

Why Don't Graduation Incentives Work?

Match Quality and Financial Aid Design*

Jori Barash[†]

April 2025

Click [here](#) for the most recent version.

Abstract

Recent policy changes limit the scope for university admission decisions to equitably ration spots. I investigate whether selective universities can instead use graduation-contingent loan forgiveness to allocate spots to the students who most benefit from attendance. Identifying variation comes from an existing loan forgiveness program that incentivized greater on-time graduation. Exploiting a discontinuity from Pell Grant eligibility, I find no detectable effect of loan take-up on course load, course completion, part-time work, on-time graduation, or earnings, so participation mostly reflects selection on unobserved ability. I incorporate selection into a structural model of students' college choice, loan forgiveness take-up, and graduation. Using model estimates, I show how counterfactual loan forgiveness could target advantageous selection rather than moral hazard, leading to greater welfare, statewide graduation, and demographic equity.

*Comments from participants at UT Austin seminars are gratefully acknowledged. The research presented here uses confidential data from the State of Texas supplied by the Education Research Center (ERC) at The University of Texas at Austin. The views expressed are those of the author and should not be attributed to the ERC or any of the funders or supporting organizations mentioned herein, including The University of Texas at Austin or the State of Texas. The conclusions of this research do not reflect the opinion or official position of the Texas Education Agency, the Texas Higher Education Coordinating Board, the Texas Workforce Commission, or the State of Texas. Any errors are my own.

[†]Department of Economics, University of Texas at Austin. Email: joribarash@utexas.edu.

1 Introduction

Some universities are better than others at improving students' long-run outcomes (Dale and Krueger, 2002, 2014; Cunha, Miller and Weisburst, 2017; Chetty et al., 2022). Limited capacity at these high-quality universities presents an allocation problem. In settings with decentralized admissions like the United States, this challenge is compounded by private information. Some students may benefit more than others from attending a particular university (“high match quality”), but the university cannot fully observe students' preferences and abilities. Universities might use demographics as a signal of match quality, but recent policy changes restrict this practice.¹ When screened out of high-quality institutions, students take on debt to attend universities where they may struggle to graduate and experience a meaningful boost in earnings (Looney and Yannelis, 2024). The costs of such inefficiency may be large: much of the \$36.5 billion of annual federal spending subsidizes the enrollment of students who do not ultimately graduate (Board, 2019).

Universities screen prospective students through both admissions and financial aid, but financial aid is far less studied as a tool to improve equity and efficiency. This paper sheds light on how financial aid can improve match quality between students and colleges. I exploit variation from a novel program in Texas that incentivized on-time graduation. Conditional on a rich set of observed characteristics, participants in this program graduated on time at much higher rates. However, this apparent effect is not causal if observed characteristics do not fully capture graduation chances. With a difference-in-discontinuity design, I find no evidence that the program improved on-time graduation or proxies for short-term effort. As a result, program participation helps identify the distribution of students' expectations about graduation chances. I estimate a model of students' college choice, incentive take-up, and graduation outcomes, and I find that counterfactual graduation incentives could substantially improve welfare and equity by influencing students' choice of university.²

I use data from Texas public universities, a setting particularly well-suited to studying unobserved variation in graduation chances. First, I can measure how students with diverse backgrounds and abilities respond to financial aid. All Texas students in the top ranks of their high school by GPA are guaranteed admission to public universities. Second, comprehensive administrative records allow me to precisely measure student ability and the returns to selective universities. For the universe of Texas students,

¹For example, the Supreme Court case *Students for Fair Admissions v. Harvard* prohibited universities from considering race when determining admission.

²By contrast, the incentive program as implemented had no measurable impact on ex-ante welfare. I use a social objective that includes student welfare, universities' priority over student demographics, universities' profit, and a social externality from student graduation.

from kindergarten through early careers, I observe demographics, educational achievement, employment, and earnings. Third, the B-On-Time (“BOT”) program attracted students who expected to graduate by lowering the cost of attending college. BOT was established through an act of the state legislature in 2003 to increase timely degree attainment. Funded by a statewide tuition set-aside, the program offered interest-free loans of up to \$8000 per year to students at 4-year public universities. Virtually all Texas residents were eligible, regardless of ability, need, or program of study. BOT loans were converted to grants for students who completed their degree in four years with a 3.0 GPA (“on-time graduation”). Students who did not graduate on time could still benefit by using the program’s zero-interest loans instead of high-interest private debt. Despite BOT’s generous incentives, only 8 percent of eligible students opted in. Overall, participants were twice as likely to graduate on time than non-participants. Conditional on observed characteristics like demographics, financial need, and proxies of ability, participants were 9 percentage points more likely to graduate on time, over a mean of 25 percent.

Higher graduation among BOT participants appears to mostly reflect selection rather than reduced moral hazard. To identify the causal effect of BOT, I leverage the discontinuity in Pell Grant eligibility. Pell Grants are the largest federal aid program and discontinuously phase out with family income. Several other institutional and private financial aid programs use the same thresholds, so relative to marginally lower-income students, Pell-ineligible students typically have a \$2000 greater net cost of attendance. I apply a difference-in-RD design, comparing students just above versus just below the cutoff, in cohorts with access to BOT versus cohorts without access. The difference across cohorts controls for a time-invariant effect of discontinuously higher net cost at the cutoff. Pell ineligibility doubles the probability of BOT participation for eligible cohorts. Despite participants’ high graduation rates, I do not find evidence that BOT causally affects 4-year graduation or proxies for effort like course load, part-time work, and persistence. The null effect for on-time graduation is similar across several alternative specifications and dimensions of heterogeneity. The lack of evidence suggests that, despite greater incentives to graduate on time, the marginal BOT participant cannot exert more effort towards academic progress. Instead, I find suggestive evidence that BOT leads students to choose fields of study with higher rates of graduation, which may be less challenging. Although BOT increases eventual graduation (within six years), it fails to increase early-career earnings.

The null effects of BOT participation suggest that graduation chances vary across students at a given university, but do graduation chances also vary across universities for a given student? I document that enrollment at the two most selective public universities (the “flagships”) increases 4-year graduation by 7

percentage points, 6-year graduation by 10 percentage points, and early career annual earnings by \$2510. Point estimates vary across subsamples by family education, family income, GPA, and high school quality. By controlling for the set of universities that admit a student, I capture the admission offices' collective information about student ability (as in Dale and Krueger, 2002; Mountjoy and Hickman, 2021). This heterogeneity builds on prior evidence that the returns to selective universities vary within university by students' characteristics like income, race, gender, and high school (for a recent review, see Lovenheim and Smith, 2023).

BOT loans may have failed to increase graduation because the program targeted students' effort (moral hazard) rather than students' choice of university (advantageous selection). For example, BOT was voluntary and available at all universities, so it did not shift the *difference* in net cost between universities. One alternative is to implement mandatory graduation-contingent loan forgiveness, but only at selective universities. This alternative could help selective universities attract students with high chances of graduation by lowering these students' effective net cost of attendance. Replacing low-graduation students at selective universities with high-graduation students has ambiguous welfare effects, depending on the joint distribution of students' tastes, graduation at selective universities, and graduation at other universities. For example, a student with high graduation chances at a selective university might have similar chances at other universities. Statewide graduation can fall if this high-graduation student displaces another who has medium graduation chances at the selective university and low chances at any other school.

To explore how counterfactual graduation incentives can improve the selection of students across universities, I present a model of students' college choice, incentive take-up, and graduation. Students first choose a university, trading off the net cost of attendance, graduation likelihood, and distance.³ Next, students apply for loan forgiveness if the expected savings outweigh an idiosyncratic application cost. Finally, each student graduates according to a stochastic process. Anticipating student choices, each university designs a net cost schedule. At this stage, universities only observe students' financial need and maximize a weighted sum of profit and idiosyncratic preferences over student composition. A regulator designs subsidies and loan forgiveness to maximize expected consumer surplus plus the net social value of graduation, subject to budget, capacity, and incentive compatibility constraints. Match quality is the contribution of a student-university match to this social objective. Students may not fully value the benefits of graduation, e.g., additional tax revenue on higher earnings.

³Students at the same university face different net costs of attendance, graduation likelihoods, and distances from their high school.

Using estimates of the model parameters, I document substantial variation in flagship match quality across students, so full-information financial aid offers could lead to large welfare gains. Match quality varies due to both student preferences and idiosyncratic effects of flagship enrollment on graduation (“value-added”).⁴ Both types of heterogeneity are only partially explained by the student characteristics that universities observe. With full information, a regulator could offer each student a personalized financial aid offer that incorporates both observed characteristics and graduation chances at each university.⁵ These first-best financial aid offers increase welfare by \$27k per student without lowering revenue or increasing flagship enrollment.

Universities also play a role in the large welfare gain from full information. Their private objectives do not necessarily coincide with the social objective. I document a revealed preference for middle-income students using deviations from profit-maximizing financial aid offers. Moreover, the graduation chances of middle-income students do not depend as much on attending a selective university. Also, due to status-quo institutional constraints like staffing and state legislation, universities only consider financial need when allocating most financial aid. By considering additional observed characteristics, a regulator could increase welfare by 80 percent of the first-best gain, even without graduation incentives.

For comparison, I find that BOT did not affect student-university match quality. Most students perceived BOT application costs as too high for loan forgiveness to impact college choice.⁶ Without effects on effort or college choice, BOT could only impact welfare as an ex-post transfer to students with relatively high ex-ante graduation. BOT’s funding might be deployed more efficiently by lowering the net cost of flagship attendance for a broad set of students. Its limited efficacy might explain why the program was discontinued in 2016.

On the other hand, counterfactual loan forgiveness can improve the selection of students into flagship universities by increasing the price *difference* between flagships and alternatives for targeted students. I first combine screening on observed characteristics with a universal loan forgiveness policy comparable in size to BOT. This combined policy increases welfare by 92 percent of the first-best gain. Counter-intuitively, screening financial aid based on the full set of observed characteristics decreases graduation by 3.5 percentage points while increasing welfare. Increasing state-wide graduation requires allocating

⁴Relative to regression estimates in prior studies of heterogeneous college quality, these graduation effects rely on weaker assumptions about how students choose colleges (see, e.g., Dale and Krueger, 2002; Andrews, Li and Lovenheim, 2016; Mountjoy and Hickman, 2021).

⁵Counterfactuals vary the cost of attending a flagship university, holding other prices fixed. For every counterfactual, universities’ producer surplus increases by \$400-800 per student, consistent with institutional frictions and incentive compatibility. Full information does not include information about student’s idiosyncratic preferences.

⁶In the model, when students choose a university, BOT lowers net cost if the expected forgiven amount outweighs the expected application cost.

capacity at a selective university to students with the largest graduation value-added. However, student characteristics that increase value-added often also increase price sensitivity or weaken tastes for graduation. In these cases, it may be efficient to offer a low net cost at a less productive university. If the regulator cannot target financial aid based on student characteristics, universal graduation incentives alone increase welfare by 25 percent of the first-best gain.

Counterfactual financial aid schedules also tend to increase flagship enrollment for students from historically underrepresented backgrounds. Screening with graduation incentives increases the flagships' share of students of color, female students, first-generation students, and students from high schools with historically low college graduation rates. Consistent with cross-subsidization, average family income becomes nearly identical between flagship and non-flagship universities.

This study contributes most to the literature on price discrimination in higher education (Waldfogel, 2015; Epple et al., 2017, 2019; Epple, Martinez-Mora and Romano, 2021). These studies focus primarily on competition between colleges and inferring the colleges' objective functions from financial aid offers with third-degree price discrimination. Empirical results typically suggest that colleges target student composition, offering relative discounts to high-ability and under-represented minority students. This paper instead investigates how second and third-degree price discrimination can shift the allocation of students across colleges to maximize a social objective. The emphasis on market design is most closely related to Fillmore (2022), which considers information disclosure through FAFSA and its role in price discrimination.

A literature in public finance derives optimal income-based financial aid schedules, focusing on the extensive margin of college enrollment (Lans Bovenberg and Jacobs, 2005; Lawson, 2017; Colas, Findeisen and Sachs, 2021). I show how the intensive margin of college match quality can also generate fiscal externalities and affect the optimal schedule.⁷ To do this, I adapt intuition from the study of insurance. In selection markets, it may be efficient to price-discriminate on both ex-ante (e.g., health risk scores or test scores) and ex-post characteristics (e.g., filing a medical claim or graduation).

The null reduced-form effects of status quo graduation incentives are somewhat unusual relative to prior literature.⁸ Several papers use variation from eligibility discontinuities and typically find that grant aid affects student outcomes (van der Klaauw, 2002; Evans and Nguyen, 2019; Denning, Marx

⁷By characterizing an optimal aid schedule, I also build on related studies that decompose the welfare and composition effects of targeted financial aid (e.g., Dobbin, Barahona and Otero, 2022).

