An Economic Approach to Generalizing Findings from Regression-Discontinuity Designs

by

Nirav Mehta

Working Paper # 2015-6 **November 2015**

Updated May 21, 2016; March 29, 2017

CIBC Working Paper Series

Department of Economics
Social Science Centre
Western University
London, Ontario, N6A 5C2
Canada

An Economic Approach to Generalizing Findings from Regression-Discontinuity Designs

Nirav Mehta*

March 29, 2017

Abstract

Regression-discontinuity (RD) designs estimate treatment effects at a cutoff. This paper shows what can be learned about average treatment effects for the treated (ATT), untreated (ATUT), and population (ATE) if the cutoff was chosen to maximize the net gain from treatment. The ATT must be positive. Without capacity constraints, the RD estimate bounds the ATT from below and the ATUT from above, implying bounds for the ATE, and optimality of the cutoff rules out constant treatment effects. Testable implications of cutoff optimality are derived. Bounds are looser if the capacity constraint binds. The results are applied to existing RD studies.

*Department of Economics, University of Western Ontario. Email: nirav.mehta@uwo.ca.

1 Introduction

Regression-discontinuity (RD) designs are a popular tool for program evaluation due to the ubiquity of cutoff-based treatment assignment and agnosticism they afford researchers. However, it is not always clear how best to use such estimates to inform policy. Suppose an impact evaluation of a program using an RD design found the treatment effect at the cutoff to be positive, but small. Is this evidence the program should be terminated? Because treatment effects likely vary, it is useful to extend findings from RD designs to people away from the cutoff.

The goal of this paper is to demonstrate how combining an RD estimate with a simple economic model can deliver useful information about treatment effects in certain contexts. Researchers using RD designs typically focus on treatment effects at the cutoff, an approach that has the ostensible benefit of imposing minimal structure. Though such estimates can help decide whether to extend the treatment at the margin, this approach implicitly assumes that the treatment effect at the cutoff is completely uninformative about treatment effects elsewhere. This assumption may not always be appropriate. This paper considers a program administrator interested in maximizing the gain from treatment, net of treatment costs, but who, as is often the case in real-world applications, has been constrained to assign treatment using a cutoff rule. By imposing structure not on treatment effects, but on the economic environment, I show that we can learn about important features of treatment effects in cases where there is reason to believe the administrator has information about the costs and benefits of treatment. That is, the choice of cutoff may reveal information key to understanding the overall costs and benefits of a program.[1]

Combining an estimate of the treatment effect at the cutoff with a simple model of cutoff choice yields many insights. The most basic inference we can draw about treatment effects is that the average effect of treatment on the treated (ATT) must be positive if the marginal cost of treatment is positive. There is also a basic testable implication of the model: if the treatment effect at the cutoff is negative, then we can reject cutoff optimality. If the marginal cost is nonincreasing and the cutoff does not reflect a binding capacity constraint, optimality implies that the treatment effect is increasing at the cutoff and several additional results relate the treatment effect at the cutoff to treatment effects elsewhere. First, the RD estimate at the treatment cutoff provides a lower bound for the ATT; if this were not true the administrator could have obtained higher utility by moving the cutoff. Intuitively, the administrator will not place the cutoff where the gain from treatment is very large if, as is

[1] See Heckman et al. (1997) for a discussion of heterogeneous treatment effects in the context of an experimental setting. See Heckman and Smith (1998) for a discussion of how to link information about program benefits with conventional cost-benefit analysis and welfare calculations.

commonly assumed by practitioners, treatment effects are smooth. Second, the fact that the administrator chose not to extend treatment to certain students provides an upper bound for the average effect of treatment on the untreated (ATUT). These bounds only require the RD estimate and qualitative information about treatment costs, not actual estimates of the cost of treatment. Additionally, we can rule out a constant treatment effect, a finding that relates to the literature comparing findings from RD studies with those from experiments (e.g., Black et al. (2007) and Buddelmeyer and Skoufias (2004)). Finally, the bounds on the ATT and ATUT provide informative bounds on the average treatment effect (ATE). Bounds are looser if the marginal cost of treatment is strictly increasing or the chosen cutoff reflects a binding capacity constraint. However, a new testable implication also emerges in the latter case: If the program is subsequently expanded until the constraint no longer binds, the RD estimate will be lower than when it had been when the constraint was binding.

These results have implications for the use of RD estimates in policymaking. Perhaps the most striking result is that, because an unconstrained administrator is unlikely to choose a cutoff where the gain is quite large, one may incorrectly surmise from RD estimates that certain programs are ineffective and eliminate them, even though in reality they are quite effective for the treated population. In fact, such a mistake would be more likely for a program with a very low marginal cost, holding constant the ATT, because an unconstrained optimizing administrator would extend treatment to units until the gain, i.e., marginal benefit, equaled this low marginal cost. If the cutoff reflects a binding capacity constraint, then the RD estimate will exceed the marginal cost of treatment, which may help explain why it is sometimes difficult to "scale up" successful interventions to larger populations (see, e.g., Elmore (1996) and Sternberg et al. (2006)).

To make these findings more concrete, the results are illustrated using two recent empirical applications. The studies used in the applications, Hoekstra (2009) and Lindo et al. (2010), exploit discontinuities in treatment assignment rules to study questions in the economics of education, and cover cases where the capacity constraint likely does and does not bind. I first show that bounds obtained for the sharp RD design can be extended to the "fuzzy" design used in one of the applications. I also formally test the necessary conditions of optimality for both applications and find that I cannot reject the assumption that cutoffs were chosen optimally by informed program administrators.[2]

There is a long tradition in economics, starting with Roy (1951), of using revealed preferences to inform empirical work about information unobservable to the econometrician. This

[2]Applications need not be restricted to education. For example, the findings from this paper might apply to a job training program in which the program officer receives a bonus based on the increase in wages. I reiterate that one could test whether the environment studied in this paper was applicable for this, or any other context by checking that the treatment effect is nonnegative and nondecreasing at the cutoff.

paper simply makes clear what we could learn by embedding the choice of treatment cutoff within a larger decision problem. To most clearly demonstrate what can be learned by taking into account the administrator's context, I assume she knows both treatment effects—which may be heterogeneous—and her cost function. Though for a different context, the insights of this paper are similar to those in Heckman and Vytlacil (2007), who build on the work of Björklund and Moffitt (1987) by using the optimality of individual decision-making in the context of a Roy model and the marginal treatment effect (MTE) to bound treatment effects for non-marginal units.

