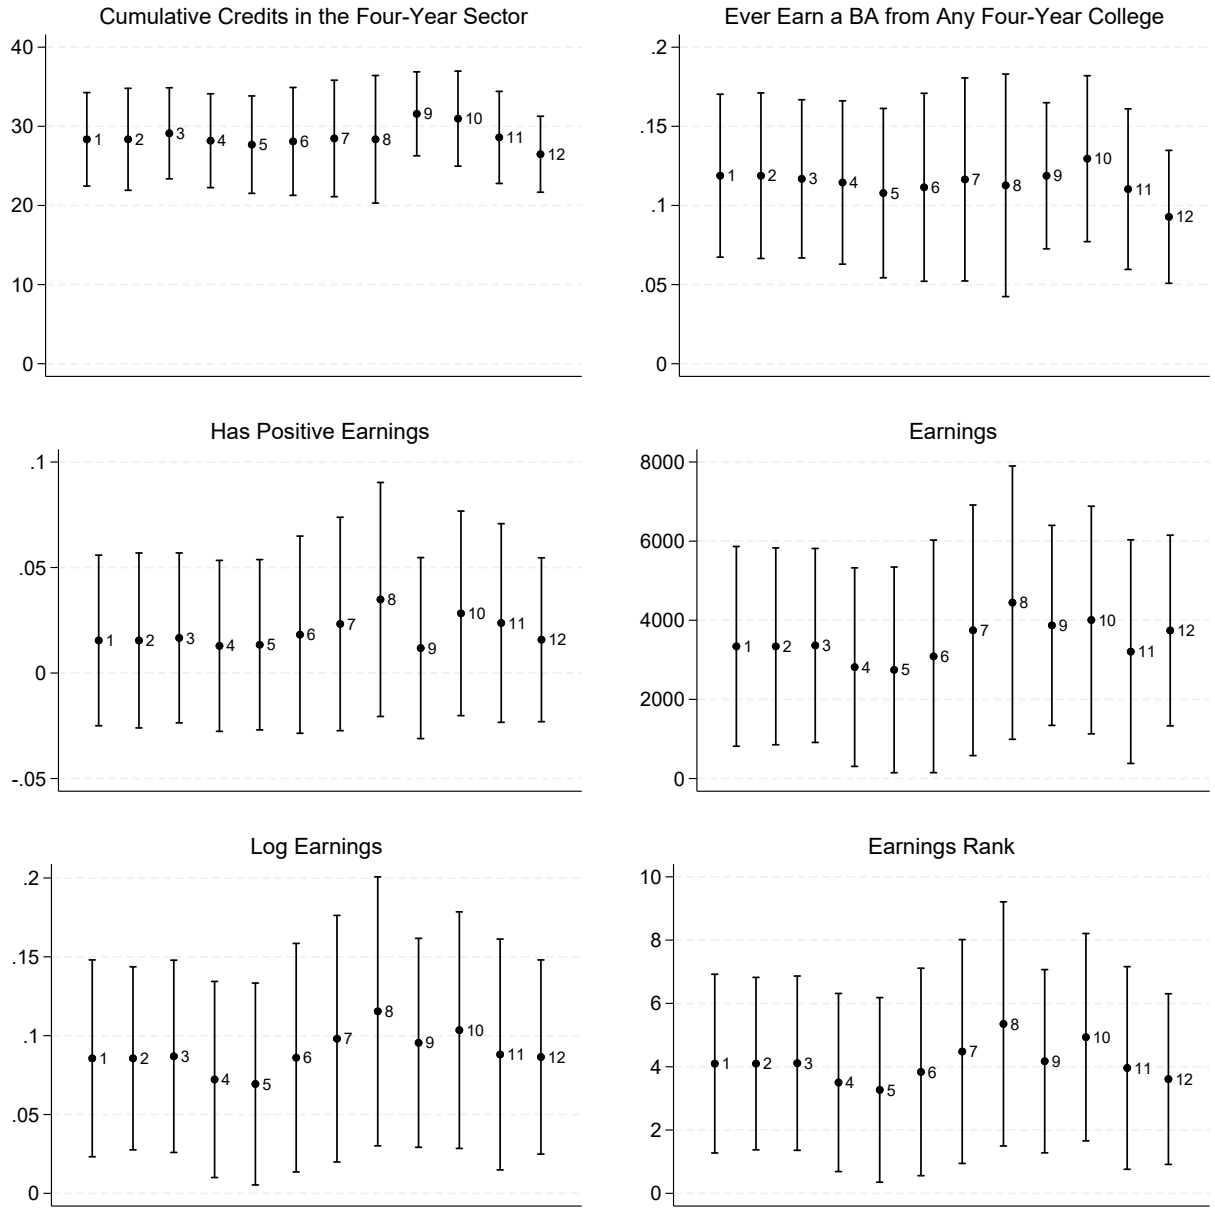

Notes: Each LATE point estimate and 95 percent confidence interval come from a different fuzzy RD specification. Specification 1 reproduces the main estimates from the baseline specification described in Section 2.4. Specification 2 clusters the standard errors at the application cell level in addition to the applicant level, yielding very similar confidence intervals. Specification 3 adds the full suite of pre-college covariates described in footnote 22 to the main specification. Specification 4 removes the application cell-specific slopes from the main specification, controlling only for cell fixed effects and common slopes above and below the cutoff. Specification 5 further removes the application cell fixed effects, leaving only a common intercept and common slopes above and below the cutoff. Specification 6 changes the local polynomial functional form from linear to quadratic in the running variable on either side of the cutoff, first with a bandwidth of 6 concorded ACT points. Specification 7 narrows the quadratic specification bandwidth to 5 ACT points, and Specification 8 further narrows it to 4 ACT points. Specification 9 includes rather than disqualifies the roughly 20% of initially inferred cutoffs that exhibit statistically significant discontinuities in the log density of applicants or (more rarely) covariate-predicted BA completion or earnings. Specification 10 restricts the main set of cutoffs to those with statistically significant (at the 1% level) jumps in admission and enrollment, and Specification 11 does the same using a 5% significance threshold. Specification 12 allows each application cell to potentially contribute multiple cutoffs, keeping all distinct cutoffs within a cell with statistically significant (at the 5% level) jumps in admission and enrollment.