⁸There are some exceptions to the positive effects of grant aid. Rubin (2011) finds no effect on initial enrollment at the Pell eligibility threshold. Likewise, Rattini (2023) finds that aid generosity in Italy lowers attempted credit hours and increases time-to-graduation.

and Turner, 2019).⁹ Related graduation incentives often improve outcomes, e.g., universal state grants in Norway (Gunnes, Kirkebøen and Rønning, 2013), need-based privately-funded grants in Wisconsin (Goldrick-Rab et al., 2016), penalties in Germany (Mathias, Martin and Normann, 2006), and short-term incentives in the Netherlands (Leuven, Oosterbeek and van der Klaauw, 2010), but not loan tax rebates in Finland (Hämäläinen, Koerselman and Uusitalo, 2016) or relaxing exam requirements (Malacrino, Nocito and Saggio, 2024). Overall, this literature suggests that increasing the cost of late college completion sometimes incentivizes more effort during the early years of a program.

Going forward, Section 2 summarizes the process of admissions and financial aid. Section 3 uses complementary reduced-form designs to evaluate the effect of BOT loans. Section 4 presents a theoretical model of students’ college choice, loan take-up, and graduation, using it to interpret the effects of BOT. Section 5 describes the dataset and empirical model used to estimate structural parameters. This section also tests the underlying assumption that flagship enrollment benefits some students more than others. Section 6 characterizes parameter estimates. Section 7 shows how counterfactual financial aid offers affect welfare, graduation, and equity effects. Section 8 concludes.

2 Background

In the U.S., university admissions and financial aid are largely decentralized. Before students apply, each university sets the tuition and fees that compose the list price. Universities may also determine capacity, the targeted number of enrolled students in the subsequent year. At public universities, list price is lower for students from the same state. List price may vary across majors within a university, but variation in list price across universities is relatively large. Students can use online calculators to estimate financial aid for each program before applying. Applications for admission generally involve submitting essays, an intended major, and measures of academic performance like high school transcripts or test scores. Universities’ admission decisions balance academic preparedness, student interest, and in a broad sense, diversity of student composition. Anticipating that not all new students will ultimately enroll, universities admit more students than capacity.

After admitting students, universities determine financial aid largely based on financial need. Historically, financial need is measured relative to the expected family contribution (“EFC”), a function of family income and wealth. To determine their EFC, students complete the centralized Free Application

⁹See Nguyen, Kramer and Evans (2019) for a review. Not all studies use discontinuity evidence. For example, Murphy and Wyness (2022) finds that unexpected means-tested aid increases persistence, test scores, and graduation with honors.

for Federal Student Aid. Students may also immediately qualify for merit-based aid, e.g., the National Merit Scholarship or Texas Grant, but need-based aid tends to dominate. Students' initial net cost of enrollment for each program is the difference between a list price and financial aid from federal, state, and institution sources. The final net cost can vary if students apply to special programs like private scholarships. After students choose which program to enroll in, they pay their final net cost through a combination of family resources and loans.

In Texas, state regulations constrain universities' admission and financial aid decisions. Public universities must automatically admit students in the top 10 percent of their high school cohort. Since 2018, the University of Texas at Austin must admit 75 percent of students through this rule, so the effective threshold lowered to approximately 6 percent. Since 2003, Texas public universities could raise tuition above a previously enforced upper bound. However, a fixed share of incremental tuition revenue must subsidize the tuition of low-income students.

Also in 2003, Texas introduced a state loan program that provides identifying variation. The program, named B-On-Time ("BOT"), aimed to increase on-time graduation through financial incentives. BOT consisted of zero-interest loans, and the state forgave debt if students graduated within four years of first enrollment. As part of this requirement, students needed to primarily enroll in courses required for their major. All in-state students attending a public university could apply for BOT to supplement their initial financial aid offer. Unlike most state programs, BOT did not require high financial need.

In other settings, graduation incentive programs often consist of a transfer after graduation. The effects of such programs may differ with liquidity constraints, hyperbolic discounting, or loss aversion. Also, over four years, BOT loans can add up to as much as \$32,000, which may be relatively salient. BOT is also quite different than recent federal policies forgiving loans of students long after graduation. Students' choices of university enrollment might change when they anticipate loan forgiveness – both the initial enrollment and how long to work towards graduation.

3 Effects of B-On-Time Loans

B-On-Time Loans increased incentives for on-time graduation, but few students participated. This section explores whether such incentives affected student behavior. Participants were much more likely to graduate on time than non-participants with identical characteristics. However, evidence from a discontinuity in financial need suggests that the program did not increase on-time graduation for marginal

participants. Together, these facts imply that students selected into the program based on unobserved heterogeneity in predicted graduation.

3.1 Econometric Models

I apply two complementary research designs to disentangle the treatment effect of graduation incentives from Roy-style selection. The first estimating equation assumes selection on observables:

$$Y_{ijt} = \beta_1 \text{BOT}_{ijt} + \beta_2 \mathbf{X}_i + \gamma_{jt} + \epsilon_{ijt}, \quad (1)$$

where Y_{ijt} is an outcome of interest for student i , with a financial aid application at school j in fiscal year t . β_1 measures the effect of BOT take-up on Y under a strict assumption: conditional on students covariates \mathbf{X}_i , BOT take-up is uncorrelated with unobserved determinants of the outcome. This assumption may be violated, e.g., in a regression with on-time graduation as the outcome. Students with high unobserved ability may be more likely to graduate on-time, while implies greater expected loan forgiveness, leading to greater selection into BOT. In this case, the coefficient of interest would be biased upward.

The second estimating equation applies a difference-in-discontinuity design:

$$Y_{ijt} = \beta_1 \text{Above}_i + \beta_2 \text{Avail}_{jt} \text{Above}_i + \beta_3 f(\text{EFC}_i) + \beta_4 \text{Above}_i f(\text{EFC}_i) \\ + \beta_5 \text{Avail}_{jt} f(\text{EFC}_i) + \beta_6 \text{Avail}_{jt} f(\text{EFC}_i) \text{Above}_i + \epsilon_{ijt}, \quad (2)$$

where Y_{ijt} is an outcome of interest for student i , with a financial aid application at school j in fiscal year t . EFC_i is expected family contribution relative to the cutoff, the running variable. Above_i indicates Pell ineligibility. Avail_{jt} indicates years when BOT was available. Some alternative specifications include student-level covariates \mathbf{X}_i and fixed effects γ_{jt} .

The coefficient β_1 represents the discontinuity in the outcome in years when BOT was unavailable. It should not necessarily be interpreted as the effect of Pell eligibility, because potential outcomes may not be continuous at the Pell eligibility cutoff. Students in the sample have already applied and been accepted by universities. If Pell-ineligible students face systematically different costs of attendance or choice sets ex-ante, those in the sample may have greater unobserved tastes or resources for attending college. For example, if students with greater unobserved resources are more likely to graduate, then β_1 would be biased upwards.

If the discontinuity in potential outcomes is time-invariant, then the coefficient of interest β_2 is an unbiased estimate of the incremental discontinuity in the outcome in years during which BOT was available (“BOT years”) relative to when it was unavailable. Equivalently, other determinants of the outcome must not systematically vary at the cutoff in BOT years. I consider two potential violations of this assumption. First, BOT may directly crowd out other sources of financial aid. If students’ enrollment choices depend on BOT, universities could offer less grant aid and achieve similar enrollment patterns. In this case, if non-BOT aid improves outcomes, β_2 would be biased down. Second, the Pell cutoff increases over time, and BOT years are on average later than non-BOT years, so if potential outcomes vary in the level of EFC, then β_2 may reflect that difference. I cluster standard errors at the level of university-year.

Comparing estimates across research designs is suggestive of selection into BOT on unobserved graduation likelihood. For on-time graduation, the most credible estimate of BOT’s effect is the local average treatment effect from Equation (2), calculated by dividing the reduced-form estimate by the first-stage estimate. In the reduced form, on-time graduation is the outcome, and in the first stage, BOT take-up is the outcome. The average treatment on the treated from Equation (1) should be biased upwards relative to the LATE if there is substantial selection into BOT on unobserved graduation likelihood. To address the possibility that potential outcomes vary with EFC, I also estimate Equation (1) among students within a narrow bandwidth of the Pell eligibility threshold.

3.2 Data

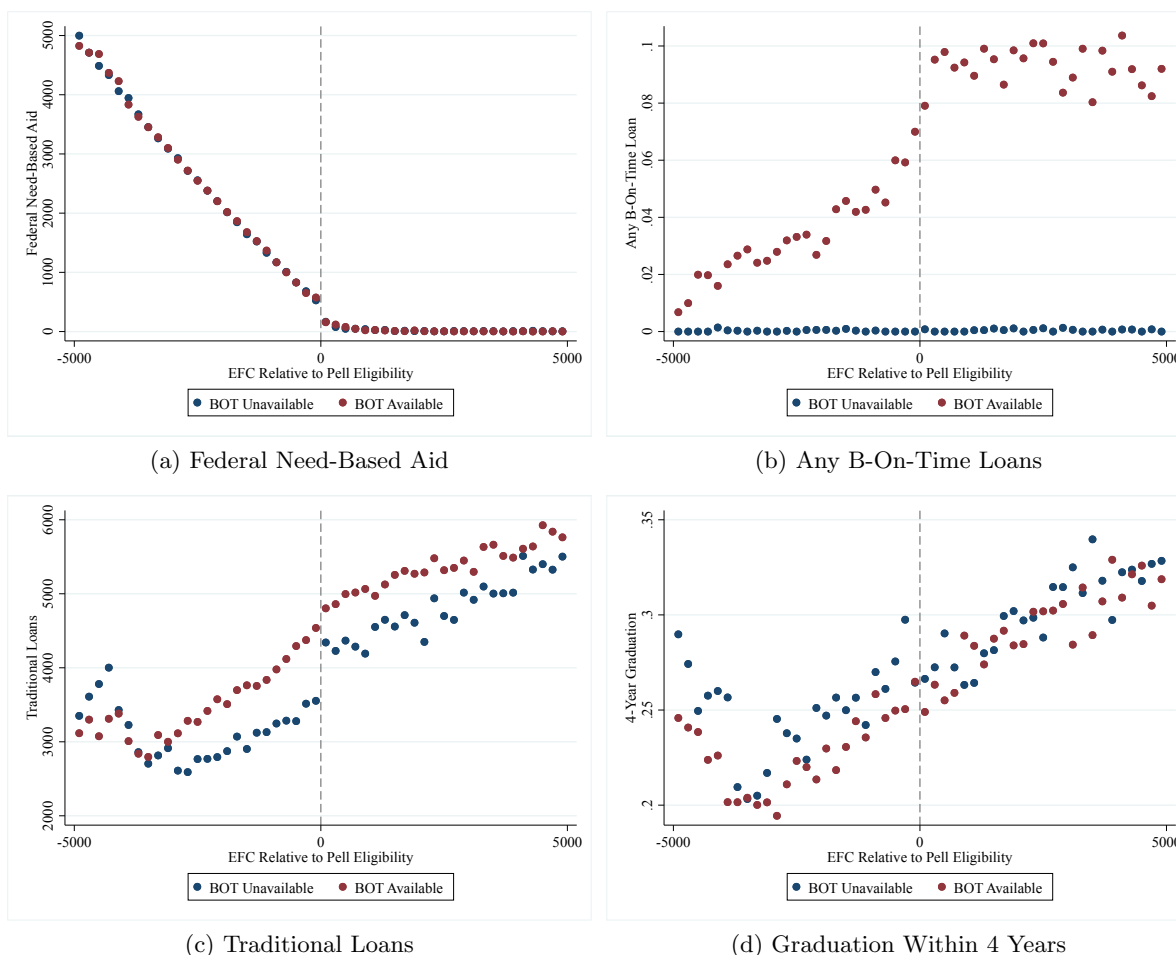
To estimate the effect of B-On-Time Loans on student outcomes, I focus on Texas high school graduates who submitted a financial aid application to a Texas public university between 2001 and 2017. Before 2017, I can measure 4-year graduation, the primary outcome of interest. Universities must be present in all years and each have at least one student with BOT loans during the sample period. The discontinuity design uses family resources as the running variable, so I remove the mass point of students who have an EFC of \$0. The Pell eligibility threshold is different for students with unusually high gross cost of attendance, so I remove these students too. To reduce measurement error, students in the sample must eventually enroll in a public 4-year university.

Figure 1 plots the discontinuities in raw mean outcomes at the Pell eligibility threshold. In both sets of years, federal aid drops off discontinuously at the threshold and generates unmet financial need (Panel A).¹⁰ Pell-ineligible students are approximately twice as likely as Pell-eligible students to take up BOT

¹⁰Federal aid mostly consists of Pell grants.

loans when BOT is available (Panel B). BOT appears to substitute for traditional loans, which vary smoothly at the cutoff when BOT is available, but increase discontinuously in other years (Panel C). Unlike other outcomes, on-time graduation varies smoothly at the cutoff for all years (Panel D). Figures A.1 and A.2 plot additional outcomes.

Figure 1: Pell Eligibility and BOT Loans



Notes: The x-axis shows expected family contribution (EFC) minus the maximum EFC for Pell Grant eligibility, which varies by year. Points represent averages of students within \$500 bins of EFC. BOT Available includes students with initial financial aid offers in years when BOT loans were widely available.