By embedding an RD design within a simple economic model, this paper contributes to several literatures. First, it adds to the literature examining technical features of RD designs (Hahn et al. (2001), Van der Klaauw (2008)) by demonstrating how inferences from RD designs can be generalized by using a simple theoretical framework. This paper also relates to the debate about the usefulness of discontinuity and other estimators of treatment effects (Heckman et al. (1999), Heckman and Urzua (2010), Imbens (2010)). In adopting a bounding approach, this paper has a similarity to Manski and Pepper (2000), which uses monotone instrumental variables to bound treatment effects. This paper takes a different approach by assuming optimality of assignment to treatment status, while making minimal assumptions about the responses of agents to the treatment.

This paper also contributes to a literature seeking to extend results from RD designs. Angrist and Rokkanen (2015) share a similar motivation and goal to this paper, invoking a conditional independence assumption to generalize findings from RD studies. Specifically, their approach exploits additional covariates which, when conditioned upon, eliminate the relationship between the running variable and outcome. This is testable for units near the cutoff, suggesting a way to confirm that extrapolation away from the cutoff would be reasonable. Due to the different type of assumption made (i.e., statistical versus economic), this paper complements their work. Dong and Lewbel (2015) show that the differentiability assumptions typically invoked to estimate RD models can be exploited to estimate the derivative of the treatment effect. In a similar vein, DiNardo and Lee (2011) show how a Taylor expansion around the cutoff can be used to estimate the ATT. There is also an extensive literature studying the validity of RD designs.[3] This paper treats the RD design as valid and instead examines how findings from such studies can be generalized to other parts of the population.

Section 2 lays out the model of the administrator's problem, which is used to obtain theoretical results in Section 3. Section 4 illustrates the results using empirical applications.

[3]This literature is reviewed in Imbens and Lemieux (2008). See Lee and Card (2008) and McCrary (2008) for examples.

Section 5 discusses policy implications as well as variations on the informational assumptions made in this paper.

2 Model

Consider a program administrator who can assign students to a training program. The administrator knows how effective the program would be for any given student and also knows the cost of enrolling students in the program. Due to institutional reasons, she is constrained to choose a cutoff rule for assigning the treatment, above which students are enrolled.[4] The choice of cutoff-based treatment assignment captures the fact that many real-world policies are discrete in nature (Ferrall and Shearer (1999)).

There is a measure one of students, indexed by $x \in [0, 1]$; this index doubles as the running variable in the discontinuity design. For example, the running variable could be student SAT scores. For simplicity, assume students are uniformly distributed over $[0, 1]$. Let $\tau(x) = 1$ if x is given the treatment and 0 otherwise. The administrator is constrained to choose a cutoff rule where $\tau(x) = 1$ if and only if $x \geq \kappa$ for some $\kappa \in (0, 1)$.[5] To simplify exposition, I assume a "sharp" RD design and perfect compliance, which means that students with indices of κ or greater receive the treatment (i.e., participate in the program) and students with indices less than κ don't receive the treatment (i.e., don't participate in the program).[6] Let κ^* denote the treatment cutoff chosen by the administrator. The measure of students receiving treatment is $\mu = \int_0^1 \tau(x)dx = 1 - \kappa^*$. The administrator faces a cost of treating μ students, $c(\mu)$, which may capture an implicit budget constraint. Note that $c(\mu)$ could capture either primitive nonlinearity in the cost function or even some forms of cost heterogeneity with respect to x (see Appendix D). If $c(\mu)$ cannot capture such heterogeneity then treatment effect, defined shortly, could be interpreted as being net of such heterogeneous costs. Results from the administrator's unconstrained problem, the model developed in this section, are presented in Section 3.1. Section 3.2 introduces capacity constraints and then analyzes that

[4]Allowing the administrator to choose which side of the cutoff to treat does not affect most results. In particular, this would not change bounds on the mean effect of treatment on the treated, untreated, or population. In the following, I indicate where this assumption would affect a result.

[5] Note the implicit assumption that the administrator chooses whether to treat students above (or below) the cutoff. As will be clear after the next section, if the administrator could choose precisely which $x \in [0, 1]$ to treat we can make the inference that the gain from serving those students was at least as large as the marginal cost of serving them, point-wise. This assumption would imply that gains were positive for *all* treated students, as opposed to positive *on average*. Therefore, we could also bound from below the share of students who would gain from treatment: $\int_0^1 1\{\Delta(x) > 0\}dx \geq \mu$. This would be relevant if, say, the population voted on whether to implement the treatment.

[6]Section 4 shows that the theoretical results are identical under a "fuzzy" design where, instead of perfect compliance, the probability of participation discontinuously changes at the treatment cutoff.

problem.

Let $Y_\tau(x)$ denote student x's outcome under treatment group τ. For example, this may be the wage earned as a function of being enrolled in a training program. The treatment effect is $\Delta(x) \equiv Y_1(x) - Y_0(x)$, and, as will be made explicit by Assumption 1, is known by the administrator. As is common in studies employing discontinuity designs, the stable-unit-treatment-value-assumption (SUTVA) is maintained here (Rubin (1980)), ruling out general equilibrium effects and other interactions between other units' treatment status and one's own treatment effect, such as endogenous social interactions.

The fundamental problem of causal inference is that we only observe each student in one treatment condition, making it difficult to recover the entire function $\Delta(\cdot)$. What can we say about $\Delta(\cdot)$ knowing 1) that κ^*, the treatment cutoff, was chosen by the administrator and 2) the value of $\Delta(\kappa^*)$, from a RD design?[7] Though I find that we can not say much about $\Delta(\cdot)$ for particular x who are not at the cutoff, we will be able to bound averages of $\Delta(\cdot)$ over different intervals, listed in Definition 1.

Definition 1 (Treatment effects of interest)**.** *I focus on the:*

- *Average Effect of Treatment on the Treated (ATT):* $\int_{\kappa^*}^{1} \frac{\Delta(x)}{1-\kappa^*} dx$*,*
- *Average Effect of Treatment on the Untreated (ATUT):* $\int_{0}^{\kappa^*} \frac{\Delta(x)}{\kappa^*} dx$*, and*
- *Average Treatment Effect among all units (ATE):* $\int_0^1 \Delta(x) dx$*.*

Note that the "local average treatment effect" (LATE) at the treatment cutoff is simply $\Delta(\kappa^*)$.