Figure A.10: Impacts on Long-Run Educational Costs: RD Plots

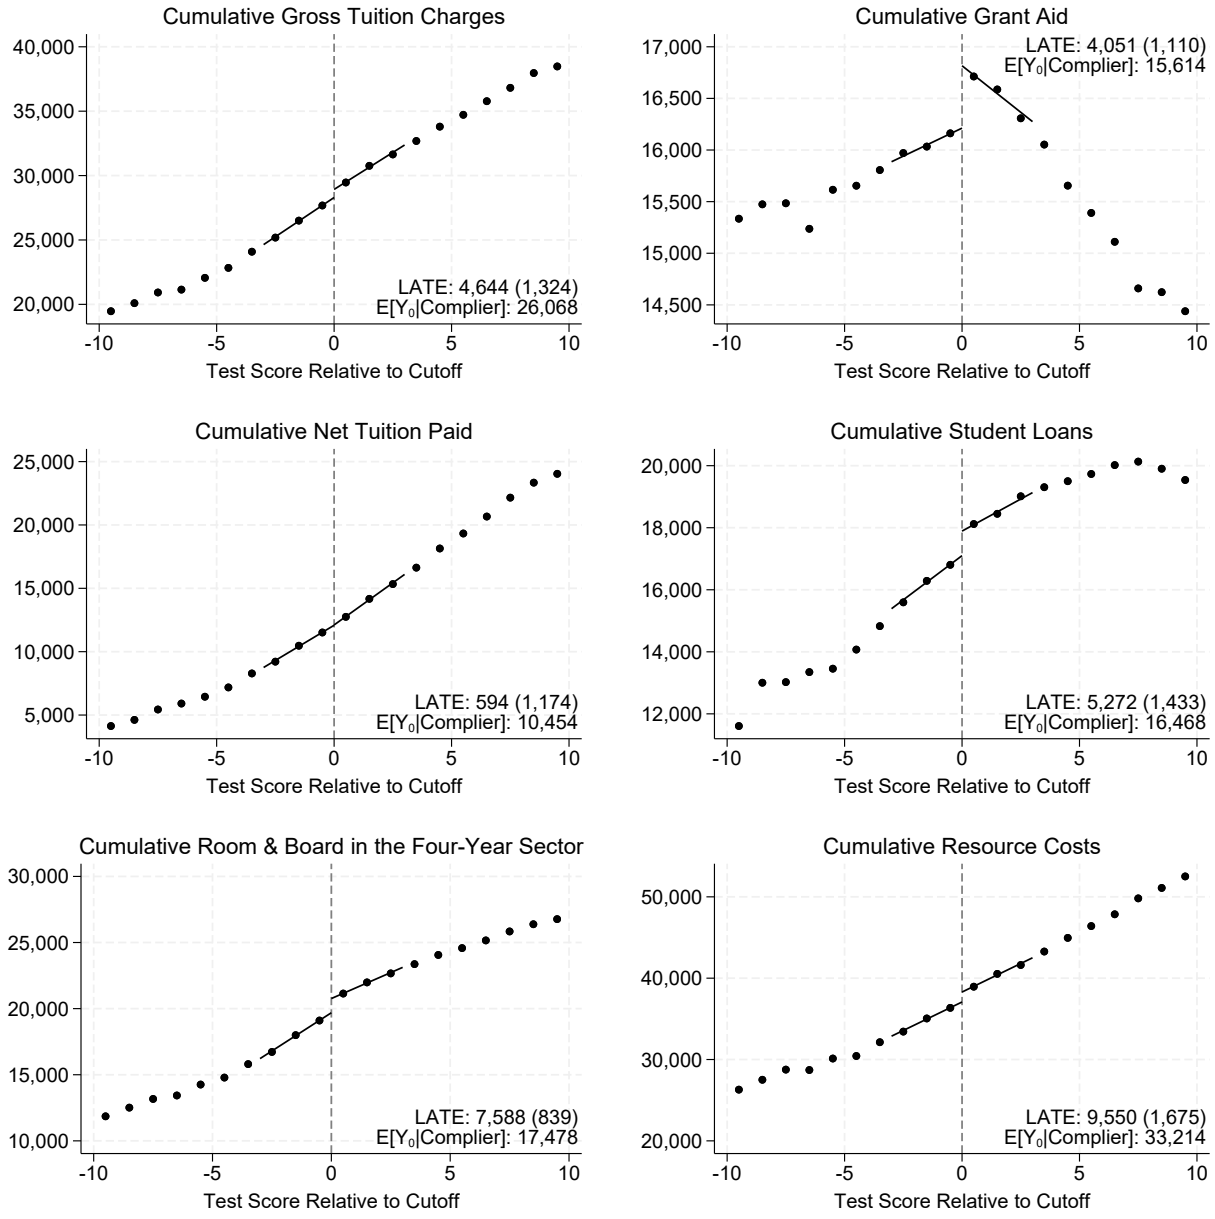

Notes: These RD plots correspond to the cost outcomes in Figure 12 measured 9 years out from application. The LATE estimates come from the main fuzzy RD specification described in Section 2.4.

Figure A.11: Cost-Benefit Calculations Using the Texas Imputed-Zero Earnings Measure