3.3 Results

Graduation incentives are associated with substantially higher graduation rates. First, Panel A of Table 1 shows estimates for on-time graduation. Regardless of specification or sample, BOT participants have 12 percentage points higher on-time graduation. Column (1) shows a raw correlation among students

that apply for financial aid. Relative to this sample's mean graduation of 26 percent, the estimate of 12 percent is almost implausibly large. To ease comparisons across research designs, Columns (2)-(3) focus on students near the Pell eligibility threshold at bandwidths of \$5000 and \$1000, for which mean graduation is 41 percent. These specifications also add controls for students' initial covariates and fixed effects for year, major, high school, and admission set. The coefficient on BOT participation is nearly identical across specifications. If unobserved selection into BOT varied by financial need, then these point estimates would likely also vary.

Estimates from the difference-in-discontinuity design suggest that BOT loans did not impact on-time graduation for marginal students. If the selection-on-observables estimate reflected a treatment effect, then on-time graduation should increase discontinuously at the Pell ineligibility threshold for BOT-eligible students relative to BOT-ineligible students. Columns (4)-(6) of Table 1 show point estimates are nearly zero. Although imprecise, confidence intervals exclude large reduced-form increases in graduation. The local average treatment effect of BOT – which normalizes these estimates by the first-stage estimates – is orders of magnitude smaller than the selection-on-observables estimate. The difference is consistent with substantial selection on unobserved gains into BOT loans. Importantly, BOT take-up is a relevant instrument. The first-stage estimate is large and precise: Pell ineligibility increases BOT take-up by 3 percentage points over a mean of 10 percent (Panel C of Table 1). This estimate is similar across specifications that add control variables or exclude observations near the cutoff.

Estimates for 6-year graduation are precise, closing the gap generated by Pell ineligibility in years when BOT is unavailable. Anticipating loan forgiveness, students may make incremental degree progress during the first four years of a program. The type of student making such progress may not be on the margin of on-time graduation. After four years of study, students may perceive a lower cost of finishing their degrees despite needing to repay BOT loans.

Similarly, Table 2 shows that although BOT impacts financial aid, there is limited evidence of reduced-form effects on students' program choice, effort, or long-run outcomes. If flagship universities increase graduation rates, then graduation incentives lower the difference in expected net cost between flagships and alternatives. Price-sensitive students might therefore respond to BOT by increasing flagship enrollment, but I do not detect an effect. Once at a university, BOT participants might choose less effort-intensive fields of study to that they can graduate on time. I find that BOT leads to choices of major with higher mean graduation. The point estimate for majors' mean earnings is positive but imprecise. Once committed to a program, BOT participants might exert more effort toward degree progress

if such effort is effective. Students might register for summer courses or refocus time towards studying, which in turn should contribute to degree progress. I do not find evidence that BOT increases attempted credit hours, decreases part-time work, or increases degree progress. Other studies find that students respond to grant aid by reducing part-time work. The lack of a precise effect here might suggest that the reduction in loans and net cost was not large enough to relax students' budget constraints. Since BOT does not seem to increase effort or on-time graduation, it is surprising that BOT increases 6-year graduation. However, there is no effect on early-career earnings. Typically, college graduation carries some signaling value that corresponds to an earnings increase. Perhaps, the 6-year effect is an artifact of choosing less difficult majors while not exerting incremental effort.

Robustness specifications validate the interpretation of the DRD coefficients: BOT does not affect graduation. First, Table 1, Column (5) shows that estimates are similar without controls for students covariates. This helps address concerns about sample selection at the Pell eligibility threshold. Second, Column (6) shows that estimates are similar when excluding observations near the threshold. Third, estimates are nearly identical to Column (5) when I do not control for students' net cost of attendance. The baseline specification controls for net cost because Table 2 shows a \$8900 lower net cost for marginal students. If net cost directly affects graduation, that effect appears to be uncorrelated with the direct effect of BOT.

Fourth, I find limited evidence that BOT directly crowds out other sources of financial aid. Universities would need to anticipate BOT receipt based on student characteristics because students apply for BOT after receiving an initial financial aid offer. Table 2 shows no evidence for incrementally higher federal aid from Pell ineligibility in BOT years. The point estimate is approximately zero and the confidence interval rules out large changes. Similarly, point estimates are small and imprecise for merit and non-discretionary aid. Small increases in the gross cost of attendance and discretionary aid are precise, but point estimates add to approximately zero, so these margins do not explain the effect on net cost. Fifth, I do not find evidence that the effect is confounded by increases in the Pell eligibility threshold over time. In alternative specifications, interactions with the cutoff level or share of students receiving BOT are imprecise.

Table 1: B-On-Time Loans Take-Up and Graduation

	(1)	(2)	(3)	(4)	(5)	(6)
(a) Four-Year Graduation						
Any B-On-Time Loans	0.120 (0.005) [<0.001]	0.119 (0.010) [<0.001]	0.112 (0.023) [<0.001]			
Pell Ineligible \times BOT Available				-0.000 (0.011) [0.996]	-0.001 (0.011) [0.932]	0.010 (0.012) [0.439]
Outcome Mean	0.255	0.408	0.407	0.243	0.255	0.254
R ²	0.002	0.176	0.304	0.089	0.074	0.075
Observations	284,039	54,247	9,792	256,979	284,039	272,689
(b) Six-Year Graduation						
Any B-On-Time Loans	0.135 (0.005) [<0.001]	0.047 (0.008) [<0.001]	0.039 (0.017) [0.023]			
Pell Ineligible \times BOT Available				0.032 (0.014) [0.027]	0.027 (0.013) [0.048]	0.035 (0.016) [0.034]
Outcome Mean	0.515	0.828	0.821	0.513	0.515	0.514
R ²	0.003	0.157	0.304	0.119	0.094	0.095
Observations	253,267	46,894	8,629	227,552	253,267	242,962
(c) Any B-On-Time Loans						
Pell Ineligible \times BOT Available				0.030 (0.009) [<0.001]	0.027 (0.008) [<0.001]	0.036 (0.010) [<0.001]
Outcome Mean				0.036	0.033	0.033
R ²				0.102	0.097	0.097
Observations	242,962	242,962	242,962	256,984	284,046	272,695

Notes: This table shows regression estimates relating B-On-Time Loans and 4-year graduation from the Selection On Observables design – Columns (1)-(3) from Equation (1) – and the Difference-in-Discontinuity Design – Columns (4)-(6) from Equation (2). Column (1) is a raw correlation. Column (2) controls for students’ initial covariates and fixed effects for year, major, high school, and admission set, among students with an EFC within \$5000 of the Pell eligibility threshold. Column (3) restricts to a bandwidth of \$1000. Column (4) is a local quadratic difference-in-discontinuity with controls for students’ initial covariates, final net cost of attendance, and fixed effects for the financial aid year, among students with an EFC within \$5000 of the Pell eligibility threshold. Column (5) only controls for the net cost and year. Column (6) is identical to (4) but excludes observations within \$200 of the cutoff. Panel names refer to the outcome variable. Panel C represents the first stage effect, so dividing by this coefficient produces the local average treatment effect of Any B-On-Time Loans on Graduation. Although BOT only rewards 4-year graduation, it also predicts overall degree completion. Panel B shows results for 6-year graduation. With controls, the raw correlation of 13 percentage points attenuates to 5 percentage points, over a mean of 83 percent.

Table 2: Difference-in-Discontinuity: Pell Ineligible \times BOT Available

	Estimate	Std. Err.	P-Value	Outcome Mean	R ²	Obs.
Financial Aid						
Federal Aid	0.006	(0.021)	[0.789]	1.661	0.837	256,984
Any B-On-Time Loans	0.030	(0.009)	[<0.001]	0.036	0.102	256,984
Traditional Loans	-0.418	(0.123)	[<0.001]	3.987	0.231	256,984
Net Cost of Attendance	-0.889	(0.196)	[<0.001]	12.115	0.609	256,984
Program Choice						
Enrolled at a Flagship	-0.001	(0.002)	[0.617]	0.169	0.933	256,984
E[Graduation Major]	0.011	(0.005)	[0.053]	0.288	0.104	256,984
Effect of Major on Earnings	0.029	(0.144)	[0.839]	-0.248	0.043	256,981
Effort Proxies						
Credit Hours Attempted	1.052	(0.974)	[0.281]	68.077	0.587	256,984
Contemporary Employment	0.003	(0.010)	[0.770]	0.759	0.039	256,984
Max Class Rank	0.031	(0.025)	[0.227]	2.301	0.671	256,796
Outcomes						
Graduation Within 4 Years	-0.000	(0.011)	[0.996]	0.243	0.089	256,979
Graduation Within 6 Years	0.032	(0.014)	[0.027]	0.513	0.119	227,552
Earnings After 8-10 Years	1.119	(0.950)	[0.240]	73.061	0.365	256,984
Balance						
Top 10% Within HS	-0.013	(0.009)	[0.159]	0.213	0.210	269,037
Top 25% Within HS	-0.023	(0.011)	[0.038]	0.402	0.140	269,037
Advanced HS Courses	0.009	(0.069)	[0.897]	3.906	0.252	284,046
White	-0.027	(0.010)	[0.009]	0.421	0.221	284,046
Free-Lunch	0.019	(0.010)	[0.054]	0.244	0.127	284,046
Discipline Days	-0.014	(0.161)	[0.931]	1.834	0.027	284,046

Notes: This table shows estimates of the difference in discontinuities at the Pell eligibility threshold in years with B-On-Time loans, i.e., β_2 from Equation (2). The specification corresponds to Column (4) in Table 1: a local quadratic difference-in-discontinuity with controls for students' initial covariates, final net cost of attendance, and fixed effects for the financial aid year, among students with an EFC within \$5000 of the Pell eligibility threshold. All monetary measures are scaled down by 1000.

4 Theoretical Framework

4.1 Model

I model the choices of students, universities, and a regulator to rationalize the effects of BOT loans and quantify welfare and graduation under counterfactual schedules. The regulator first sets a policy, then universities set financial aid schedules, then students choose a university, and then students exert costly effort toward degree completion and human capital accumulation. I describe these components in reverse order.

THE STUDENT. Student i is characterized by characteristics at the time of college application, $\mathcal{X} \equiv \{X, h_0\}$, and preferences over university characteristics θ_i . X consists of characteristics observed by the econometrician such as demographics and measures of ability. h_0 is the stock of human capital.

After college, students' earnings increase in their human capital stock and graduation status, which is used as a signal by employers. During college, anticipating future earnings, students repeatedly choose how much costly effort to exert studying. Effort increases human capital accumulation and the probability and speed of graduation. Both outcomes may vary across universities based on educational quality and accessibility. Effort below a threshold is equivalent to dropping out and forgoing graduation to enter the labor market. Effort and part-time paid work are substitutes and bounded above by a time constraint. Students value current consumption and leisure as well as the present discounted value of future earnings. Each semester, idiosyncratic shocks to income and demands for time can influence students' choices. The sequence of shocks induce an ex-ante probability of graduation g_{ij} . For example, small changes to family income can discontinuously change financial aid, making attendance unaffordable. Alternatively, the combination of current human capital and local labor market conditions creates variation in the value of full-time work without degree completion relative to current college enrollment.

At the start of college, after choosing a program, students can pay a cost to accept graduation incentives. The random idiosyncratic portion of this cost is realized after choosing a program.¹¹ To help smooth future consumption, students accept graduation incentives if the incentives lower the expected net cost of attendance. At this stage, students anticipate the probability of graduation. I assume that students are risk-neutral.

Students choose the university program that maximizes expected indirect utility. Indirect utility is a linear function of the program's mean quality, net cost, distance from family, the continuation

¹¹With BOT loans, students choose to accept graduation incentives. In contrast, counterfactual loan forgiveness may be universal, so the upfront cost is zero.

value of enrollment, and idiosyncratic taste shocks. Preferences over program characteristics can vary across students. Students maximize indirect utility in expectation because of uncertainty about future choices. First, net cost depends on a student’s anticipated choice of BOT take-up and future graduation. At this point, the idiosyncratic BOT application cost is not yet realized. Second, the continuation value of enrollment depends on a student’s anticipated choices of costly effort. Counterfactuals may understate changes to indirect utility if students value other time-varying program characteristics, e.g., the demographics of other enrolled students.

THE UNIVERSITY. University j is characterized by its preferences over student characteristics, θ_j , capacity Q_j , per-student budget R_j , and two production functions, human capital $f_{h,j}$ and degree completion $f_{g,j}$. At the time of financial aid determination, the university only observes a subset of characteristics $X_0 \in \mathcal{X}$. Anticipating students’ enrollment choices, the university chooses a financial aid schedule to maximize the expected sum of revenue and valuations of enrolled students’ characteristics, subject to budget and capacity constraints. Revenue equals the student-paid net cost plus federal and state subsidies. Marginal valuations may be higher for students with underrepresented demographics if universities seek to maintain diversity. Diversity might be valued due to its downstream effects on reputation and educational quality. Universities cannot increase enrollment beyond capacity without substantial fixed costs, so financial aid helps ration spots among admitted students.¹²

Generally, a financial aid schedule P_j maps student characteristics and outcomes to a student-specific net cost of attendance for each university j . In the status quo, financial aid schedules depend only on university-observed X_0 , i.e. expected family contribution. Students can apply for small additional scholarships, but at the time of university choice, they anticipate paying the net cost from the schedule.