The administrator's problem is to choose a cutoff to maximize the total treatment effect, net treatment cost:[8]

$$\max_{\widetilde{\kappa}} \ \beta \left(\int_{\widetilde{\kappa}}^{1} \Delta(x) dx \right) - c\left(1 - \widetilde{\kappa}\right), \tag{1}$$

where β measures how much the administrator values the effect of the program in terms of the cost of treatment. Though in principle identified when the cost function is known to the researcher, β is normalized to one to simplify exposition. The administrator has an outside option of zero. This objective function is similar to those studied in Manski (2003, 2004, 2011), where a utilitarian social planner takes an action to maximize expected welfare (i.e., the gain net the cost of treatment), as well as those in studies of statistical discrimination

[7]Note that $\kappa^* \in (0, 1)$; otherwise all units have the same treatment status and an RD design cannot be implemented.

[8]For an example of a slightly different objective function, see Heinrich et al. (2002), who study treatment decisions when administrators face performance standards.

such as Knowles et al. (2001), Anwar and Fang (2006), and Brock et al. (2011), where police officers face a cost of pulling over motorists to maximize expected hit rates. The economic rationale for studying a utilitarian social planner is that a system of lump-sum transfers could then be designed to redistribute total output in such a manner as the social planner saw fit; that is, the utilitarian objective corresponds to the efficient allocation.[9]

Assumption 1. *The following assumptions about costs and benefits of treatment are maintained throughout this section:*

(i) The cost of treatment is known by the administrator, and is strictly increasing and linear in the number of units treated, i.e., $c(\mu) = \chi\mu$*, where* $\chi = c'(\cdot) > 0$ *denotes the constant marginal cost of treatment.*

(ii) Treatment effects $\Delta(\cdot)$ *are differentiable in* x *and known by the administrator.*

(iii) There exist finite lower and upper bounds of $\Delta(\cdot)$*. Denote these by* $\underline{\Delta} \in \mathcal{R}$ *and* $\overline{\Delta} \in \mathcal{R}$*, respectively.*

Assumption 1(i) implies that the marginal cost of providing treatment is known and strictly positive; the assumption of a linear cost function is made to simplify exposition. The assumption that the cost function is linear is also made in Manski (2011)'s analysis of optimal treatment choices, who assumes costs are separable across treated units. All the following results would obtain in the more general case where the marginal cost of treatment was nonincreasing in μ, i.e., where the cost of treatment is weakly concave in μ.[10] The first part of Assumption 1(ii), i.e., differentiability of $\Delta(\cdot)$, is typically invoked in applications of RD designs, which control for a smooth (typically polynomial or smoothed non-parametric) function of the running variable.[11] The second part of Assumption 1(ii), that $\Delta(\cdot)$ is known by the administrator, produces a testable implication (as is shown in the next section). The administrator need not be perfectly informed about students' potential outcomes; so long as the administrator has an unbiased signal of $\Delta(x)$, uncertainty about treatment effects does not affect the analysis, as the administrator's objective is linear.[12] In general, there may be multiple students with the index x, and heterogeneous treatment

[9]In the baseline case presented here, the administrator is a utilitarian who weighs gains for all students equally. See Appendix C for a case where gains are not weighed equally; the results derived in Section 3.1 are also obtained there.

[10]Note that $\frac{\partial c}{\partial \mu} > 0 \Rightarrow \frac{\partial c}{\partial \kappa} < 0$, as $\frac{\partial \mu}{\partial \kappa} < 0$. If one thought marginal costs were increasing in a context of interest, the bounds would have to be adjusted accordingly (see Appendix D for details).

[11]Note that $\Delta(\cdot)$ only needs to be smooth local to the chosen cutoff. The assumption that it is globally smooth is only made to simplify exposition.

[12] This is shown in Appendix B, which also examines the case of biased beliefs about $\Delta(\cdot)$.

effects among these students at x. In this case, $\Delta(x)$ would represent the expected gain from treating students at x, i.e., $\Delta(x) = \int \Delta(z) f(z|x) dz$, where $f(\cdot)$ characterizes treatment effect heterogeneity among students with index x. Assumption 1(iii) means that the set of outcomes $Y_\tau(\cdot)$ has finite support, which makes sense for outcomes such as wages, test scores, or probabilities.[13] In what follows, I refer to a bound as "uninformative" when the bound on the object of interest cannot be tightened relative to the bound specified in Assumption 1(iii). An informative lower bound for $\Delta(x)$ is a lower bound above $\underline{\Delta}$, an informative upper bound is an upper bound below $\overline{\Delta}$, and the uninformative bound for $\Delta(x)$ is $[\underline{\Delta}, \overline{\Delta}]$, $\forall x \in [0, 1]$, which has a width of $\overline{\Delta} - \underline{\Delta}$.[14]

Unless superseded by another assumption, assumptions are maintained after introduced. For example, Assumption 1 is maintained until Assumption 1(i) is superseded by Assumption 1′(i) in Appendix D, which studies variable marginal costs of treatment.

3 Results

Section 3.1 develops results for the administrator's problem when there is no capacity constraint. Section 3.2 develops results for the administrator's problem in the presence of a capacity constraint.

3.1 Results without Capacity Constraints

The goal of this paper is to link the chosen cutoff κ^* to treatment effects $\Delta(\cdot)$. Therefore, I first derive necessary and sufficient conditions to characterize κ^*, in terms of $\Delta(\cdot)$ and the cost function $c(\mu) = \chi\mu$. Note that, throughout this paper, I study interior (of the the unit interval) κ^*, which is not restrictive if there exist both treated and untreated units (i.e., we are analyzing results from an RD design).

Condition 1 (Necessity). *The following necessary conditions must hold for κ^*:*

(i) MB=MC: $\Delta(\kappa^*) = c'(1 - \kappa^*) = \chi$

(ii) Increasing MB: $\Delta'(\kappa^*) \geq 0$.

Proof. Differentiate the administrator's problem (1) with respect to $\tilde{\kappa}$ to obtain (i). Note that if the derivative is negative at a candidate solution satisfying (i), the administrator would gain by not treating students just above κ^*, thereby obtaining (ii). □

[13] Note that $Y_\tau(\cdot)$ may contain negative values, which may capture a negative treatment effect or a positive opportunity cost of participating in an ineffective treatment.

[14] Notice that all bounds are "sharp", i.e., they are the smallest possible bounds given the data, i.e., the RD-estimated LATE.