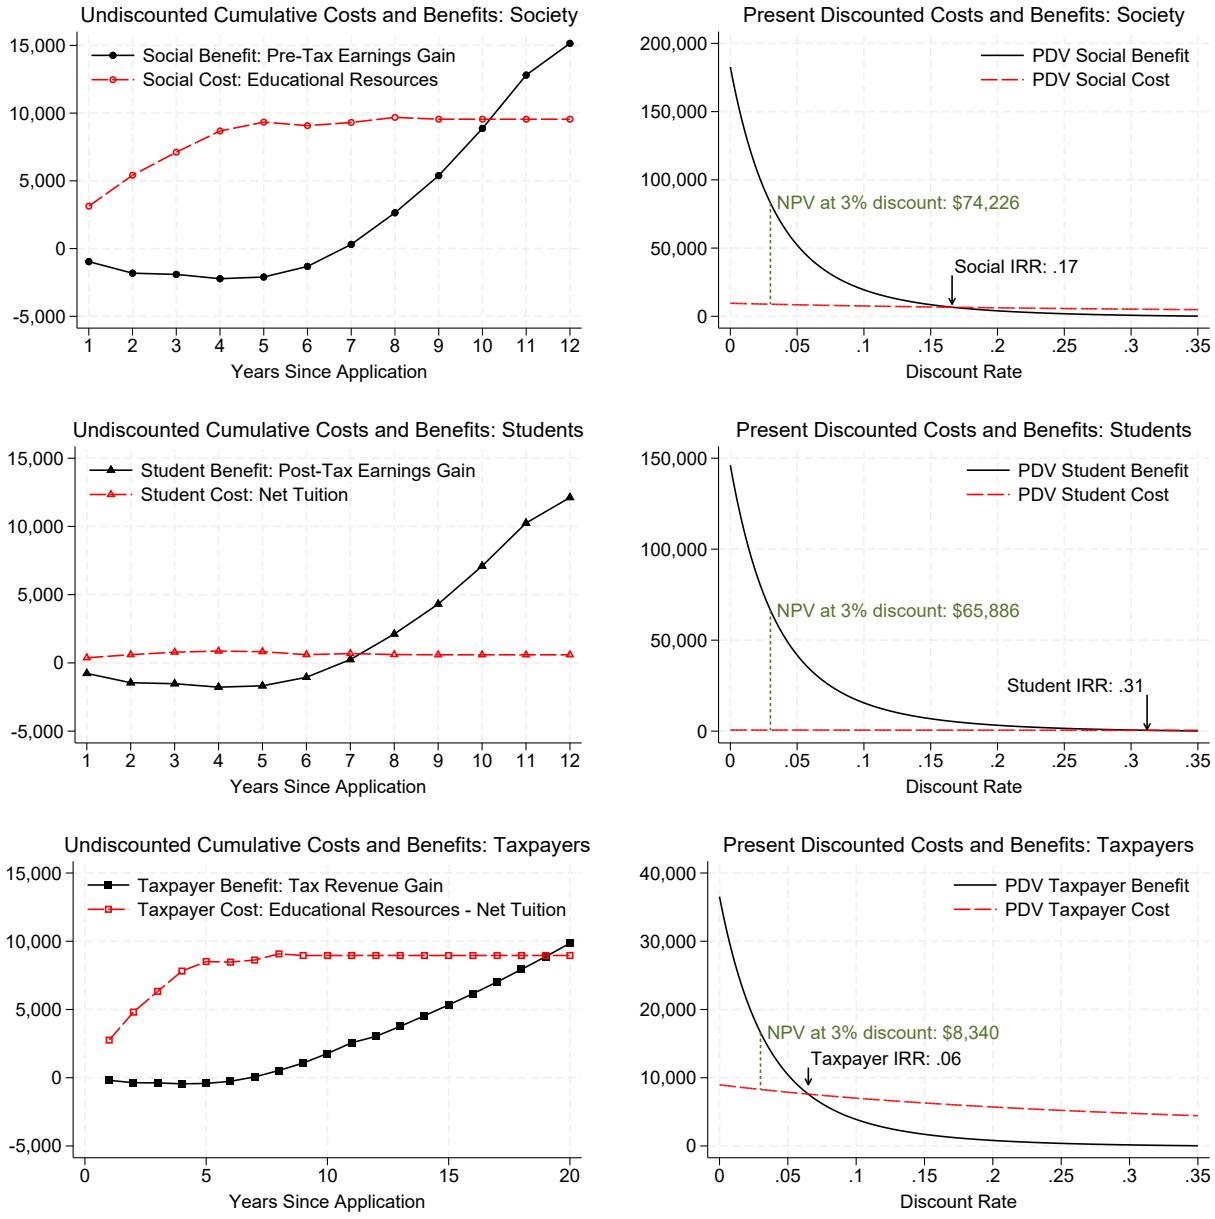

Notes: This figure is an alternative to Figure 13 that uses the Texas earnings measure in the bottom panels of Figure A.6 that imputes all missing earnings quarters as zeros.

Figure A.12: Cost-Benefit Calculations That Include Room and Board Charges

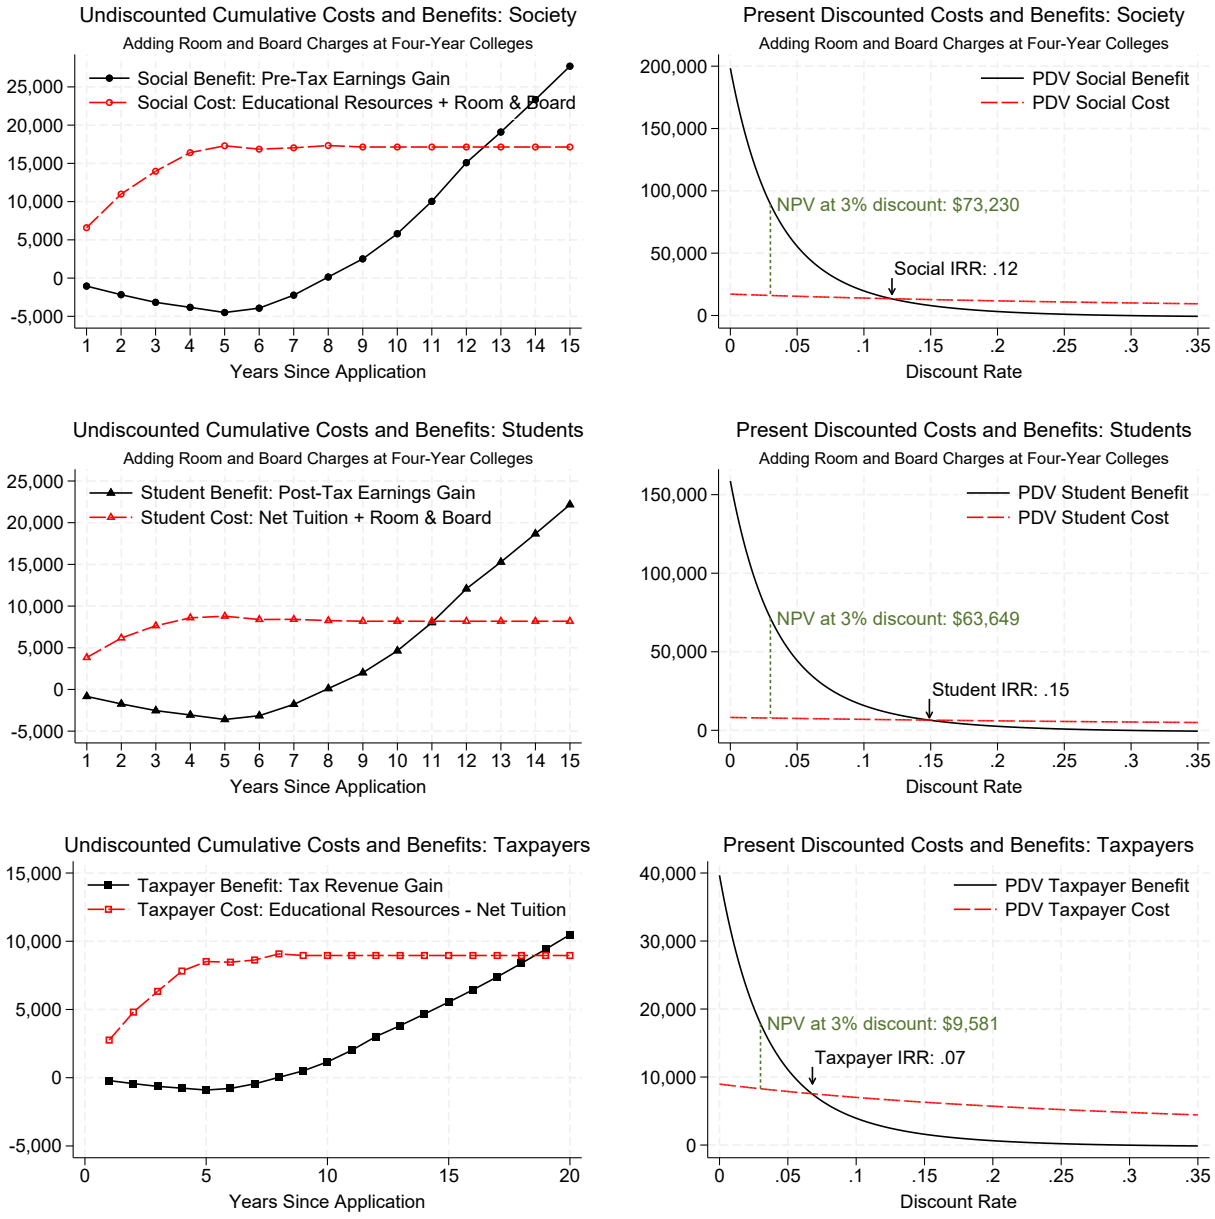

Notes: This figure is an alternative to Figure 13 that includes the additional room and board costs estimated in Figure 12, which are likely an upper bound on marginal students' actual increased room and board expenditures.