THE REGULATOR. The regulator values a weighted sum of students’ consumer surplus, universities’ producer surplus, and externalities of higher education, in expectation. Subject to a statewide budget constraint, regulation can affect students’ net cost through restrictions on list price, in-kind tuition subsidies, or direct transfers to students. Through such regulation, a financial aid schedule might condition on a broad set of student characteristics – i.e., those in administrative records – as well as graduation outcomes.

Generally, an equilibrium is characterized by a set of financial aid schedules that clears the market subject to subsidies, regulation, constraints, and students’ privately optimal choices of enrollment. To illustrate the effects of loan forgiveness, I simplify the problem: the regulator directly sets financial aid

¹²In Texas, flagship university capacity is regulated.

schedules, valuing consumer surplus and externalities, subject to universities' constraints and an added participation constraint. In expectation, each university must weakly prefer a counterfactual financial aid schedule to the status quo. The regulator may be able to improve producer surplus because institutional constraints limit universities' information and which financial aid schedules are feasible.

4.2 Discussion

The model clarifies that BOT loans might have limited effects on proxies for effort because the effort of marginal participants is constrained at an upper bound. Otherwise, BOT should increase effort through two paths. First, BOT increases the value of on-time graduation, which raises the marginal benefit of effort above the marginal cost. Second, BOT lowers the expected net cost of college, so students can consume the same amount with less paid work, and this frees up time for educational effort and leisure.

As a caveat, the estimated null effects of BOT loans on effort are local to students at the threshold of Pell eligibility. I interpret the constraint on effort as binding for students with other levels of expected family contribution ("EFC") because proxies for effort like attempted credit hours are similar on average.¹³ Below, when I analyze the impacts of counterfactual financial aid offers, I continue to assume that BOT has no causal effect on graduation among a sample of students with high observed effort. This assumption implies that financial aid does not affect students' upper bound for effort.¹⁴

Even if graduation incentives do not change effort (reduce moral hazard), they can induce students with high graduation chances to choose the flagship university (advantageous selection). As a benchmark, suppose flagship financial aid optimally depends on all observed student characteristics – when would adding a graduation incentive improve selection? Unlike BOT, in this alternative program design, *all* flagship students face a baseline net cost and a fixed reward for graduation, and *no* non-flagship students have a graduation incentive. If budget-neutral, this program lowers the flagship net cost for students with high flagship graduation chances and increases the flagship cost for others.¹⁵ The wedge in net cost replaces low-graduation students with high-graduation students at the flagship.

The aggregate impacts of these alternative graduation incentives are theoretically ambiguous. For

¹³To make this comparison, I also assume that EFC does not vary systematically with students' upper bound on effort or marginal cost of effort. The exception is that students with very low EFC have lower attempted credit hours (less effort), but also lower contemporary employment (more time remaining for effort). This combination could be rationalized by a greater cost of effort.

¹⁴Policies outside the scope of this paper might allow financial aid to increase graduation chances for a broad set of students. Improving the initial stock of human capital or mitigating students' external constraints could increase the upper bound or cost of effort.

¹⁵For simplicity, fix other schools' net costs. The average net cost of the flagship does not change, so other schools have, at most, a small incentive to change prices in a Bertrand-Nash equilibrium. All else equal, if students displaced to these schools have relatively low price sensitivity, the prices at these schools will rise, which amplifies aggregate effects.

example, if each student has the same graduation chance at all schools, then the reallocation does not impact statewide graduation. Otherwise, attending the flagship has a causal effect on graduation (“value-added”), which may vary across students. Statewide graduation only increases if the total gain in value-added among students marginally enrolled at the flagship outweighs the value-lost among displaced students. Students who expect to graduate at high rates at the flagship need to also have relatively large graduation value-added. Moreover, average student welfare can fall as average graduation rises. For example, suppose the marginal utility of graduation is steeply decreasing in graduation. Then, students marginally enrolled at the flagship receive a small gain from low net cost, and displaced students receive a large loss from high net cost and low graduation.

5 Estimation

5.1 Data

I focus on students admitted to at least one flagship university and one other public university. Students in the estimation sample both graduated from a Texas high school in the top 10 percent of their class and submitted a FAFSA between 2001 and 2017. Relative to the sample used in Section 3, this estimation sample includes greater variation in family income and relatively high-achievement students. I observe one financial aid offer per student per year, typically for the university at which they first enroll. I use the first observed offer and calculate the net cost as the cost of attendance minus all forms of grant aid, scholarship aid, and work-study funding. Cost of attendance is typically greater than tuition due to estimated costs of living like housing. Cost of attendance may vary across students within a school-year due to state rules governing tuition discounts which are separate from aid.¹⁶ I predict net cost for all students by regressing observed net cost on interactions of EFC bins and continuous EFC percentile, separately for each year.¹⁷ I cap net cost for students in the top bin of EFC. I measure miles between each student’s K-12 district and each university. I measure college graduation for public universities and compare the college graduation year to the high school graduation year.

¹⁶I impute cost of attendance using the school-year median for some students, bound net cost between \$0 and \$100,000, and impute missing values of net cost as \$0.

¹⁷I group students into 31 bins; the lowest bin includes those with an EFC of \$0 and the remainder are binned by the within-year EFC percentile.

Table 3: Mean Estimates of Value-Added

	(1)	(2)	(3)	(4)
(a) Four-Year Graduation				
Enroll Flagship	0.130 (0.003) [<0.001]	0.064 (0.003) [<0.001]	0.070 (0.007) [<0.001]	0.070 (0.007) [<0.001]
R ²	0.048	0.155	0.157	0.157
Outcome Mean	0.414	0.414	0.414	0.414
Observations	94,839	94,839	94,839	94,839
(b) Six-Year Graduation				
Enroll Flagship	0.174 (0.003) [<0.001]	0.097 (0.003) [<0.001]	0.101 (0.005) [<0.001]	0.102 (0.005) [<0.001]
R ²	0.056	0.197	0.217	0.217
Outcome Mean	0.703	0.703	0.704	0.704
Observations	77,705	77,705	77,705	77,705
(c) Annual Earnings 8-10 Years Post-HS (Thousands)				
Enroll Flagship	2.191 (0.314) [<0.001]	2.621 (0.319) [<0.001]	2.454 (0.486) [<0.001]	2.510 (0.484) [<0.001]
R ²	0.025	0.033	0.066	0.067
Outcome Mean	54.587	54.587	54.680	54.680
Observations	56,833	56,833	56,833	56,833

Notes: This table shows the effects of flagship enrollment on medium-run outcomes, with estimates from Equations (3). Column (1) includes basic initial student characteristics observed by every university and a fixed effect for year of enrollment. Column (2) adds average graduation rates conditional on each of high school, major, and admission portfolio. Column (3) adds an admission portfolio fixed effect, which absorbs the corresponding mean graduation rate, and clusters standard errors by admission portfolio. Column (4) adds a control for distance to the nearest flagship.

5.2 Heterogeneity in Value-Added

To motivate key features of the structural model, I test whether the returns to selective universities vary across students. Consistent with prior studies, average returns to flagship enrollment are large for graduation and earnings. Effects vary substantially across subsamples of student demographics. For example, flagship enrollment most increases 4-year graduation for first-generation students, free-lunch-eligible students, and students of color.

I estimate the value-added of flagship universities as the difference in outcomes among students with similar ability and choice sets, following Mountjoy and Hickman (2021) and Dale and Krueger (2002):

$$Y_{ij} = \beta_1 1[\text{Enrolled Flagship}] + \gamma_{a(i)} + \beta_2 \mathbf{X}_i + \epsilon_{ij}, \quad (3)$$

Table 4: Heterogeneity in Value-Added

	HS \geq 60 miles		High-Graduation HS		First-Generation	
	(N)	(Y)	(N)	(Y)	(N)	(Y)
(a) Four-Year Graduation						
Enroll Flagship	0.071 (0.009) [<0.001]	0.070 (0.006) [<0.001]	0.076 (0.007) [<0.001]	0.063 (0.008) [<0.001]	0.062 (0.007) [<0.001]	0.078 (0.008) [<0.001]
Outcome Mean	0.419	0.410	0.337	0.492	0.441	0.387
R ²	0.152	0.172	0.141	0.143	0.160	0.162
Observations	48,269	45,085	46,621	46,695	48,303	44,887
(b) Six-Year Graduation						
Enroll Flagship	0.111 (0.007) [<0.001]	0.094 (0.006) [<0.001]	0.092 (0.006) [<0.001]	0.112 (0.007) [<0.001]	0.113 (0.006) [<0.001]	0.091 (0.007) [<0.001]
Outcome Mean	0.709	0.698	0.648	0.761	0.738	0.669
R ²	0.225	0.219	0.197	0.233	0.242	0.203
Observations	39,265	36,975	38,857	37,383	39,008	37,094
(c) Annual Earnings 8-10 Years Post-HS (Thousands)						
Enroll Flagship	3.867 (0.530) [<0.001]	1.135 (0.601) [0.059]	1.798 (0.477) [<0.001]	3.122 (0.818) [<0.001]	2.971 (0.818) [<0.001]	1.896 (0.442) [<0.001]
Outcome Mean	55.239	54.159	51.240	58.556	57.081	52.439
R ²	0.074	0.077	0.074	0.057	0.066	0.082
Observations	28,818	26,680	29,176	26,297	27,459	27,877

Notes: This table shows heterogeneity in the effects of flagship enrollment on medium-run outcomes, with estimates from Equations (3). The specification corresponds to Column (3) from Table 3. It includes controls for basic initial characteristics, fixed effects for the admission portfolio and year of enrollment, and average graduation by high school and major. Each column represents a separate regression on a subsample of students.

where Y_{ij} is an outcome of interest for student i with initial enrollment at school j . The key assumption is that, conditional on observed characteristics, students symmetrically sort on unobserved characteristics into a flagship university or an alternative. In that case, β_1 represents the effect of initial flagship enrollment on the outcome. $\gamma_{a(i)}$ is a fixed effect for the set of universities that admitted student i , which may capture both students' and universities' private information about unobserved ability through application and admission decisions. I estimate this equation on the full sample before comparing estimates across subsamples.

First, Table 3 shows a large average treatment effect of flagship enrollment on 4-year graduation (Panel A), 6-year graduation (Panel B), and early-career earnings. Column (1) only controls for basic student characteristics that are observed by the university such as financial need and test scores. Columns (2)-(4)

Table 5: Heterogeneity in Value-Added: Additional Characteristics

	Top 10%		Free-Lunch		White	
	(N)	(Y)	(N)	(Y)	(N)	(Y)
(a) Four-Year Graduation						
Enroll Flagship	0.073 (0.008) [<0.001]	0.070 (0.009) [<0.001]	0.067 (0.007) [<0.001]	0.078 (0.008) [<0.001]	0.070 (0.009) [<0.001]	0.068 (0.007) [<0.001]
Outcome Mean	0.414	0.416	0.437	0.350	0.390	0.441
R ²	0.149	0.183	0.151	0.183	0.173	0.148
Observations	46,211	47,019	69,872	23,306	48,069	45,364
(b) Six-Year Graduation						
Enroll Flagship	0.111 (0.007) [<0.001]	0.093 (0.006) [<0.001]	0.107 (0.006) [<0.001]	0.085 (0.008) [<0.001]	0.088 (0.007) [<0.001]	0.113 (0.006) [<0.001]
Outcome Mean	0.701	0.707	0.725	0.639	0.675	0.732
R ²	0.200	0.258	0.222	0.217	0.218	0.225
Observations	39,350	36,787	57,912	18,196	37,579	38,743
(c) Annual Earnings 8-10 Years Post-HS (Thousands)						
Enroll Flagship	2.541 (0.488) [<0.001]	2.631 (0.659) [<0.001]	2.571 (0.583) [<0.001]	2.233 (0.735) [0.002]	1.799 (0.665) [0.007]	2.997 (0.797) [<0.001]
Outcome Mean	54.861	54.602	56.607	48.755	51.210	57.909
R ²	0.068	0.084	0.060	0.105	0.075	0.066
Observations	29,681	25,705	42,343	13,018	26,388	29,137

Notes: This table shows heterogeneity in the effects of flagship enrollment on medium-run outcomes, with estimates from Equations (3). The specification corresponds to Column (3) from Table 3. It includes controls for basic initial characteristics, fixed effects for the admission portfolio and year of enrollment, and average graduation by high school and major. Each column represents a separate regression on a subsample of students.

add covariates that are not directly observed by the university, which attenuates the coefficient of interest by almost half. This finding is consistent with prior work suggesting students select a university based on their unobserved ability. Estimates are still large: flagship enrollment increases 4-year graduation by 6.4 percentage points over a mean of 41.4 percent; this drives three-quarters of the increase in 6-year graduation; and earnings increase by \$2,621 over a mean of \$54,587. Universities do not observe students' other admissions which proxy ability through selection. Column (3) takes the conventional approach of an admission set fixed effect, while Column (2) instead controls for mean graduation rate by admission set. Both approaches result in similar estimates, which helps reduce the parameter set for estimating the structural model below. In prior studies, estimates are similar when using high-dimensional fixed effects or a low-dimensional approximation (Dale and Krueger, 2002; Mountjoy and Hickman, 2021). Column

(4) adds distance as a control, producing similar estimates and nearly identical R^2 , easing concerns about remaining unobserved variation. Across several settings, distance shifts college enrollment and studies frequently assume that distance is independent of ability.