Condition 1(i) will play a key role in bounding treatment effects and also provides testable implication of cutoff optimality, in that a negative LATE at the cutoff (i.e., $\Delta(\kappa^*) < 0$) would contradict Assumption 1, because $\chi > 0$. Condition 1(ii) is another testable implication of the model's assumptions that the administrator is acting optimally and with knowledge of $\Delta(\cdot)$. It can be tested using methods developed in Dong and Lewbel (2015). That is, the model's maintained Assumption 1 would be falsified if one rejected that $\Delta'(\kappa^*) \geq 0$. Condition 1 need not be sufficient; there can be multiple cutoffs satisfying it.

Assumption 2 (Unique maximand)**.** *κ^* uniquely maximizes the administrator's problem (1).*

Assumption 2 implies that $\Delta(\cdot)$ crosses $c'(\cdot)$ finitely many times and is made to simplify exposition. Note that uniqueness of κ^* implies that Condition 1(ii) should be strict (i.e., $\Delta'(\kappa^*) > 0$). To guarantee uniqueness, inspection of (1) implies two additional conditions sufficient for characterizing κ^*.

Condition 2 (Sufficiency)**.** *The fact the program was implemented implies that the total gain from treating those units was at least as large as the total costs, i.e.:*

$$\textit{Participation:} \int_{\kappa^*}^{1} \Delta(x)dx \geq c(1-\kappa^*) = \chi(1-\kappa^*). \tag{2}$$

The fact the program was not extended to $\hat{\kappa} < \kappa^*$ implies that treating these units would be sub-optimal, i.e.:

$$\int_{\hat{\kappa}}^{\kappa^*} \Delta(x)dx < c(1-\hat{\kappa}) - c(1-\kappa^*) = \chi(\kappa^* - \hat{\kappa}). \tag{3}$$

Intuitively, Condition 2 uses revealed preferences to make statements about the gains and costs of treating students who are either treated or untreated. It must be worthwhile to have treated the treated students, and it could not have been worthwhile to treat the untreated. A corollary immediately follows.

Corollary 1. *The following are globally true about $\Delta(\cdot)$:*

(i) $\Delta(\cdot)$ cannot be constant.

(ii) $\Delta(\cdot)$ is not globally monotonically decreasing in x.

Proof. If $\Delta(\cdot)$ were constant then κ^* would either be at a corner or violate Assumption 2. Condition 1(ii) already rules out $\Delta(\cdot)$ decreasing at κ^*. Consider the behavior of $\Delta(\cdot)$ for positive measures of units away from the cutoff. The second part of Condition 2 says that it must be the case that $\int_{\hat{k}}^{\kappa^*} (\Delta(x) - \chi)\, dx < 0$. Moreover, we know that $\Delta(\kappa^*) = \chi$ and that $\Delta(\cdot)$ is continuous by Assumption 1(ii). Therefore, if $\Delta(\cdot)$ were monotonically decreasing in x in any interval $[\hat{k}, k^*]$, this inequality would be violated. Similar reasoning, using the first part of Condition 2, shows that $\Delta(\cdot)$ also cannot be monotonically decreasing above κ^*. □

It is often said that there is no reason to believe treatment effects would be the same for students away from the cutoff, though this notion is not always reflected in empirical implementations. Corollary 1(i) strengthens this statement by ruling out constant treatment effects. Corollary 1(ii) is consistent with the administrator treating students above, rather than below, her chosen cutoff. Corollary 1 provides fairly weak statements about the global behavior of $\Delta(\cdot)$. Therefore, I next examine what can be deduced about averages of treatment effects for subsets of students.

Proposition 1. *The ATT is bounded below by the LATE at the treatment cutoff.*

Proof. Divide (2) by the measure of treated students $(1-\kappa^*)$ and combine this with Condition 1(i) to obtain

$$\underbrace{\frac{\int_{\kappa^*}^{1} \Delta(x)dx}{1-\kappa^*}}_{\text{ATT}} \geq \frac{\chi(1-\kappa^*)}{1-\kappa^*} = \chi = \underbrace{\Delta(\kappa^*)}_{\text{LATE at } \kappa^*} .$$

□

Note that a positive fixed cost of treatment, if known, would increase the lower bound on the ATT. I focus on the case with no fixed cost because it results in more conservative bounds and also because this section's results only require qualitative information about the marginal cost function (i.e., that is constant), not its level. A corollary immediately follows.

Corollary 2. *The ATT is positive.*

Proof. This follows directly from Proposition 1 because $\chi > 0$, by Assumption 1(i). □

Proposition 1 shows that the discontinuity-based estimate provides a lower bound for the average effect of treatment on the treated. *In other words, estimates from discontinuity-based designs will understate the effect of treatment on the treated.* Intuitively, the administrator chooses κ^* to set the marginal benefit from providing the treatment equal to the marginal cost, which is lower than the average cost of providing treatment to treated students (i.e., the marginal cost). The fact that the administrator chose to implement the program, however, implies that the gain to treating those students must have been at least as large as the total

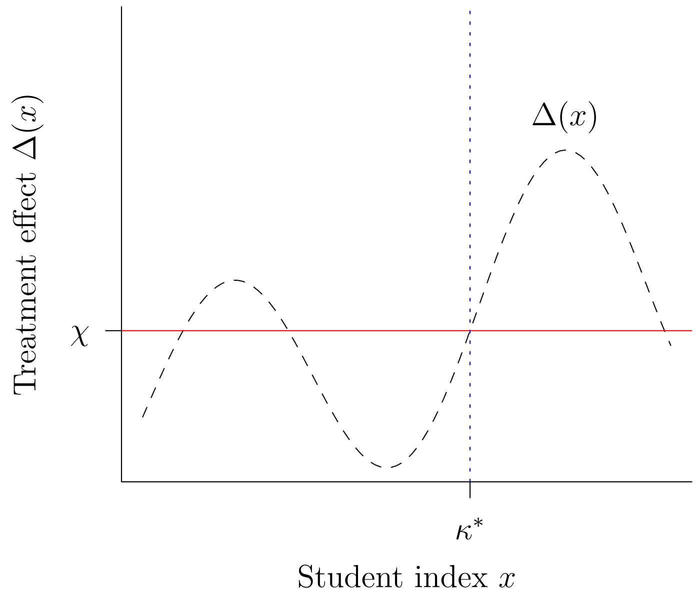

Figure 1: Example $\Delta(\cdot)$ with optimal cutoff κ^*

cost of treating them. Note that the level of the marginal cost does not need to be known by the researcher.