Figure A.13: Cost-Benefit Calculations Assuming a Constant Dollar Earnings Gain

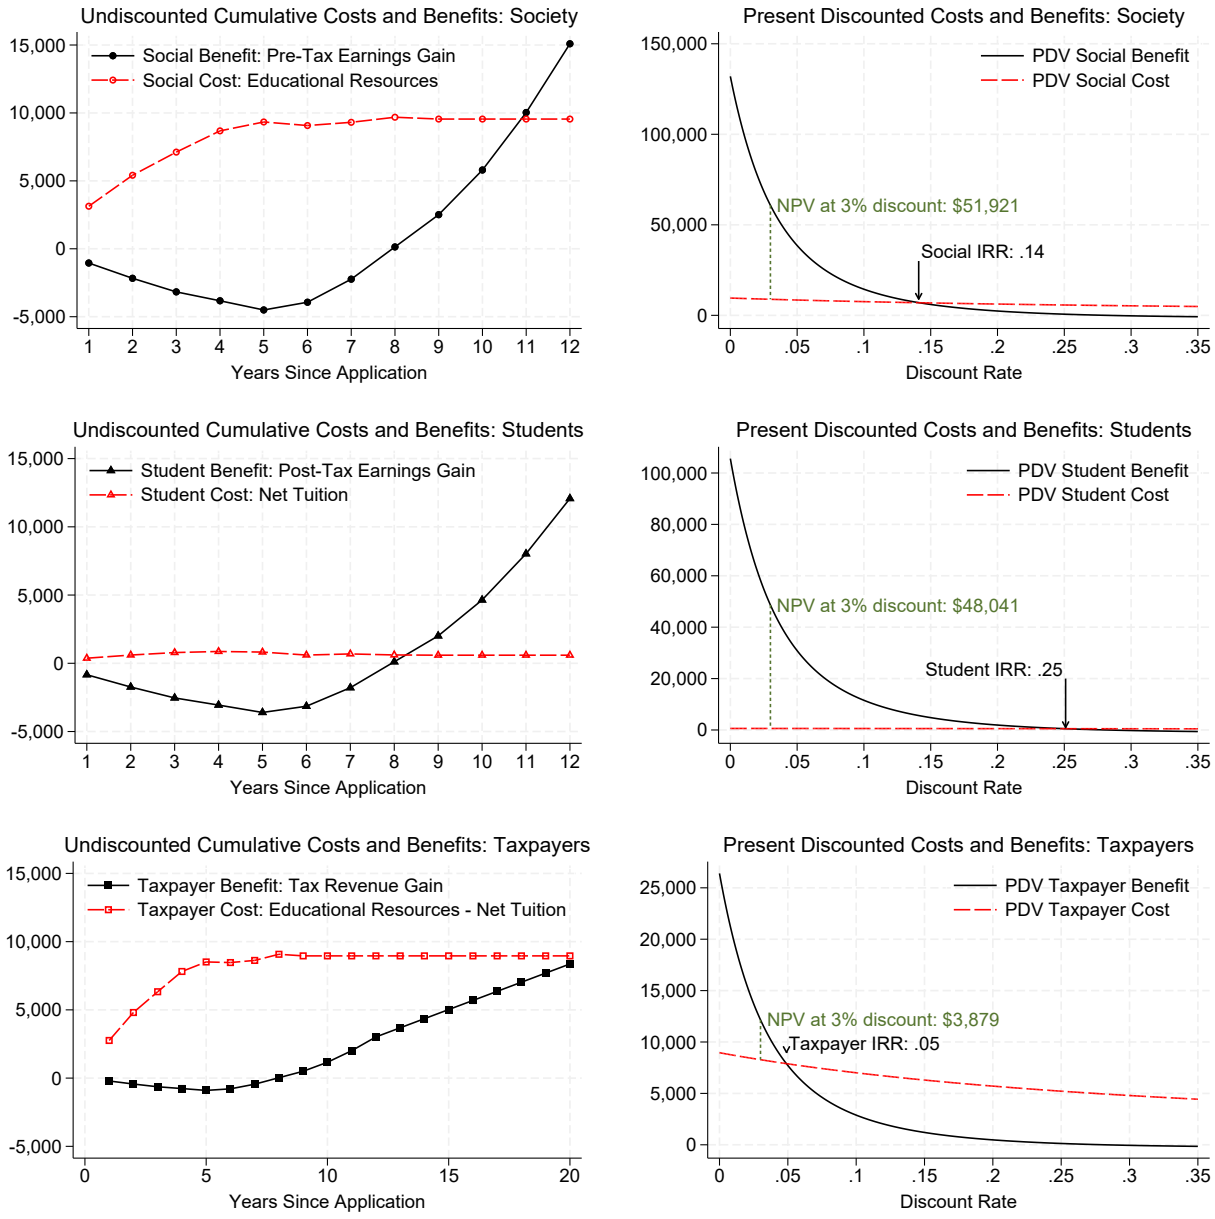

Notes: This figure is an alternative to Figure 13 that dispenses with the lifecycle extrapolation and more conservatively assumes that the annual earnings gain stays constant at $3,341 (as estimated in Figure 10) from year 13 until retirement, rather than allowing the dollar amount of the earnings gain to rise proportionally over the lifecycle.