Second, Tables 4 and 5 show that returns to flagship enrollment vary across student characteristics, particularly relative to mean outcomes. For example, I compare students by the average 4-year college graduations rates of their high school. Students from schools with below-median graduation have a larger point estimate of the effect of flagship enrollment on their own 4-year graduation: 7.6 percent versus 6.3 percent. As a share of mean outcomes, the estimate for high-graduation high schools is nearly twice as large: 23 percent versus 12 percent. The relative returns are similar for 6-year graduation and much smaller for earnings (3.5 percent versus 5.3 percent). For other subsamples too, larger relative returns on 4-year graduation do not necessarily translate to larger returns for 6-year graduation or earnings. First-generation students, free-lunch-eligible students, and students of color all have relatively large returns to 4-year graduation. On the other hand, I do not find evidence of heterogeneity by distance or high school rank. Heterogeneity in returns to education is broadly consistent with the literature. For example, Andrews, Li and Lovenheim (2016) use quantile regression to estimate the returns to Texas flagship enrollment. They find a dispersed distribution of earnings effects across students with heterogeneity by race and ethnicity.¹⁸

5.3 Empirical Model

I estimate preferences and graduation value-added by maximizing the joint likelihood of the college choice, BOT take-up, and the probability of graduation:

$$l_i(Y_i, B_i, G_i, X_i) = \int_{g_{f,i}} Pr(G_i | X_i, Y_i) Pr(B_i | X_i, Y_i, g_{f,i}) Pr(Y_i | X_i, g_{f,i}) dF(g_{f,i} | X_i)$$

where G_i indicates observed graduation within 4 years for student i , g is the latent ex-ante probability of on-time graduation, X_i includes observed characteristics (e.g., demographics, financial need, and the admission set), Y_i indicates the choice of university, and B_i indicates take-up of BOT loans. Students can choose UT Austin, Texas A&M, or an outside option.¹⁹

The key timing assumption is that students do not observe the exact cost of applying for BOT until after university enrollment. The university choice incorporates expectations over the distribution of BOT

¹⁸At UT Austin, the earnings premium ranges from 2.7 percent at the 9th percentile to 31.7 percent at the 97th percentile.

¹⁹Some students are not admitted to both UT Austin and Texas A&M. $f \in \{0, 1\}$ indicates enrollment at a flagship university.

application cost. Since effort is largely unobserved, I also simplify the broader theoretical framework by parameterizing the continuation value of enrollment as linear in ex-ante graduation probability $g_{f,i}$. As a slight abuse of notation, I refer to $g_{f,i}$ as a student-university characteristic. Graduation probabilities vary flexibly depending on whether students enroll at a flagship university. This flexibility reflects that the underlying production functions may vary with university quality. I parameterize $g_{f,i}$ as a logistic function of student characteristics and a normal random variable $\epsilon_{f,i}^g$.

$$g_{f(j),i} = \Phi\left(\beta_x^g X_i + \epsilon_f^g\right)$$

$$Pr(G_i | X_i, f_i) = G_i g_{f,i} + (1 - G_i) (1 - g_{f,i})$$

Students take up BOT loans if the expected reduction in net cost of attendance outweighs an idiosyncratic application cost, which is a function of student characteristics and a normally distributed error. Likewise, the expected net cost is reduced by expected loan forgiveness.

$$Pr(B_{ijkt} = 1) = Pr(g_{f(j),i} \bar{B} - a_{ijkt} > 0), a_{ijkt} \sim N(a_{jkt}, \sigma_B)$$

$$E[P_{ijkt}^B | g] = \max\{0, a_{jkt} - g_{f(j),i} \bar{B}\}$$

Students choose a college to maximize indirect utility, a function of net cost of attendance, graduation likelihood, residual mean quality, and a taste shock:

$$u_{ij} = -\alpha_i P_{ij} + \beta_{g,i}^u g_{ij} + \beta_{d,i}^u d_{ij} + \delta_{ij} + \epsilon_{ij}^u$$

$$u_{i0} = -\alpha_i \bar{P}_{i0} + \beta_{g,i}^u g_{i0} + f(|\mathcal{J}_0|) + \epsilon_{i0}^u$$

$$P_{ij} = P_{ij}^0 - E[P_{ijkt}^B | g]$$

$$\theta_i = f_\theta(\Pi_\theta X_i + \epsilon_\theta), \theta \in \{\alpha, \beta_g, \beta_d, \delta\}$$

I normalize the outside option intercept δ_{i0} to 0. All else equal, students who are admitted to more non-flagship universities may have a stronger preference for the outside option. To address this, I include a quadratic function of the count of other admissions. $\beta_{g,i}$ and α_i are log transformations of the linear index because indirect utility should not increase in net cost or decrease in graduation. I allow for correlation between ϵ_0^g and ϵ_1^g . The preference shocks are distributed Type-I extreme value. I integrate over ϵ^g using quadrature to calculate the likelihood.

Based on reduced-form findings, I assume that BOT has no direct impact on graduation. As a result, observed selection into BOT – among students with identical observed characteristics – identifies the variance in ex-ante graduation likelihood. Conditioning on BOT take-up and university choice, residual variation in graduation rates identifies $E[g_{f(j)}]$. Residual correlation between BOT take-up and observed characteristics identifies mean BOT application cost. The correlation between selection into flagship universities and the corresponding ex-ante graduation probability g_1 separately identifies $\beta_{g,ik}^u$ and δ_{jt} as long as g_1 and δ_{jt} have different functional forms. I assume financial need is exogenous so formulaic financial aid is also exogenous of preference shocks. Then the correlation between residual enrollment and financial aid identifies α_{ik} .

6 Estimates

This section provides context for how parameter estimates shape outcomes across counterfactual financial aid schedules. Counterfactual financial aid can increase welfare through selection by improving student-university match. First, I describe the heterogeneous effects of flagship enrollment on enrollment. Graduation value-added correlates sensibly with observed characteristics. Conditional on characteristics, there is substantial variation in graduation chances driving selection into flagship enrollment and BOT take-up. Second, I describe heterogeneity in preferences for flagship universities, even among students with identical observed characteristics. Graduation value-added and preferences are correlated due to demographic variation. Third, model estimates imply heterogeneity in universities’ marginal costs across students. I interpret this as a relative valuation for middle-income students. In the status quo, graduation is relatively low because these students face large discounts despite their low value-added.

Parameter estimates in Table 6 show how ex-ante graduation correlates with student characteristics. Relative to regression estimates, these correlations correct for bias from asymmetric selection. Graduation at any university is greater for female, white, high family education, Advanced Placement, and Gifted program students. Average graduation rates by admission portfolio, high school, and college major are highly predictive of graduation. Notably, family income is not particularly predictive of graduation chances after controlling for other characteristics and unobserved selection. The one exception is that at non-flagship universities, students with the lowest family income, i.e. zero expected family contribution, have precisely greater graduation chances than students with the highest family income. In estimation, I make the relatively strong assumption that financial aid does not directly impact graduation. This

assumption may be reasonable because financial aid is a function of income, and generally, income is not precisely correlated with graduation. Estimates vary for flagship and non-flagship graduation chances so observed characteristics also drive variation in graduation value-added, the increase in graduation chances from enrolling at a flagship. These correlations may not be causal if characteristics correlate with the unobserved component of graduation.

The bottom rows show substantial variation in the unobserved component of graduation. For flagships, a one-standard-deviation increase is roughly equivalent to either attending a high school with 25 percentage points greater college graduation or the combined impact of being female, white, having two college-educated parents, and attending a high school gifted program. Second, residual graduation likelihoods are highly correlated between flagship and non-flagship universities: the confidence interval includes 1. Counterfactual financial aid that targets selection on ex-ante observed characteristics may be almost as efficient as a policy that targets selection on unobserved characteristics via graduation incentives.

Although family income is not precisely correlated with graduation, higher-income students place more weight on graduation when choosing a university. Several other demographics that correlate with graduation and value-added – female, white, family education – correspond to lower marginal utility from graduation. Advanced courses and expected college graduation by admission portfolio, high school, and major are positively correlated.

Consistent with prior research, net cost plays a large role in determining college enrollment. Net cost sensitivity is higher for students with more positive point estimates: lower income, white, lower family education, fewer AP courses, and no gifted program. Estimates contrast prior literature in suggesting that preferences for nearby universities are small and generally uncorrelated with student characteristics. The exception is that students with high family education are less deterred by distance.

The intercept term of university indirect utility absorbs remaining differences in the perceived quality of flagship universities relative to alternatives. Generally, demographics that correlate with high graduation correspond to lower mean quality. This pattern rationalizes relatively low enrollment at the high-graduation alternative. A few additional terms discipline the outside option. As expected, students with more non-flagship admissions have a greater outside option value. On the other hand, I do not find evidence of correlation with average net cost of students' non-flagship alternatives, which mitigates concern that students' preferences for non-flagship universities are incompletely modeled.

Figure 2 plots the distributions of students' valuations of university characteristics by normalizing