However, although we have quite a bit of information about averages of treatment effects $\Delta(\cdot)$ over some intervals of interest, we cannot make statements about $\Delta(x)$ for students $x \neq \kappa^*$. Figure 1 plots an example treatment effect function $\Delta(x)$ (dashed black curve) and the marginal cost of treatment (solid red horizontal line) against the student index x, which ranges from 0 to 1, and the optimal cutoff κ^* (dotted blue vertical line). This figure shows a case satisfying Conditions 1-2 where there are also untreated students with gains greater than their cost of treatment and treated students with gains smaller than their cost of treatment. Although Corollary 1(ii) rules out a treatment effect that is decreasing everywhere , it could be the case that $\Delta(\cdot)$ is decreasing for some x on either side of κ^*. Therefore, it is useful to make a statement about the average effect of extending treatment to the untreated. In particular, we can bound averages of $\Delta(\cdot)$ itself for subsets of untreated students.

Proposition 2. *There exists an informative upper bound for* $\int_a^b \Delta(x)dx$ *for* $0 \leq a < b \leq \kappa^*$.

Proof. Suppose we would like to characterize $\Delta(\cdot)$ for values less than $\hat{x} \leq \kappa^*$. Let $\hat{\mu}$ be the measure of units under consideration and split (3) into two parts at $\hat{x}$ and rearrange terms:

$$\int_{\hat{x}-\hat{\mu}}^{\hat{x}} \Delta(x)dx < c(1-(\hat{x}-\hat{\mu}))-c(1-\kappa^*)-\int_{\hat{x}}^{\kappa^*} \Delta(x)dx \Rightarrow \int_{\hat{x}-\hat{\mu}}^{\hat{x}} \Delta(x)dx < \chi(\kappa^*-(\hat{x}-\hat{\mu}))-\underline{\Delta}\,(\kappa^*-\hat{x})\,, \tag{4}$$

□

where the implication follows from Assumption 1(iii).[15]

The right side of (4) in Proposition 2 provides an upper bound for the gain from treating students $x \in [\hat{x}-\hat{\mu}, \hat{x}]$. Because we do not know $\Delta(\cdot)$, by assuming the worst possible treatment effect ($\underline{\Delta}$) we can find an upper bound for how large it could be for a measure of students $\hat{\mu}$ and satisfying the individual rationality constraint from Condition 2. Intuitively, this upper bound grows the further below the cutoff we go. To gain more intuition for Proposition 2, rearrange (4) and divide by the measure of students under consideration $\hat{\mu}$ to obtain:

$$\int_{\hat{x}-\hat{\mu}}^{\hat{x}} \frac{\Delta(x)}{\hat{\mu}}dx < (\chi-\underline{\Delta})\left(\frac{\kappa^*-\hat{x}}{\hat{\mu}}\right)+\chi. \tag{5}$$

The left side of (5) is the average treatment effect among students $x \in [\hat{x}-\hat{\mu}, \hat{x}]$. First, consider the extreme scenario where we want an upper bound for the treatment effect for student $\hat{x}$, $\Delta(\hat{x})$. Take the limit of (5) as the additional treated students go to zero:

$$\underbrace{\lim_{\hat{\mu}\to 0}\left(\int_{\hat{x}-\hat{\mu}}^{\hat{x}} \frac{\Delta(x)}{\hat{\mu}}dx\right)}_{\Delta(\hat{x})} < \lim_{\hat{\mu}\to 0}\left((\chi-\underline{\Delta})\left(\frac{\kappa^*-\hat{x}}{\hat{\mu}}\right)+\chi\right)=\infty,$$

i.e., the expression becomes uninformative when we evaluate it for measure zero of students to bound $\Delta(\cdot)$ at a point. However, consider the other extreme where $\hat{\mu}=\hat{x}$, i.e., the administrator is considering extending treatment to all students below $\hat{x}$:

$$\int_{0}^{\hat{x}} \frac{\Delta(x)}{\hat{x}}dx < (\chi-\underline{\Delta})\left(\frac{\kappa^*}{\hat{x}}\right)+\underline{\Delta}. \tag{6}$$

Equation (6) says that the upper bound on the average treatment effect among students $x \leq \hat{x}$ (the left side) grows the further $\hat{x}$ goes below κ^*, the higher is the marginal cost χ, and the lower is the lower bound $\underline{\Delta}$.

Setting the measure of students to whom the treatment is extended equal to κ^* provides the following result about the ATUT.

[15] Note this bound will be informative for all but very low values of $\underline{\Delta}$, i.e., those satisfying $\underline{\Delta} > \chi-(\overline{\Delta}-\chi)\hat{\mu}/(\kappa^*-\hat{x})$.

Corollary 3. *The ATUT is bounded above by the LATE at the treatment cutoff.*

Proof. Let $\hat{x} = \hat{\mu} = \kappa^*$ in (4) and divide through by κ^* to obtain the result:

$$\underbrace{\int_0^{\kappa^*} \frac{\Delta(x)}{\kappa^*} dx}_{\text{ATUT}} < \underbrace{\frac{c(1) - c(1-\kappa^*)}{\kappa^*} = \chi}_{>0, <\infty} = \Delta(\kappa^*), \tag{7}$$

□

where the middle term is positive from Assumption 1(i) and the last equality obtains because $\Delta(\kappa^*) = \chi$ by Condition 1(i).

Analogously to the upper bound for the ATT, although Corollary 3 bounds the average of treatment effects for all untreated students, there is no informative (i.e., greater than $\underline{\Delta}$) lower bound.

Finally, the next result shows how the prior results can be used to bound the average treatment effect (ATE).

Corollary 4. *The ATE has informative bounds.*

Proof. To form the lower bound for the ATE, note that measure κ^* units are untreated, and, by Assumption 1(iii), the treatment effect for each unit cannot be worse than $\underline{\Delta}$. Analogously, $1-\kappa^*$ units are treated, and Proposition 1 says the ATT is no smaller than $\Delta(\kappa^*)$. Integrate and sum the two parts to form $\Delta^{LB} \equiv \underline{\Delta}\kappa^* + \Delta(\kappa^*)(1-\kappa^*)$. To form the upper bound for the ATE, note that Corollary 3 says that the ATUT is no larger than $\Delta(\kappa^*)$. By Assumption 1(iii), the treatment effect for any treated unit cannot exceed $\overline{\Delta}$. Integrate and sum to form $\Delta^{UB} \equiv \Delta(\kappa^*)\kappa^* + \overline{\Delta}(1-\kappa^*)$. □

Corollary 4 shows that higher values of κ^* tighten the upper bound on the ATE while loosening the lower bound on the ATE. Intuitively, treating fewer students increases the share of untreated students, who have an upper bound of $\Delta(\kappa^*)$, while increasing the share of students with very low values of $\Delta(\cdot)$, i.e., $\underline{\Delta}$.