Figure A.14: Across Institutional Spending: Effects on Earnings

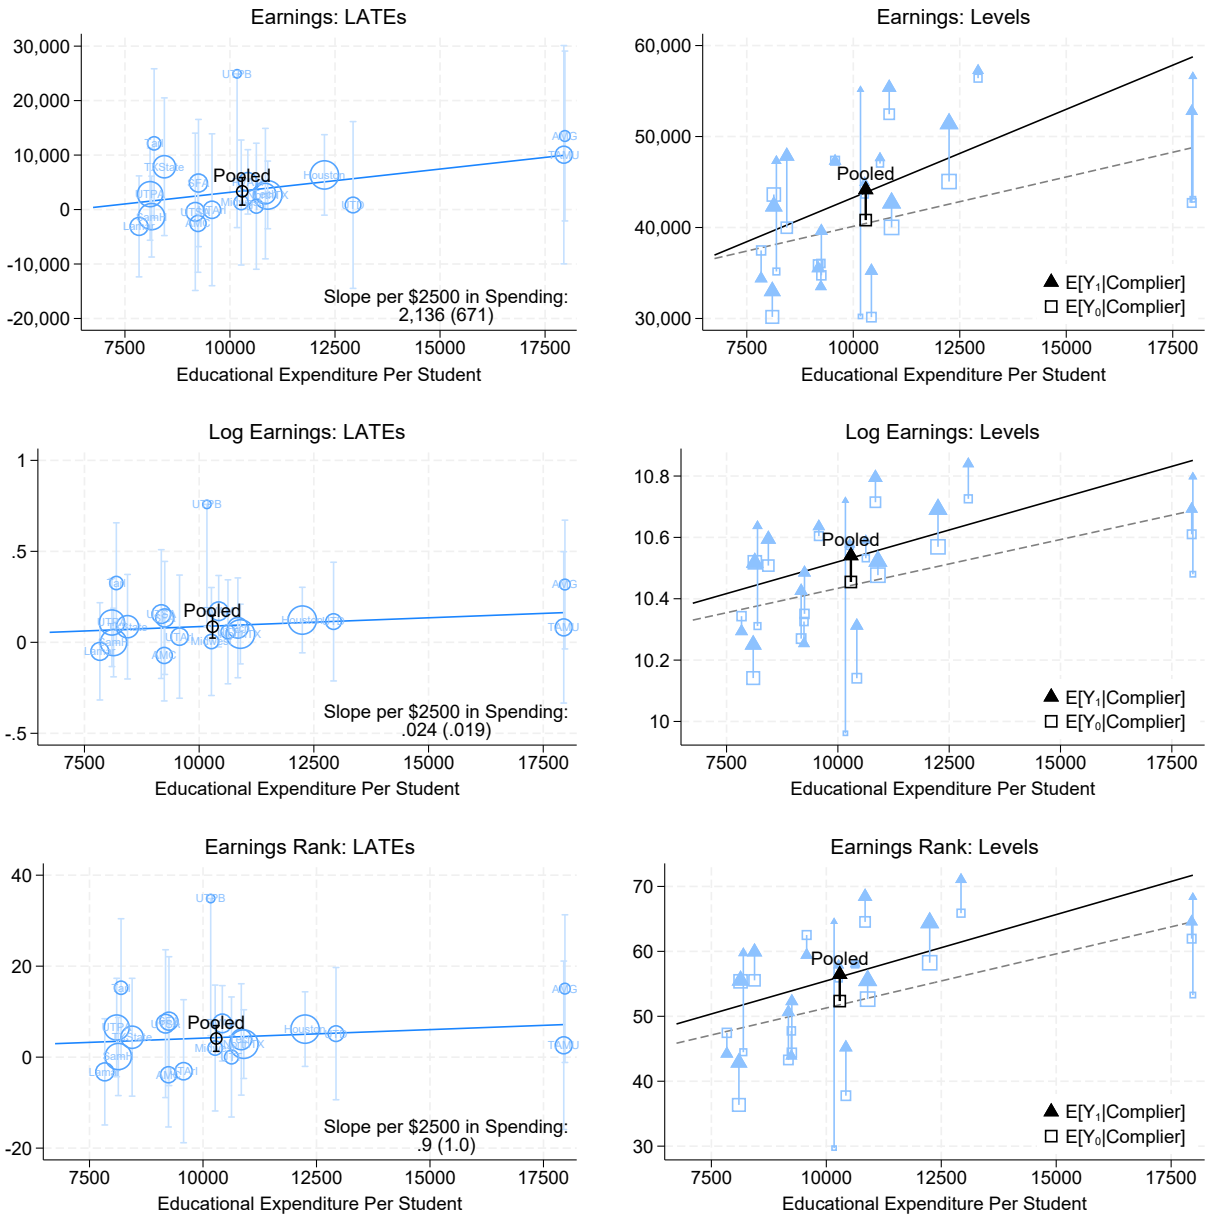

Notes: This figure has the same structure as Figure 16, but changes the horizontal axis from the institution's average SAT score to its educational expenditure per student, as defined in Section 4.1. The upper half of the confidence interval for UT-Permian Basin (UTPB) is omitted from the plot for readability, since it extends far beyond the current plot border.