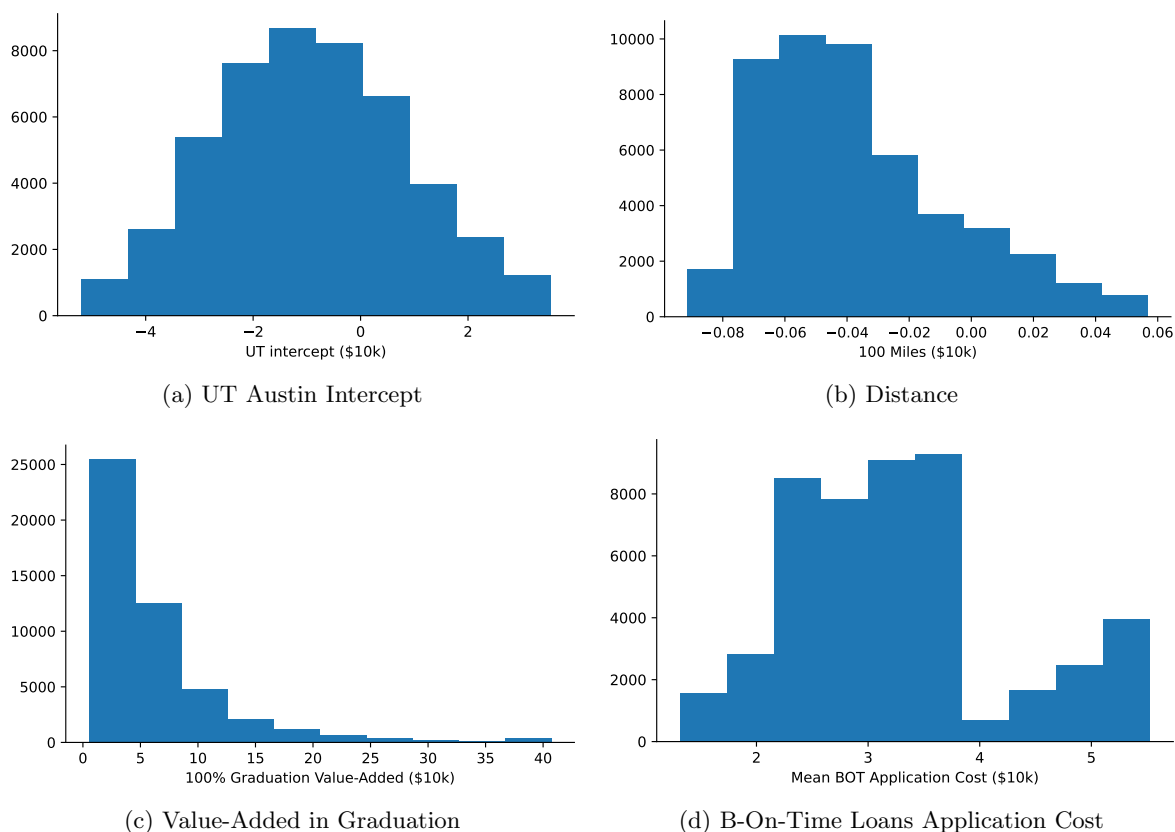
Table 6: Demand Model Estimates

	Intercept	Net Cost	Miles	Value-Added	BOT Cost	$P(Grad f = 0)$	$P(Grad f = 1)$
Intercept	9.388 (0.309)	1.499 (0.022)	-0.344 (0.042)	0.811 (0.131)	2.807 (1.048)	-7.677 (0.124)	-6.641 (0.092)
EFC Percentile	-3.552 (0.919)	-0.343 (0.092)	0.168 (0.237)	0.710 (0.028)	-0.144 (0.332)	-0.054 (0.202)	0.322 (0.172)
Female	-0.527 (0.107)	0.013 (0.010)	-0.010 (0.027)	-0.183 (0.013)	0.039 (0.041)	0.179 (0.025)	0.417 (0.022)
White	0.841 (0.129)	0.115 (0.012)	-0.039 (0.033)	-0.068 (0.013)	0.162 (0.063)	0.114 (0.027)	0.233 (0.025)
Family Education	-0.480 (0.078)	-0.053 (0.007)	0.070 (0.019)	-0.012 (0.007)	-0.001 (0.027)	0.059 (0.016)	0.085 (0.014)
Advanced Courses	-0.198 (0.018)	-0.007 (0.002)	0.005 (0.005)	0.123 (0.004)	-0.019 (0.009)	0.047 (0.003)	0.019 (0.003)
Gifted Program	-0.243 (0.109)	-0.016 (0.011)	0.033 (0.028)	-0.021 (0.011)	0.044 (0.041)	0.156 (0.024)	0.130 (0.021)
E[Graduation Admissions]				0.592 (0.047)	0.387 (0.179)	4.347 (0.115)	3.708 (0.097)
E[Graduation HS]				0.551 (0.078)	0.219 (0.225)	4.811 (0.125)	3.927 (0.102)
E[Graduation Major]				1.916 (0.104)	1.046 (0.175)	6.238 (0.137)	5.863 (0.114)
HS Residual Family Education				0.190 (0.063)	-0.095 (0.245)	0.150 (0.125)	0.679 (0.113)
EFC = 0	3.618 (0.291)	0.349 (0.030)	-0.024 (0.070)	-0.619 (0.039)	-1.023 (0.351)	0.150 (0.066)	0.064 (0.055)
EFC Percentile 1-20	4.887 (0.479)	0.328 (0.048)	-0.034 (0.114)	-0.769 (0.036)	-0.964 (0.357)	0.141 (0.101)	0.085 (0.085)
EFC Percentile 21-40	4.637 (0.678)	0.293 (0.065)	-0.103 (0.165)	-0.209 (0.019)	-0.357 (0.249)	0.135 (0.138)	0.083 (0.117)
EFC Percentile 41-60	4.253 (0.878)	0.291 (0.084)	-0.172 (0.219)	0.292 (0.018)	0.181 (0.289)	0.133 (0.175)	-0.043 (0.151)
EFC Percentile 61-80	2.885 (0.965)	0.189 (0.095)	-0.107 (0.238)		1.790 (0.717)	0.157 (0.196)	0.154 (0.167)
UT Austin	11.504 (0.319)						
Other 4-Year Net Cost	-0.008 (0.172)						
Other Admissions	-0.325 (0.093)						
Other Admissions ²	-0.049 (0.015)						
Standard Deviation					1.338 (0.438)	0.000 (0.580)	0.977 (0.027)
Correlation P(Grad f=1)						0.973 (0.004)	

Notes: This table shows model estimates for the Top 10 percent sample with asymptotic standard errors calculated using the approximate Hessian. The coefficient on Net Cost in indirect utility is the negative exponential of a linear combination of student characteristics. The coefficient on Value-Added in a positive exponential of a linear combination of student characteristics. $P(Grad | f = 0)$ and $P(Grad | f = 1)$ are logistic transformations of a linear combination of student characteristics. Standard Deviation refers to a normally distributed random term of each linear combination. The random components of $P(Grad | f = 0)$ and $P(Grad | f = 1)$ are correlated. An indicator for expected family contribution (“EFC”) above \$150k is omitted. Indicators for each financial aid year are not shown.

marginal utility for each characteristic by the marginal utility of net cost. Valuations for increased graduation are large and approximately log-normally distributed, consistent with the log-normal earnings potential of college graduates. By contrast, the distaste for distance explains a small fraction of the variation in flagship enrollment. As expected, most students prefer a nearby school, all else equal.

Figure 2: Distribution of Student Valuations of University Characteristics



Notes: This figure shows distributions across students in valuations of university characteristics, converted to units of \$10,000 by dividing marginal utility estimates by the marginal (dis)utility of net cost of attendance.

Large and dispersed application costs rationalize why so few students apply for BOT loans. These costs are perhaps too large to be plausible but might be rationalized by information frictions outside the model. Due to federal regulations, universities could not advertise BOT loans, so a large share of students may have been unaware of the program. Application costs might be much smaller among the set of students with knowledge of the program. Also, application costs rationalize some students' distaste for the coursework requirements of BOT loans. To graduate in four years, students could not take many courses outside their majors, and could not choose certain majors that are unlikely to be completed in four years. Students' idiosyncratic preferences for such courses and majors are outside of the model

except for unobserved variation in the BOT application cost.

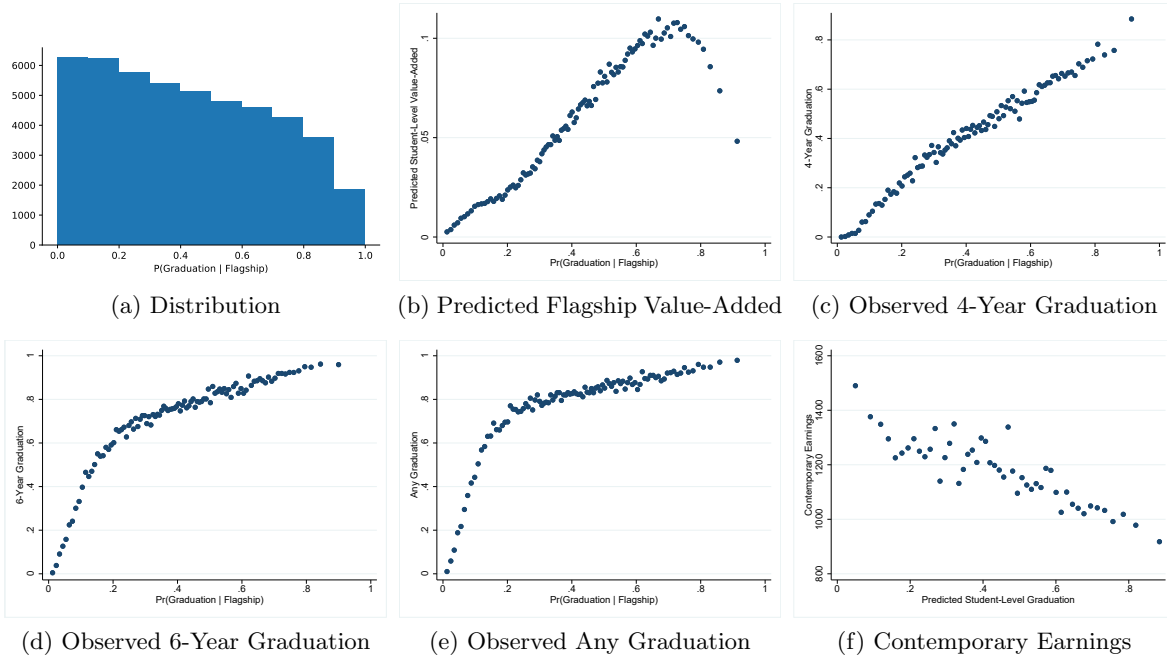
Large application costs for BOT loans imply that BOT had virtually no impact on ex-ante welfare. When choosing a university, the expected gross reduction in net cost of attendance from BOT is negligible relative to expected application costs. If application costs were much lower, some students might have chosen a different university with higher expected graduation and lower net cost via BOT. After enrolling, students realize their idiosyncratic application costs, which are sometimes low enough to apply. Some of these students graduate and enjoy the ex-post transfer of forgiven BOT loans. However, with a budget constraint, this small ex-post benefit comes at the relatively large cost of greater prices for most students who do not receive BOT loans or do not graduate on time. I rely on the timing assumption that students do not know their exact application cost until after enrolling at a university.

Dispersion in graduation value-added helps counterfactual financial aid to increase statewide graduation. Figure 3, Panel B shows that higher flagship graduation generally corresponds to higher graduation value-added. Equivalently, graduation increases in ability at all universities, but the slope is larger at flagships. This pattern is consequential for student equity and counterintuitive given prior work studying large compositional changes of selective universities (e.g., Black, Denning and Rothstein, 2020; Bleemer, 2021). Generally, in those studies, underrepresented students reap relatively large benefits from attending selective universities, while other students are similarly successful regardless of university. In this setting, targeting enrollment towards students with high graduation value-added (and improving statewide graduation rates) may also increase disparities across demographics and universities.

Panel A highlights that predicted ex-ante graduation likelihood at flagship universities is highly dispersed, most students have low graduation chances, and even the highest-ability students have meaningful uncertainty about graduation. Estimates fit the data well for both targeted and untargeted outcomes. Predicted 4-year flagship graduation is highly correlated with observed graduation over several time windows. In Panel C, the relationship is highly linear reflecting strong fit. The slope is less than 1 because idiosyncratic preferences drive some students to enroll at non-flagship universities and graduate at lower rates. The difference is largest for students with high flagship graduation chances. At 6 or more years after first enrollment, higher predicted graduation corresponds to higher realized graduation. The relationship is not linear because students with intermediate ex-ante chances graduate eventually but in more than 4 years. Finally, Panel F shows that students with high graduation chances exert more effort towards studying, proxied by less part-time work.

Estimates imply that flagship universities have strong non-pecuniary incentives to enroll middle-

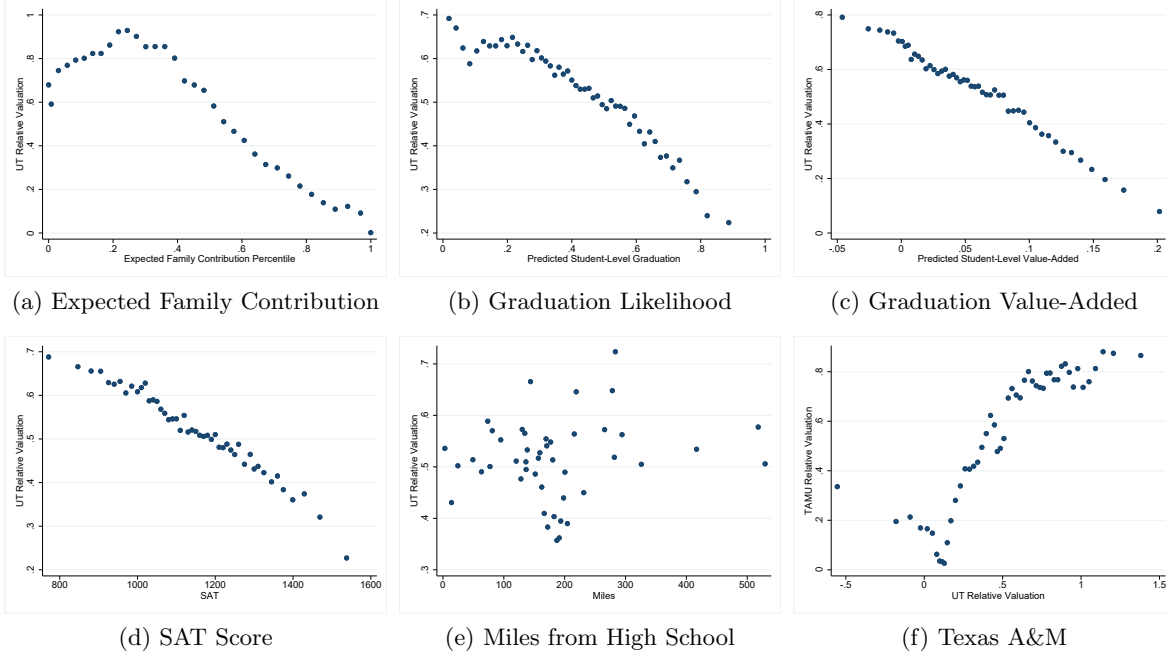
Figure 3: Predicted Ex-Ante Graduation Likelihood



Notes: This figure shows the distribution and correlates of students’ ex-ante likelihood of on-time graduation. “Predicted” indicates that a measure is calculated using estimates of the structural model. Figures (b)-(f) are scatter plots showing the means for each percentile of graduation likelihood.

income students over low-income students and low-income students over high-income students. I first calculate the effective marginal cost for each student and university by inverting the first-order condition for net cost, given a student’s net cost elasticity and federal financial aid. True marginal costs should be approximately equal across students, so I subtract the estimated marginal cost of the wealthiest students to calculate the universities’ relative valuation of each student. Figure 4 shows the UT Austin’s implied relative valuation of student expected family contribution (“EFC”). Relative valuations increase from minimum EFC to the 25th percentile and then decrease. The peak implies that enrolling a middle-income student is valued approximately \$9500 more than enrolling a high-income student. Since low-income students have lower expected graduation, value-added, and SAT scores, the first-order condition also implies that universities devalue these characteristics. Net cost schedules are a function of EFC, so these secondary correlations do not necessarily represent how universities would target composition with full information. For comparison, relative valuations are uncorrelated with distance and highly correlated between both flagship universities.

Figure 4: Flagship Universities' Relative Valuations of Student Characteristics



Notes: This figure shows the correlates of UT Austin's valuation of students (in units of \$10,000 per student) and their characteristics. Relative valuations are the difference in implied marginal cost between students with EFC above \$150,000 and a given student. Marginal costs are implied by student own-price elasticities and inverting the university's first-order condition. Figures (b)-(f) are scatter plots showing the means for each bin of the x-axis with an equal number of students. By construction, financial aid offers are purely functions of EFC, so any other correlation with relative valuation is driven by its correlation with EFC.