To summarize, optimality of the treatment cutoff κ^* implies a lower bound on the average effect of treatment on the treated (ATT) and an upper bound on the average effect of treatment on the untreated (ATUT). The average treatment effect for the population (ATE) combines the above bounds. Optimality further implies that ATUT $< \Delta(\kappa^*) \leq$ ATT. Moreover, we can contrast what can be said about $\Delta(\cdot)$ for students with indices $x > \kappa^*$ and $\Delta(\cdot)$ for student with indices $x < \kappa^*$. Because the treatment is being provided to all students in the treated group, we cannot separate how treatment effects accumulate for students $x > \kappa^*$. But the fact that the administrator is not choosing to extend (i.e., decrease) the cutoff to student $\hat{x} < \kappa^*$ provides us information for how large treatment effects can possibly be for students $x \in [\hat{x}, \kappa^*)$.

As with the constant marginal cost of treatment, optimality of κ^* implies a lower bound on the ATT and an upper bound on the ATUT when the marginal cost of treatment is instead variable, as it is in Appendix D. Most bounds remain the same if, instead of being constant, the marginal cost of treatment is nonincreasing. In particular, if the marginal cost of treatment is constant or decreasing then it must be the case that ATUT $< \Delta(\kappa^*) \leq$ ATT.

Though the ATT and ATUT are respectively bounded below and above by the cutoff LATE when marginal costs are nonincreasing, the LATE does not bound them when the marginal cost of treatment is strictly increasing. It is worthwhile to discuss the intuition behind the variable marginal cost results here. Suppose we rotated the marginal cost of treatment curve in Figure 1 clockwise about the point (κ^*, χ), to model a strictly increasing marginal cost of treatment. This would mean that the average gain to having treated the treated (left side of (2) in Condition 2) could have been smaller than the marginal cost of treatment at the cutoff (χ) and still warrant treatment. Analogously, the gain to treating the untreated (left side of (3)) in Condition 2) could have been larger than the marginal cost of treatment at the cutoff and still warrant non-treatment. Then, a strictly increasing marginal cost of treatment would decrease the lower bound on the ATT and increase the upper bound on the ATUT, leading to looser bounds on the ATE as well. The opposite would be true for a strictly decreasing marginal cost of treatment. It is important, however, to distinguish an increasing marginal cost of treatment from a binding capacity constraint, the latter of which I examine in Section 3.2.

3.2 Results with Capacity Constraints

Suppose now that the administrator faces a capacity constraint, $\overline{\mu}$. Then, the administrator solves

$$\max_{\widetilde{\kappa}} \beta \left(\int_{\widetilde{\kappa}}^{1} \Delta(x) dx \right) - c\,(1 - \widetilde{\kappa}) \qquad (\bar{1})$$

$$\text{s.t.} \quad 1 - \widetilde{\kappa} \leq \overline{\mu},$$

where we continue to maintain Assumptions 1-2 from problem (1).

If the desired (i.e., unconstrained) measure of treated students does not exceed capacity, i.e., $\mu^*(\equiv 1 - \kappa^*) \leq \overline{\mu}$, then the constraint does not bind, and the optimal cutoff κ^* and resulting analysis are unaffected. This means the results from Section 3.1 apply here as well. By definition, if the capacity constraint is binding the measure of students treated must be strictly less than the desired measure of students treated, meaning the optimal cutoff and results may differ from those in Section 3.1.