Figure A.15: Across Institutions: Effects on Costs

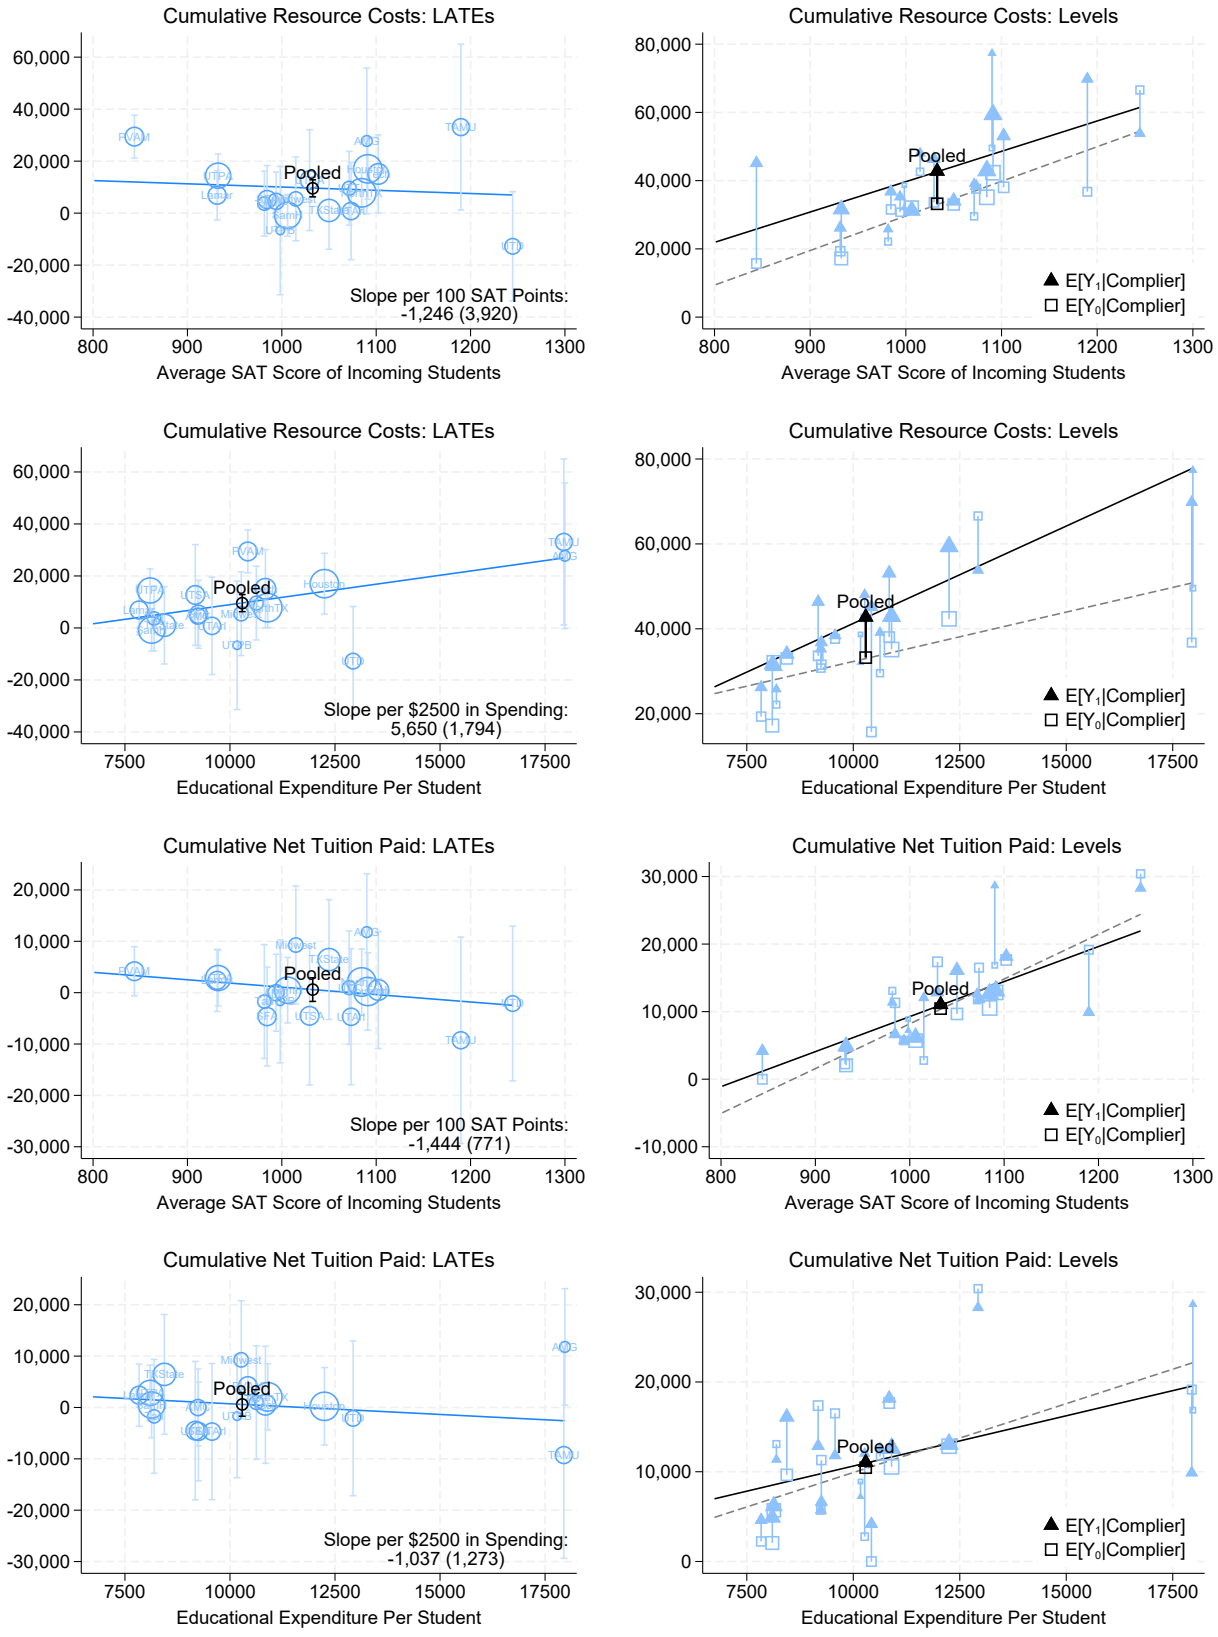

Notes: See the notes to Figures 16 and A.14.