7 Counterfactual Financial Aid

Using the estimates from Section 6, I simulate students' choices under counterfactual financial aid schedules to illustrate the welfare effects of graduation incentives. First, I quantify the total cost of information frictions by deriving students' first-best financial aid offers. In the first best, the regulator has perfect information about all observed characteristics and unobserved graduation likelihood. Second, I decompose information frictions by setting financial aid without graduation chances, instead using the full set of ex-ante characteristics. Third, I add graduation incentives which largely close the gap between these two scenarios. Fourth, I show that graduation incentives alone can improve welfare, fixing the status quo average net cost. I conclude by assessing the equity implications of optimal financial aid schedules.

With perfect information about students' graduation chances and demand for flagship enrollment, the regulator could improve welfare by \$26,820 per admitted student. Fixing the net cost of the outside option, the regulator offers each student a net cost for flagship enrollment that conditions on all observed characteristics and the student's latent graduation chance. 82 percent of the welfare increase represents

Table 7: Average Counterfactual Outcomes Per Admitted Student

	Social Surplus	Consumer Surplus	Graduation (%)	Enrollment (%)	Revenue	Producer Surplus
(1) First-Best	2.682	2.222	0.550	-1.931	0.318	0.077
(2) Screening	2.145	2.150	-3.524	-0.021	0.090	0.011
(3) Screening + Incentive	2.463	2.101	-3.548	-0.072	0.325	0.013
(4) Status Quo + Incentive	0.679	0.243	1.568	2.522	0.017	0.007

Notes: This table shows key outcomes from counterfactual contract schedules. All outcomes are based on ex-ante expectations over students. Counterfactuals vary FFS the incremental net cost of flagship universities for each student, enforcing capacity and budget constraints. All outcomes are incremental, relative to the status quo. Unless otherwise indicated, outcomes are in units of \$10,000. Revenue is a conditional average among students who counterfactually enroll at a flagship university.

student surplus. The remainder largely represents the regulator’s valuation of graduation externalities. By varying the difference in net cost between flagship universities and their alternatives, the regulator increases graduation by 0.55 percentage points. Graduation increases at all universities – not just flagships, because non-flagship students have relatively low graduation effects of switching to a flagship. Despite greater efficiency, enrollment at flagships decreases by nearly 2 percentage points. Estimated welfare does not reflect that educational quality might improve with lower student-faculty ratios.

Even without perfect information, the regulator can achieve 80 percent of the first-best welfare increase by targeting financial aid based on students’ ex-ante observed characteristics (“Screening”). In the status quo, financial aid only conditions on expected family contribution, which reflects a relatively small share of the variation in student demand. Moreover, as discussed in Section 6, universities’ privately optimal financial aid schedules target students with relatively low graduation value-added. With improved screening, consumer surplus is nearly identical to first-best. Even though the regulator values graduation, the schedule lowers graduation by 3.5 percentage points relative to the status quo. Two patterns drive this large and counterintuitive decrease. First, there is substantial unobserved variation in graduation chances. Second, conditioning on observed characteristics, students with high graduation value-added have relatively low consumer surplus or low net cost elasticity.

Combining screening with graduation incentives can achieve 92 percent of the first-best welfare increase. With incentives, students receive an additional discount on net cost when graduating. For comparability to the status quo, I fix this discount at a size similar to status quo BOT loans. Efficiency may further improve if the incentive varied optimally with student observed characteristics. With any incentive versus screening alone, the net cost is higher for students who do not graduate so that the budget balances. Average outcomes are similar to screening alone because of the correlation between graduation value-added and net cost elasticity. Graduation incentives lower flagships’ expected net cost

and improve welfare for targeted students, but these discounts may not be cost-effective. Greater welfare among targeted students with high value-added may not outweigh lower welfare among non-targeted students who have a strong preference to attend a flagship despite low value-added.²⁰ These non-targeted students need to pay a higher net cost or attend a less preferred university.

Without screening, graduation incentives alone can achieve 25 percent of the first-best welfare increase. Relative to the status quo, prior counterfactuals sometimes lead to large changes in expected net cost among students who have the same EFC. Such changes may be infeasible under institutional constraints. For example, Texas has a legal requirement that tuition set-asides benefit students below a threshold EFC. Now, I fix average net cost by EFC at the status quo level while adding a graduation incentive. This approach prevents further redistribution across income but allows redistribution across match quality. Students' consumer surplus increase is an order of magnitude smaller than prior counterfactuals, but statewide graduation rates increase three times as much as in first best: 1.6 percentage points.²¹ With limited variation in unobserved value-added, screening on observed characteristics (third-degree price discrimination) and graduation incentives (second-degree price discrimination) can be thought of as substitutes. Welfare increases from only screening or only incentives are much larger than the combination of screening and incentives.

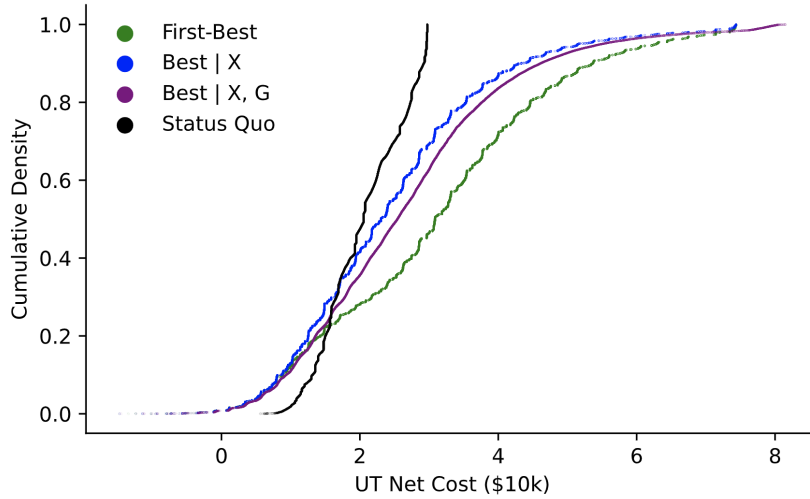
Figure 5 shows how counterfactual financial aid schedules largely improve welfare through cross-subsidization. Given the social objective, especially the budget and capacity constraints, students with relatively inelastic demand should face higher net costs. Greater revenue from these students permits much lower net costs for the relatively elastic students who most increase social welfare by enrolling at a flagship. Among counterfactual schedules, greater average welfare corresponds to a flatter cumulative distribution function and more redistribution. For example, the first-best schedule is the most progressive because perfect information enables more precise price discrimination.

Counterfactual financial aid schedules also lead to large compositional changes, which generally benefit students from historically underrepresented backgrounds. Students' incremental consumer surplus and graduation rates correlate with demographic characteristics, so counterfactual financial aid changes the composition of flagship enrollment. Table 8 compares the difference in average characteristics between flagship universities and alternatives. As before, these numbers rely on the sample of Top 10% students admitted to multiple universities, and this sample is not necessarily representative of Texas

²⁰These non-targeted students have a large surplus when the net cost does not include graduation incentives.

²¹Comparisons of this counterfactual to others should be interpreted with caution. Greater enrollment drives part of the increased graduation and welfare. After fixing expected net cost and satisfying the budget constraint, there are no degrees of freedom to bound enrollment below capacity.

Figure 5: Cumulative Distributions of Expected Net Cost Across Counterfactual Schedules



Notes: This figure shows the empirical cumulative distributions of net cost of attendance across students in the estimation sample for each counterfactual financial aid schedule. In First-Best, each student receives a personalized financial aid offer to each flagship university that maximizes the social objective given ex-ante student observed characteristics X and graduation likelihood g .

high school students or overall university enrollment. In the status quo, flagship students are relatively white, low income, female, and low family education, with fewer AP courses. Flagship students come from high schools with quite similar college graduation rates. Besides income, differences are relatively small. First best financial aid most increases flagship enrollment among students from historically underrepresented backgrounds: students of color, first-generation students, students with few AP courses, and students from low-graduation high schools. With cross-subsidization, the gap by income approaches zero. Composition is qualitatively similar across counterfactuals except for graduation incentives alone, which especially incentivize flagship enrollment among relatively low-income and male students.

Table 8: Counterfactual Demographic Composition: Flagship - Alternatives

	(SQ)	(1)	(2)	(3)	(4)
White	0.004	-0.041	-0.030	-0.025	-0.044
Wealth (EFC Percentile)	-0.056	0.003	-0.004	0.002	-0.087
Female	0.007	0.010	0.022	0.029	-0.031
College-Educated Parents	-0.007	-0.083	-0.073	-0.055	-0.108
Advanced Courses	-0.406	-1.207	-1.349	-1.281	-0.686
College Graduation by HS	-0.001	-0.051	-0.046	-0.041	-0.027

Notes: Each cell is a difference in demographic composition between flagship universities and other 4-year public universities. For example, 0.4pp more students are white at flagship universities in the status quo. Column numbers index counterfactuals: (1) First-Best, (2) Screening, (3) Screening + Incentive, and (4) Status Quo + Incentive, relative to (SQ) Status Quo. College graduation by HS is the average 4-year graduation (at public 4-year universities) among students that graduated from the same high school.

8 Conclusion

This paper demonstrates how revenue-neutral financial aid design can improve both efficiency and equity by allocating limited spots at selective universities to the students who most benefit. Graduation incentives are expedient at a moment when federal and state policies constrain universities from increasing diversity through admissions. First, graduation incentives exploit students' private information about match quality to lower the price for relatively well-matched students and increase their enrollment. Second, although graduation incentives are race-neutral (and neutral with respect to all demographics), they happen to also increase enrollment at selective universities among historically underrepresented students. These features may be increasingly useful if future policy further restricts universities from considering student race or its proxies, e.g., via essays or scholarships.

The benefits of financial aid design require changing students' choice of university. Combining evidence from a discontinuity in financial need and a structural model of college choice, I show how historical graduation incentives failed to improve welfare. B-On-Time loans could have improved match quality by changing the relative price of selective universities. Instead, it introduced uniform incentives that aimed to increase effort. In this setting, the marginal student was too constrained to change behavior.

B-On-Time's limitations are instructive for contemporary national and state policies. To target selection and improve match quality, programs must change relative incentives to attend low versus high-quality universities. For example, as of 2024, the College Cost Reduction Act aims to reduce higher education spending by expanding income-based repayment plans, limiting per-person federal loans, and reducing dispersion in list prices. Around the same time, several university systems announced zero-

tuition policies for students below an income threshold. In such cases, unintended consequences for match quality might diminish the benefits of lower average costs. Tighter bounds on net cost can limit opportunities for screening on price and for redistribution.