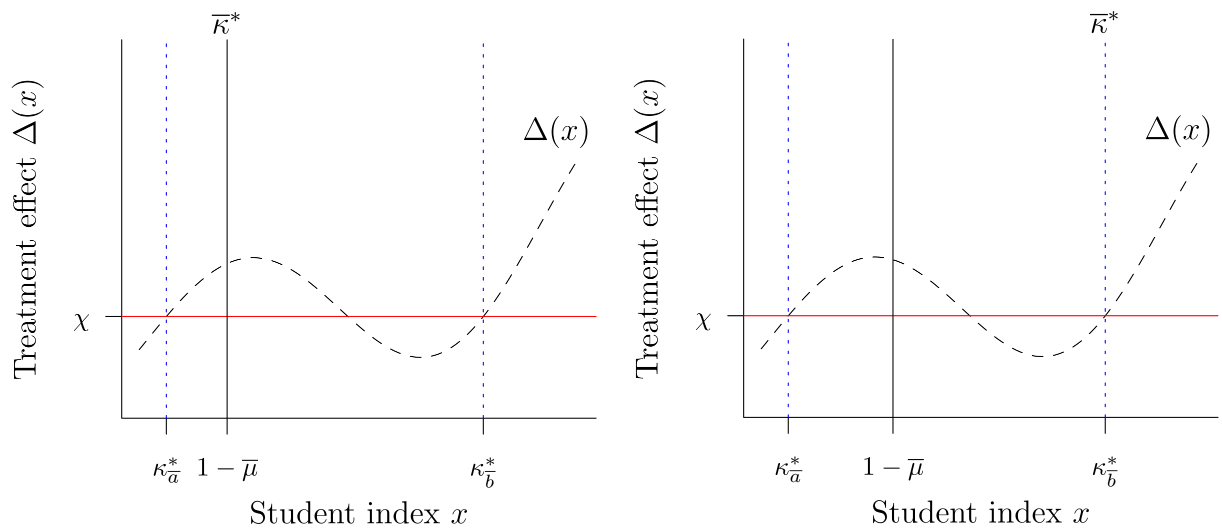

(a) Case $\overline{a}$

(b) Case $\overline{b}$

Figure 2: Types of binding capacity constraints

Let $\overline{\kappa}^*$ denote the binding solution to the constrained problem ($\bar{1}$). There are two types of potential cases corresponding to optimal cutoffs in the presence of a binding capacity constraint, depicted in Figure 2. The cutoff $\kappa^*_{\overline{a}}$ indicates the optimal cutoff were the administrator unconstrained. Figure 2a indicates the administrator's constrained-optimal cutoff in Case $\overline{a}$, where $\overline{\mu}$ is such that the administrator would set $\overline{k}^* = 1 - \overline{\mu}$, i.e., the capacity constraint is locally binding. For example, Case $\overline{a}$ would apply if $1 - \overline{\mu}$ was just above $\kappa^*_{\overline{a}}$. However, if $\overline{\mu}$ decreased by enough then the gains from extending treatment to units up until capacity may not be worth the cost. This is Case $\overline{b}$, depicted in Figure 2b, where a more severe capacity constraint would cause the administrator to instead set the constrained-optimal cutoff to $\overline{\kappa}^* = \kappa^*_{\overline{b}} > 1 - \overline{\mu}$. In Case $\overline{b}$, the capacity constraint is binding, but not locally. Note that, although the $\Delta(\cdot)$ function only has two possible solutions, $\kappa^*_{\overline{a}}$ and $\kappa^*_{\overline{b}}$, as drawn in Figure 2, this is not a necessary assumption, nor is it exploited in the following analysis. Rather, the following results rely only on knowledge of $\overline{\mu}$, $\overline{\kappa}^*$, $\Delta(\overline{\kappa}^*)$, and, potentially, knowledge of $\kappa^*_{\overline{a}}$.[16] As will be clear later, this produces more conservative bounds.

[16]For example, it could be the case that $\Delta(\cdot)$ crosses χ twice between $\kappa^*_{\overline{a}}$ and $\overline{\kappa}^*$ in Figure 2a.

For example, I do not tighten the lower bound on the ATUT by assuming $\Delta(\cdot)$ is on average above the marginal cost of treatment between $\kappa^*_{\overline{a}}$ and $1-\overline{\mu}$.

This section focuses on averages of $\Delta(\cdot)$ when the capacity constraint binds, defined as follows.

Definition $\overline{1}$ (Treatment effects of interest)**.** *When the capacity constraint is binding, define:*

- *Average Effect of Treatment on the Treated ($\overline{ATT}$): $\int_{\overline{\kappa}^*}^1 \frac{\Delta(x)}{1-\overline{\kappa}^*}dx$,*
- *Average Effect of Treatment on the Untreated ($\overline{ATUT}$): $\int_0^{\overline{\kappa}^*} \frac{\Delta(x)}{\overline{\kappa}^*}dx$, and*
- *Average Treatment Effect among all units ($\overline{ATE}$): $\int_0^1 \Delta(x)dx$.*

The conditional treatment effects in Definition $\overline{1}$ differ from those in Definition 1 because they use the capacity-constrained-optimal cutoff $\overline{\kappa}^*$, meaning they may correspond to different groups of treated students. However, the capacity-constrained ATE is the same as the unconstrained one, i.e., $\overline{\text{ATE}} = \text{ATE}$.

I begin by adapting Condition 1.

Condition $\overline{1a}$ (Necessity)**.** *If $\overline{\kappa}^* = 1-\overline{\mu}$ solves the administrator's capacity-constrained problem ($\overline{1}$), then $\Delta(\overline{\kappa}^*) > \chi$.*

Proof. Ignoring the measure-zero case(s) where $\Delta(\overline{\kappa}^*) = \chi$, this follows from Condition 1(i) and the fact that the capacity constraint is binding, i.e., the administrator would have liked to treat more students—in particular, lower the cutoff infinitesimally. □

Unlike Case $\overline{a}$, in Case $\overline{b}$ the administrator's choice of cutoff is not locally binding, resulting in the same necessary conditions as in the unconstrained problem.

Condition $\overline{1b}$ (Necessity)**.** *If $\overline{\kappa}^* = \kappa^*_{\overline{b}}$ solves the administrator's capacity-constrained problem ($\overline{1}$), then the following necessary conditions must hold:*

(i) MB=MC: $\Delta(\overline{\kappa}^*) = \chi$

(ii) Increasing MB: $\Delta'(\overline{\kappa}^*) \geq 0$.

Proof. Identical to Condition 1. □

Conditions $\overline{1a}$ and $\overline{1b}$(i) have the same testable implication as that derived from Condition 1(i)—that the model can be falsified if the RD estimate is negative. Condition 1(ii)—that the treatment effect derivative is nondecreasing at the cutoff—applies when the constraint binds in Case $\overline{b}$ but need not apply in Case $\overline{a}$. However, an alternative testable implication of optimality can be deduced by combining Condition 1(i) with Condition $\overline{1a}$: $\Delta(\overline{\kappa}^*) > \Delta(\kappa^*_{\overline{a}})$. This could be tested by using data from two years during which one thought the model parameters $\Delta(\cdot)$ and χ did not change, one where the constraint was Case $\overline{a}$-binding (allowing estimation of $\Delta(\overline{\kappa}^*)$) and another where the administrator's budget increased, say, due to a large RD estimate stemming from the binding constraint in the first year (allowing estimation of $\Delta(\kappa^*_{\overline{a}})$). What may look like a lack of "scale-up" for a program may simply reflect that the marginal benefit at the treatment cutoff is smaller if the constraint is no longer binding. Further note that one could use the variation in capacity constraints to trace out $\Delta(\cdot)$ in Case $\overline{a}$. Interestingly, these last two results are only implementable if the constraint Case-$\overline{a}$-binds in at least one period—otherwise the cutoff, and resulting LATE, would be the same for both periods. Perhaps counterintuitively, we may actually learn more when the constraint binds in at least one year.

The following participation condition must hold in Case $\overline{a}$.

Condition $\overline{2a}$. *Suppose $\overline{\kappa}^* = 1-\overline{\mu}$ solves ($\overline{1}$). The fact the program was implemented implies that the total gain from treating those units was larger than the total costs, i.e.:*

$$\textit{Participation:} \quad \int_{\overline{\kappa}^*}^{1} \Delta(x)dx > \chi(1-\overline{\kappa}^*), \tag{$\overline{2a}$}$$

where the strict inequality follows from combining Condition $\overline{1a}$ with (2) from Condition 2.

Note that Condition $\overline{2a}$ does not have an analogue to the second part of Condition 2. This is because the administrator would have wanted to treat the inframarginal students (i.e., those between the unconstrained- and constrained-optimal cutoffs); otherwise the constraint would not have been binding. In Case $\overline{b}$ the following condition characterizes the cutoff.