Figure A.16: Across Student Backgrounds: Effects on Costs

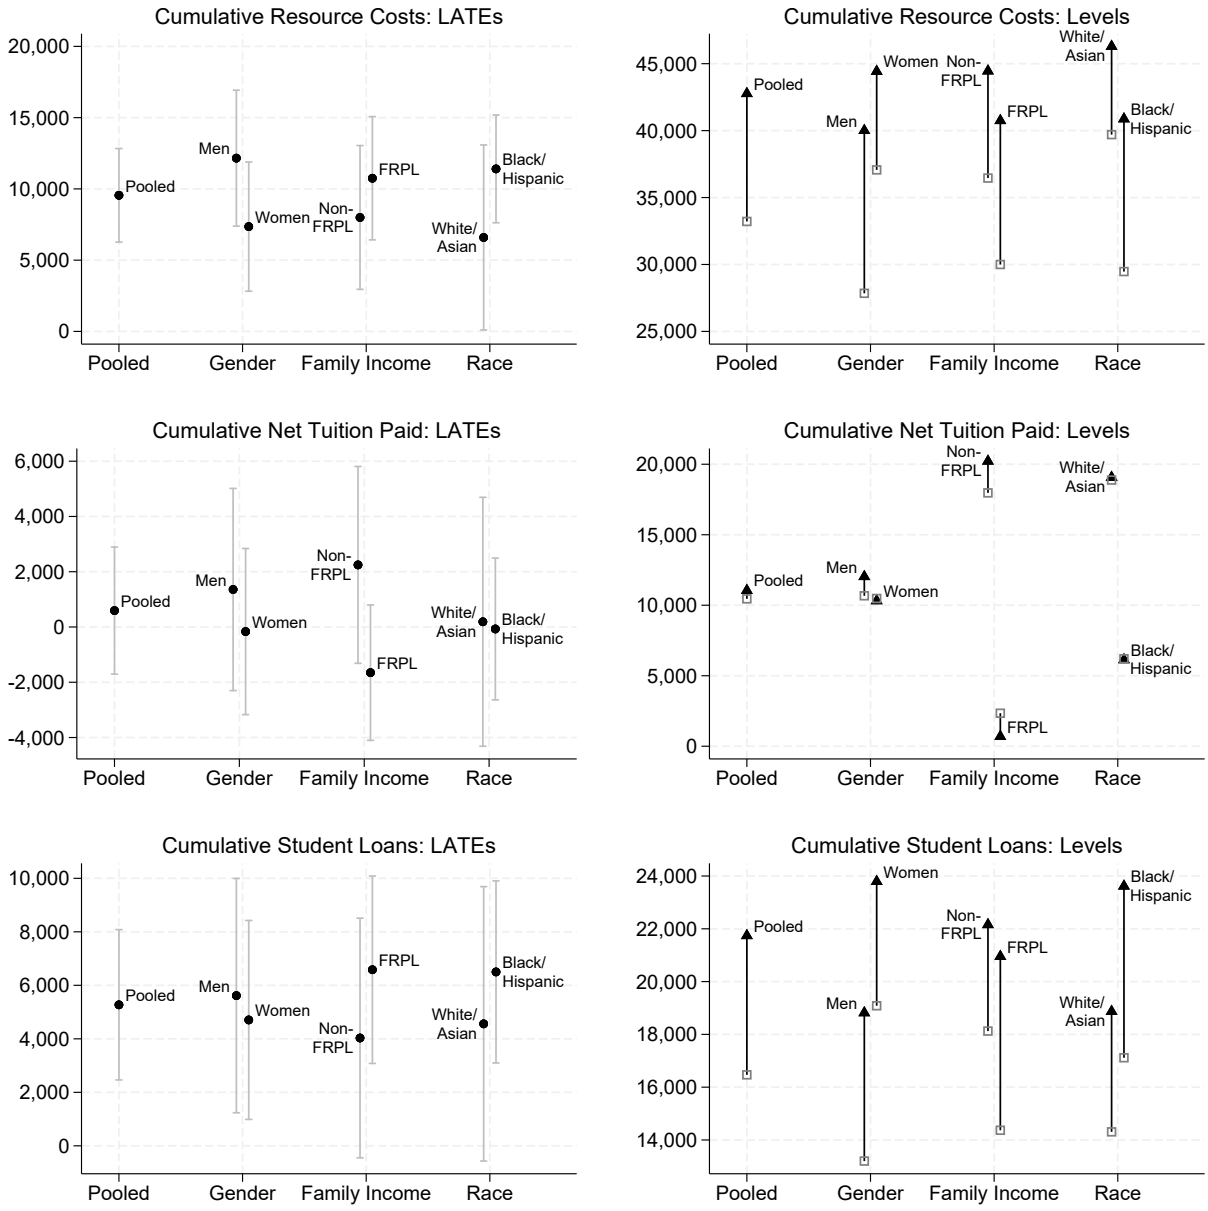

Notes: The "Pooled" estimate reproduces the pooled LATEs or potential outcomes from Section 3, while the other estimates come from running the main specification in Section 2.4 separately for each student subgroup. The "Levels" plots decompose the LATEs into the complier mean treated potential outcome level $\mathbb{E}[Y_1|\text{Complier}]$ in triangles and the complier mean untreated potential outcome level $\mathbb{E}[Y_0|\text{Complier}]$ in hollow squares. FRPL denotes eligibility for free or reduced-price lunch in high school, used as a proxy for low family income.

Figure A.17: The Other-Admissions Stratifier Is Endogenous to Cutoff-Crossing

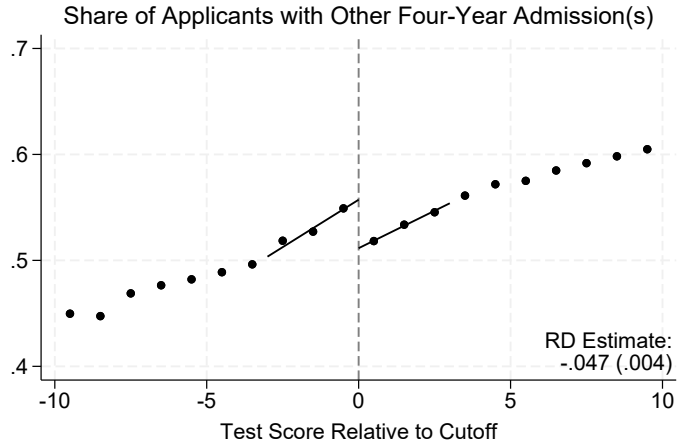

Notes: This figure plots the share of applicants at each of value of the concorded ACT running variable who have at least one other admission to a Texas public university in the year they applied to the target university. The reduced form RD estimate comes from the main RD specification described in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

Figure A.18: Disentangling the Intensive vs. Extensive Margins: Bootstrapped Confidence Intervals