References

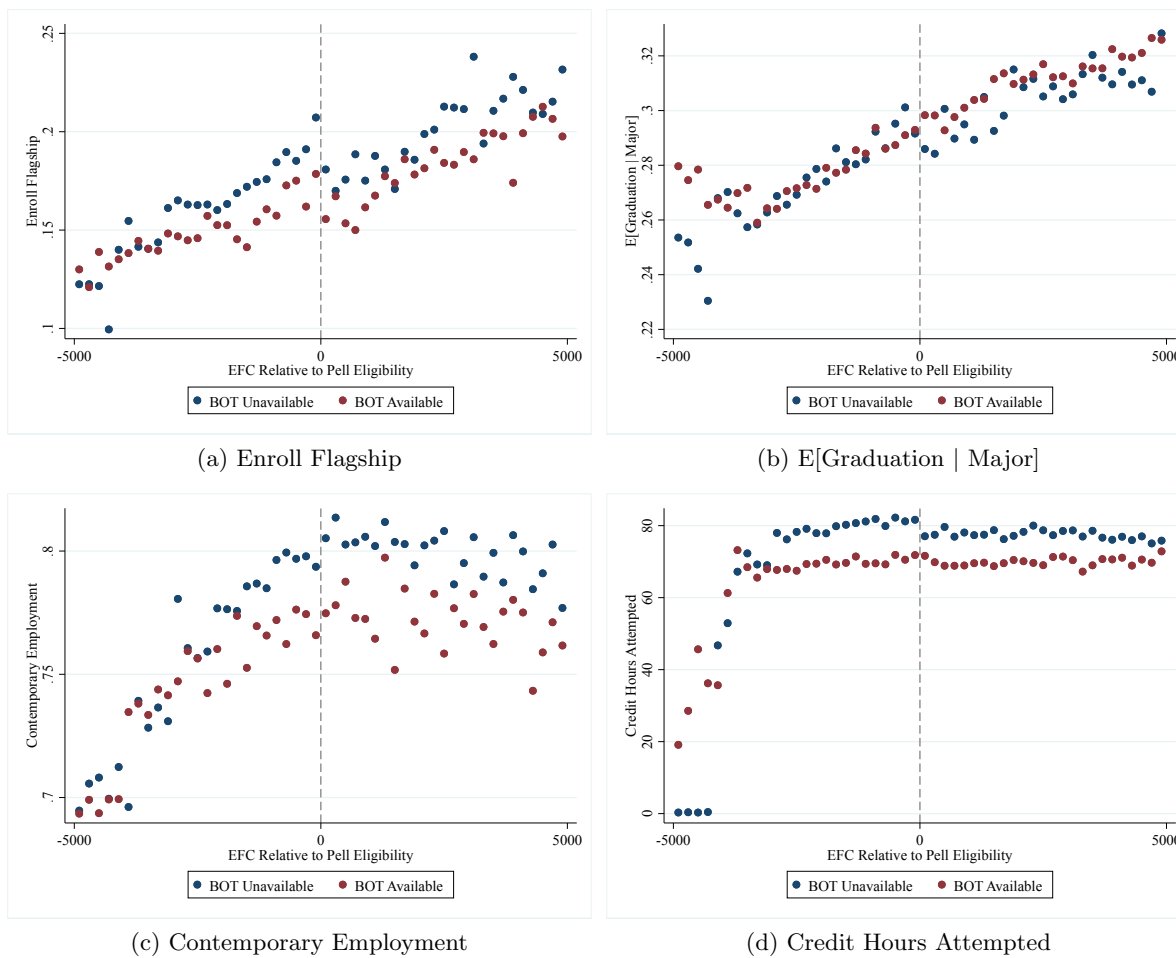
- Andrews, R. J., J. Li, and M. F. Lovenheim.** 2016. “Quantile Treatment Effects of College Quality on Earnings.” *Journal of Human Resources*, 51: 200–238.
- Black, Sandra E., Jeffrey T. Denning, and Jesse Rothstein.** 2020. “Winners and Losers? The Effect of Gaining and Losing Access to Selective Colleges on Education and Labor Market Outcomes.”
- Bleemer, Zachary.** 2021. “Affirmative Action, Mismatch, and Economic Mobility after California’s Proposition 209.” *The Quarterly Journal of Economics*, 137.
- Board, College.** 2019. “Trends in Student Aid Highlights.”
- Chetty, Raj, Matthew O. Jackson, Theresa Kuchler, Johannes Stroebel, Nathaniel Hendren, Robert B. Fluegge, Sara Gong, Federico Gonzalez, Armelle Grondin, Matthew Jacob, Drew Johnston, Martin Koenen, Eduardo Laguna-Muggenburg, Florian Mudekereza, Tom Rutter, Nicolaj Thor, Wilbur Townsend, Ruby Zhang, Mike Bailey, Pablo Barberá, Monica Bhole, and Nils Wernerfelt.** 2022. “Social capital II: determinants of economic connectedness.” *Nature*, 1–13.
- Colas, Mark, Sebastian Findeisen, and Dominik Sachs.** 2021. “Optimal Need-Based Financial Aid.” *Journal of Political Economy*, 129: 492–533.
- Cunha, Jesse M., Trey Miller, and Emily Weisburst.** 2017. “Information and College Decisions: Evidence From the Texas GO Center Project.” *Educational Evaluation and Policy Analysis*, 40: 151–170.
- 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.
- Dale, Stacy B., and Alan B. Krueger.** 2014. “Estimating the Effects of College Characteristics over the Career Using Administrative Earnings Data.” *Journal of Human Resources*, 49: 323–358.
- Denning, Jeffrey T., Benjamin M. Marx, and Lesley J. Turner.** 2019. “ProPelled: The Effects of Grants on Graduation, Earnings, and Welfare.” *American Economic Journal: Applied Economics*, 11: 193–224.
- Dobbin, Cauê, Nano Barahona, and Sebastián Otero.** 2022. “The Equilibrium Effects of Subsidized Student Loans.”
- Epplé, Dennis, Francisco Martínez-Mora, and Richard Romano.** 2021. “Charter School Practices and Student Selection: An Equilibrium Analysis.”
- Epplé, Dennis, Richard Romano, Sinan Sarpaçad, and Holger Sieg.** 2017. “A general equilibrium analysis of state and private colleges and access to higher education in the U.S.” *Journal of Public Economics*, 155: 164–178.
- Epplé, Dennis, Richard Romano, Sinan Sarpaça, Holger Sieg, and Melanie Zaber.** 2019. “Market power and price discrimination in the US market for higher education.” *RAND Journal of Economics (Wiley-Blackwell)*, 50: 201–225.
- Evans, Brent J., and Tuan D. Nguyen.** 2019. “Monetary substitution of loans, earnings, and need-based aid in postsecondary education: The impact of Pell Grant eligibility.” *Economics of Education Review*, 70: 1–19.
- Fillmore, Ian.** 2022. “Price Discrimination and Public Policy in the US College Market.” *The Review of Economic Studies*.
- Goldrick-Rab, Sara, Robert Kelchen, Douglas N. Harris, and James Benson.** 2016. “Reducing Income Inequality in Educational Attainment: Experimental Evidence on the Impact of Financial Aid on College Completion.” *American Journal of Sociology*, 121: 1762–1817.
- Gunnes, Trude, Lars J. Kirkebøen, and Marte Rønning.** 2013. “Financial incentives and study duration in higher education.” *Labour Economics*, 25: 1–11.
- Hämäläinen, Ulla, Kristian Koerselman, and Roope Uusitalo.** 2016. “Financial incentives for timely graduation.”
- Lans Bovenberg, A., and Bas Jacobs.** 2005. “Redistribution and education subsidies are Siamese twins.” *Journal of Public Economics*, 89: 2005–2035.
- Lawson, Nicholas.** 2017. “Liquidity Constraints, Fiscal Externalities and Optimal Tuition Subsidies.” *American Economic Journal: Economic Policy*, 9: 313–343.

- Leuven, Edwin, Hessel Oosterbeek, and Bas van der Klaauw.** 2010. “THE EFFECT OF FINANCIAL REWARDS ON STUDENTS’ ACHIEVEMENT: EVIDENCE FROM A RANDOMIZED EXPERIMENT.” *Journal of the European Economic Association*, 8: 1243–1265.
- Looney, Adam, and Constantine Yannelis.** 2024. “What Went Wrong with Federal Student Loans?” *The Journal of Economic Perspectives*, 38: 209–236.
- Lovenheim, Michael, and Jonathan Smith.** 2023. “Returns to different postsecondary investments: Institution type, academic programs, and credentials.” *Handbook of the economics of education*, 6: 187–318.
- Malacrino, Davide, Samuel Nocito, and Raffaele Saggio.** 2024. “Do Reforms Aimed at Reducing Time to Graduation Work? Evidence from the Italian Higher Education System.”
- Mathias, Kifmann, Heineck Martin, and Lorenz Normann.** 2006. “A Duration Analysis of the Effects of Tuition Fees for Long-Term Students in Germany.” *Journal of Economics and Statistics (Jahrbuecher fuer Nationaloekonomie und Statistik)*, 226: 82–109.
- Mountjoy, Jack, and Brent R. Hickman.** 2021. “The Returns to College(s): Relative Value-Added and Match Effects in Higher Education.”
- Murphy, Richard, and Gill Wyness.** 2022. “Testing Means-Tested Aid.” *Journal of Labor Economics*.
- Nguyen, Tuan D., Jenna W. Kramer, and Brent J. Evans.** 2019. “The Effects of Grant Aid on Student Persistence and Degree Attainment: A Systematic Review and Meta-Analysis of the Causal Evidence.” *Review of Educational Research*, 89: 831–874.
- Rattini, Veronica.** 2023. “The effects of financial aid on graduation and labor market outcomes: New evidence from matched education-labor data.” *Economics of Education Review*, 96: 102444.
- Rubin, Rachel B.** 2011. “The Pell and the Poor: A Regression-Discontinuity Analysis of On-Time College Enrollment.” *Research in Higher Education*, 52: 675–692.
- van der Klaauw, Wilbert.** 2002. “Estimating the Effect of Financial Aid Offers on College Enrollment: A Regression-Discontinuity Approach*.” *International Economic Review*, 43: 1249–1287.
- Waldfoegel, Joel.** 2015. “First Degree Price Discrimination Goes to School.” *The Journal of Industrial Economics*, 63: 569–597.

A Additional Analysis

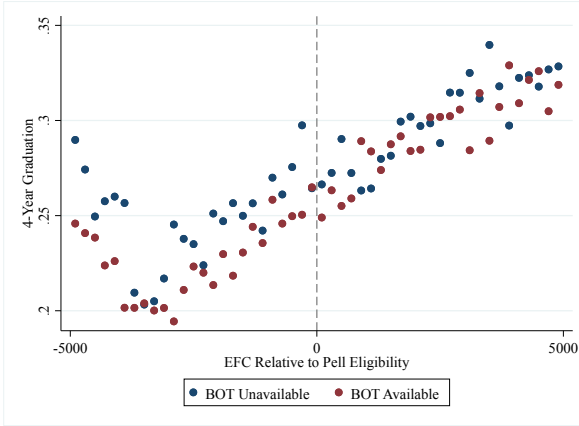
A.1 Additional Tables and Figures

Figure A.1: Pell Eligibility, Program Choice, and Effort

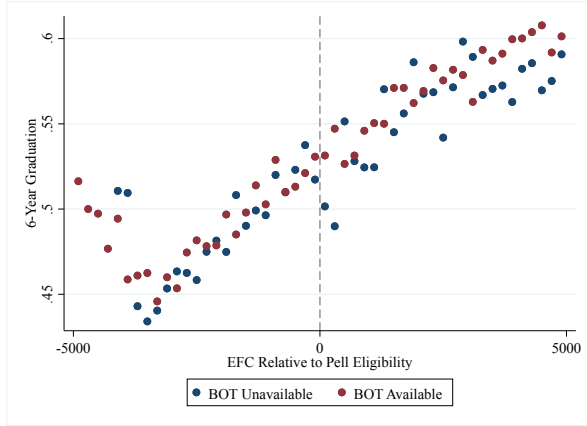


Notes: The x-axis shows expected family contribution (EFC) minus the maximum EFC for Pell Grant eligibility, which varies by year. Points represent averages of students within \$500 bins of EFC. BOT Available includes students with initial financial aid offers in years when BOT loans were widely available.

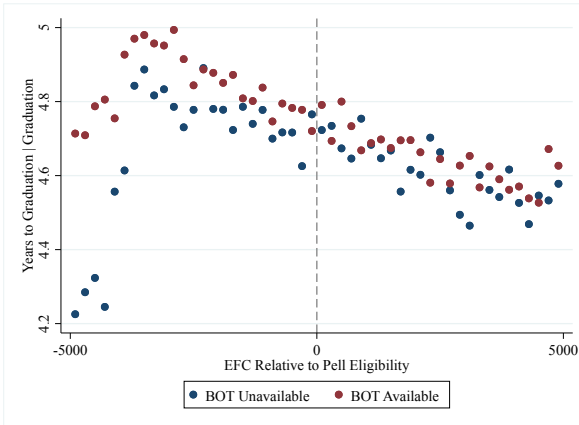
Figure A.2: Pell Eligibility and Outcomes




(a) Graduation Within 4 Years



(b) Graduation Within 6 Years



(c) Years To Graduation | Any Graduation



(d) Annual Earnings 8-10 Years After High School

Notes: The x-axis shows expected family contribution (EFC) minus the maximum EFC for Pell Grant eligibility, which varies by year. Points represent averages of students within \$500 bins of EFC. BOT Available includes students with initial financial aid offers in years when BOT loans were widely available.

Table A.1: Difference-in-Discontinuity: Pell Ineligible

	Estimate	Std. Err.	P-Value	Outcome Mean	R ²	Obs.
Financial Aid						
Federal Aid	-0.334	(0.015)	[<0.001]	1.661	0.837	256,984
Any B-On-Time Loans	-0.002	(0.000)	[<0.001]	0.036	0.102	256,984
Traditional Loans	0.614	(0.107)	[<0.001]	3.987	0.231	256,984
Net Cost of Attendance	1.284	(0.149)	[<0.001]	12.115	0.609	256,984
Program Choice						
Enrolled at a Flagship	0.001	(0.002)	[0.560]	0.169	0.933	256,984
E[Graduation Major]	-0.008	(0.004)	[0.054]	0.288	0.104	256,984
Effect of Major on Earnings	-0.196	(0.110)	[0.076]	-0.248	0.043	256,981
Effort Proxies						
Credit Hours Attempted	-1.105	(0.800)	[0.168]	68.077	0.587	256,984
Contemporary Employment	-0.005	(0.008)	[0.548]	0.759	0.039	256,984
Max Class Rank	-0.053	(0.020)	[0.009]	2.301	0.671	256,796
Outcomes						
Graduation Within 4 Years	-0.009	(0.009)	[0.334]	0.243	0.089	256,979
Graduation Within 6 Years	-0.033	(0.012)	[0.008]	0.513	0.119	227,552
Earnings After 8-10 Years	-1.224	(0.719)	[0.090]	73.061	0.365	256,984
Balance						
Top 10% Within HS	0.001	(0.007)	[0.886]	0.213	0.210	269,037
Top 25% Within HS	0.014	(0.008)	[0.105]	0.402	0.140	269,037
Advanced HS Courses	0.049	(0.056)	[0.382]	3.906	0.252	284,046
White	-0.001	(0.008)	[0.908]	0.421	0.221	284,046
Free-Lunch	0.001	(0.007)	[0.867]	0.244	0.127	284,046
Discipline Days	0.031	(0.106)	[0.768]	1.834	0.027	284,046

Notes: This table shows estimates of the uninteracted discontinuity at the Pell eligibility threshold (from years without B-On-Time loans), i.e., β_1 from Equation (2). The specification corresponds to Column (4) in Table 1: a local quadratic difference-in-discontinuity with controls for students' initial covariates, final net cost of attendance, and fixed effects for the financial aid year, among students with an EFC within \$5000 of the Pell eligibility threshold. All monetary measures are scaled down by 1000.