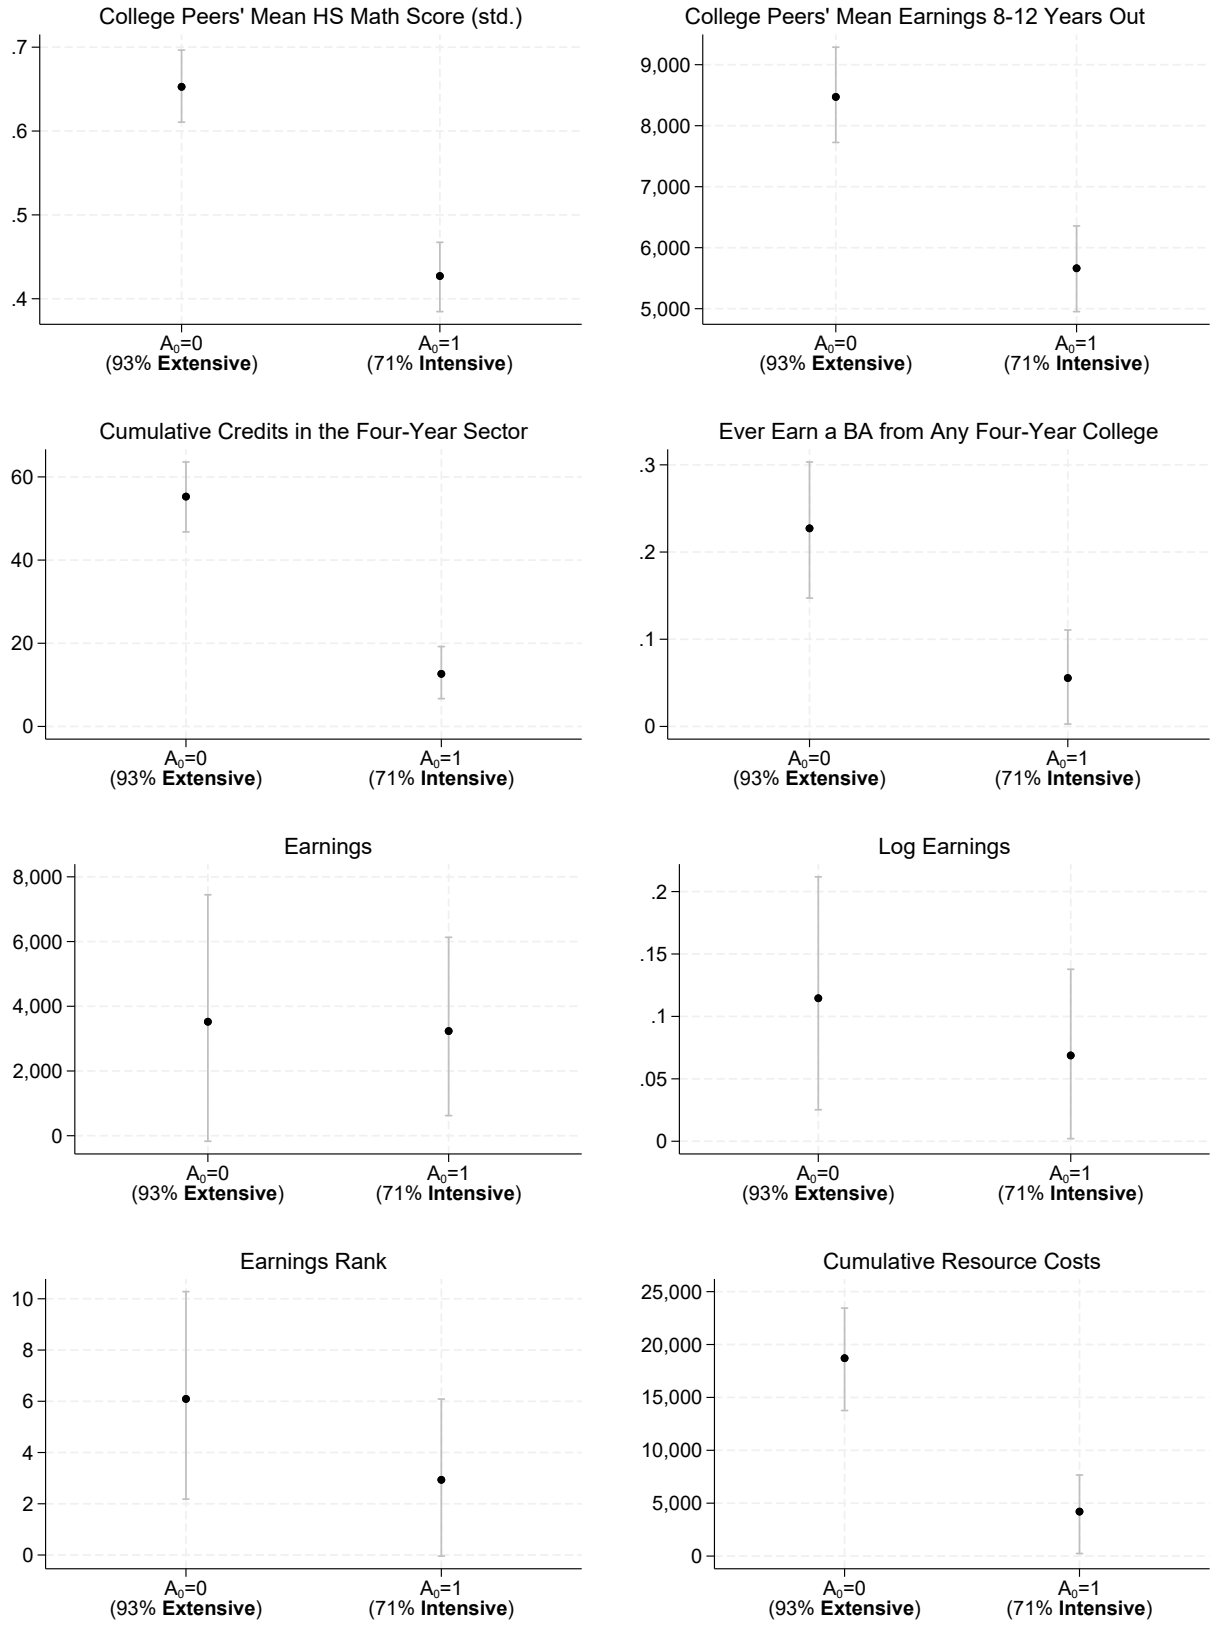

Notes: This figure plots the midpoints of the intensive versus extensive margin LATE bounds in Figure 21, along with the 90% confidence intervals that result from bootstrapping (at the applicant level) the estimation of those midpoints 1000 times.

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 23 percent of applicants to be positively top 25 and 45 percent to be non-top 25, but leaves 32 percent unknown. To classify those remaining 32 percent, I use the following pre-college covariates to estimate a logit model of top 25 status among those with known status. 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 cohort 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. Finally, I include fixed effects for high school graduation year. Using the predicted probabilities from this logit model and a cut point of 0.5, I classify 31 percent of the previously unclassified applicants as positive top 25, and the remaining 69 percent as non-top25.

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 roughly 700 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 20 percent

of the candidate cutoffs exhibit statistically significant (at the 1% level) 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. In the main results, 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.

Section 3.5 conducts a series of robustness checks across alternative ways of inferring the relevant cutoffs, including keeping the potentially publicly known cutoffs rather than dropping them, excluding cells with statistically insignificant first stage discontinuities, and allowing each cell to contribute multiple discontinuities. All of these approaches yield similar results.

The cutoffs analyzed in this paper do not appear to differ systematically by applicant race. To investigate this, I run an alternative version of the cutoff discovery algorithm in which I add race as another dimension in the definition of an application cell, such that the search algorithm runs independently within each combination of school, year, top GPA quartile status, test submitted, and race, where race is the three roughly equally sized categories of white/Asian, Black, and Hispanic. I then compare the independently discovered cutoffs across the different races within each school-year-top25-test cell. The mean difference (weighted by the overall complier mass in each cell) between the Black cutoff and the white/Asian cutoff is just .43 concorded ACT points, and the mean difference between the Hispanic cutoff and the white/Asian cutoff is just .23 concorded ACT points, both implying slightly higher cutoffs (though trivially so) for underrepresented minority applicants, and thus yielding no evidence of racial preferences in the setting of these cutoff locations.

C Identifying Separate Intensive vs. Extensive Margin LATEs

This appendix develops the identification method summarized in Section 7.1. Using the notation introduced there, expand the stratification of $A_0 = 1$ versus $A_0 = 0$ into mutually exclusive and exhaustive latent response types with respect to A:

$$\begin{aligned}
&A_0 = 0, A_1 = 0\text{: } A\text{-never-taker (never has another admission)} \\
&A_0 = 1, A_1 = 1\text{: } A\text{-always-taker (always has another admission)} \\
&A_0 = 1, A_1 = 0\text{: } A\text{-complier (secures another admission if and only if below cutoff)} \\
&A_0 = 0, A_1 = 1\text{: } A\text{-defier (secures another admission if and only if above cutoff)}
\end{aligned} \tag{3}$$

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

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

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 the proofs in Section C.1,

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

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

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

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

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

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 left panel of Figure 20 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 93 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 71 percent on the intensive margin. The shares of these response types are also identified, by the denominators of the modified fuzzy RD estimands above: 37 percent of D-compliers are A-never-takers ($A_0 = 0$), and the other 63 percent are A-compliers

(31 percent) or A-always-takers (32 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 left panel of Figure 20 are not 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, reminiscent of Manski and Pepper (2000)'s notion of monotone treatment selection. First, notice in the top left panel of Figure 20 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 top right panel of Figure 20 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, who in turn are below A-always takers.

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):

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

Assumption 3 is both empirically grounded, as shown above, and intuitive: 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 and only 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 bottom panel of Figure 20 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

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.

C.1 Proofs

First, I show that

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

The first term in the numerator is

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

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

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

so the numerator is

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

Similarly, the terms in the denominator are

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

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

which implies

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

which together with the numerator yields the desired result.

Second, I show that

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

The first term in the numerator is

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

and the second term is

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

so the numerator is

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

Similarly, the terms in the denominator are

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

and

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

which implies

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

which together with the numerator yields the desired result.

Third, I show that

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

The first term in the numerator is

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

and the second term is

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

so the numerator is

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

Similarly, the terms in the denominator are

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

and

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

which implies

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

which together with the numerator yields the desired result.

Fourth, I show that

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

The first term in the numerator is

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

and the second term is

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

so the numerator is

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

Similarly, the terms in the denominator are

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

and

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

which implies

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

which together with the numerator yields the desired result.