Condition $\overline{2b}$. *Suppose $\overline{\kappa}^* = \kappa^*_{\overline{b}}$ solves ($\overline{1}$). The fact the program was implemented implies that the total gain from treating those units was at least as large as the total costs, i.e.:*

$$\textit{Participation:} \quad \int_{\overline{\kappa}^*}^{1} \Delta(x)dx \geq \chi(1-\overline{\kappa}^*). \tag{$\overline{2b}$}$$

The fact the program was not extended to $\hat{\kappa} \in [1-\overline{\mu}, \kappa^*)$ implies that treating these units *would be sub-optimal, i.e.:*

$$\int_{\hat{\kappa}}^{\overline{\kappa}^*} \Delta(x)dx < \chi(\overline{\kappa}^* - \hat{\kappa}). \tag{$\overline{3b}$}$$

The first part of Condition $\overline{2b}$ is the same as the first part of Condition 2. The second part, ($\overline{3b}$), differs from (3) in that it only applies to candidate values $\hat{\kappa}$ no less than $1 - \overline{\mu}$. Intuitively, this reflects the fact that the administrator chose to set the cutoff at $\kappa^*_{\overline{b}}$ instead of $1 - \overline{\mu}$.

Neither Corollary 1 nor Proposition 1 obtain in Case $\overline{a}$ but they, or an analogue, can be obtained for Case $\overline{b}$.[17]

Corollary $\overline{1b}$. *Suppose $\overline{\kappa}^* = \kappa^*_{\overline{b}}$ solves ($\overline{1}$). The following are globally true about $\Delta(\cdot)$:*

(i) $\Delta(\cdot)$ cannot be constant.

(ii) $\Delta(\cdot)$ is not globally monotonically decreasing in x.

Proof. Analogous to proof of Corollary 1. □

Next, I examine what can be learned about $\Delta(\cdot)$ for subsets of students, starting with the treated.

Proposition $\overline{1b}$. *Suppose $\overline{\kappa}^* = \kappa^*_{\overline{b}}$ solves ($\overline{1}$). The $\overline{ATT}$ is bounded below by the LATE at the treatment cutoff.*

Proof. Identical to Proposition 1. □

Corollary $\overline{2a}$. *Suppose $\overline{\kappa}^* = 1 - \overline{\mu}$ solves ($\overline{1}$). The $\overline{ATT}$ is positive.*

Proof. Because the marginal cost of treatment is positive (Assumption 1(i)), ($\overline{2a}$) implies that

$$\int_{\overline{\kappa}^*}^{1} \Delta(x)dx > \chi(1 - \overline{\kappa}^*) > 0.$$

Divide through by $(1 - \overline{\kappa}^*)$ to obtain the result:

$$\underbrace{\int_{\overline{\kappa}^*}^{1} \frac{\Delta(x)}{(1 - \overline{\kappa}^*)}dx}_{\overline{\text{ATT}}} > \underbrace{\chi}_{\text{avg. cost of treating treated}} > 0.$$ □

[17]Corollary 1(ii) would also obtain in Case $\overline{a}$ if the administrator could choose which side of the cutoff to treat.

Corollary $\overline{2a}$ shows that we can bound the $\overline{\text{ATT}}$ from below by zero. This lower bound is looser than that for the ATT—the unconstrained analogue of $\overline{\text{ATT}}$—obtained in Proposition 1, which is the unconstrained RD estimate, which was shown to be positive. This because, though we know that both $\Delta(\overline{\kappa}^*)$ and the $\overline{\text{ATT}}$ are nonnegative, we do not have enough information to order them.For example, consider $\Delta(\cdot)$ such that $\Delta'(x) < 0$ for all $x \geq \overline{\kappa}^*$; in this case the RD estimate would bound the $\overline{\text{ATT}}$ *from above*. However, we could use ($\overline{2a}$) to tighten this bound if information about χ, the marginal cost of treatment, were available.[18] In contrast, in Case $\overline{b}$ we obtain the tighter bound obtained in the unconstrained problem.

Corollary $\overline{2b}$. *Suppose $\overline{\kappa}^* = \kappa^*_{\overline{b}}$ solves ($\overline{1}$). The $\overline{ATT}$ is positive.*

Proof. Identical to Corollary 2. □

Next, I focus on what can be learned about $\Delta(\cdot)$ for subsets of untreated students, starting with Case $\overline{a}$.

Proposition $\overline{2a}$. *Suppose $\overline{\kappa}^* = 1 - \overline{\mu}$ solves ($\overline{1}$). There exists an informative upper bound for $\int_a^b \Delta(x)dx$ for $0 \leq a < b \leq \overline{\kappa}^*$, if $\kappa^*_{\overline{a}}$ is known and $a < \kappa^*_{\overline{a}}$.*

Proof. If $b \leq \kappa^*_{\overline{a}}$, then by Proposition 2 the upper bound is $\chi(\kappa^*_{\overline{a}} - a) - \underline{\Delta}(\kappa^*_{\overline{a}} - b)$; if χ is unknown then apply Condition $\overline{1a}$ to form the upper bound $\Delta(\overline{\kappa}^*)(\kappa^*_{\overline{a}} - a) - \underline{\Delta}(\kappa^*_{\overline{a}} - b)$.

If $b > \kappa^*_{\overline{a}}$, then we can split the integral into two parts: $\int_a^{\kappa^*_{\overline{a}}} \Delta(x)dx$, which, applying the result for $b \leq \kappa^*_{\overline{a}}$, has an upper bound of $\Delta(\overline{\kappa}^*)(\kappa^*_{\overline{a}} - a)$, and $\int_{\kappa^*_{\overline{a}}}^b \Delta(x)dx$, which has an upper bound of $\overline{\Delta}(b - \kappa^*_{\overline{a}})$. Sum to form the upper bound $\Delta(\overline{\kappa}^*)(\kappa^*_{\overline{a}} - a) + \overline{\Delta}(b - \kappa^*_{\overline{a}})$.

Note that if $a \geq \kappa^*_{\overline{a}}$, then all students $x \in [a, b]$ are inframarginal, meaning the upper bound would be $\overline{\Delta}(b - a)$, i.e., uninformative. □

Proposition $\overline{2b}$. *Suppose $\overline{\kappa}^* = \kappa^*_{\overline{b}}$ solves ($\overline{1}$). There exists an informative upper bound for $\int_a^b \Delta(x)dx$ for $0 \leq a < b \leq \overline{\kappa}^*$ if either (i) $b \leq 1 - \overline{\mu}$ and $\kappa^*_{\overline{a}}$ is known and $a < \kappa^*_{\overline{a}}$; or (ii) $b \in (1 - \overline{\mu}, \overline{\kappa}^*]$.*

Proof. First consider the case (i), where $b \leq 1 - \overline{\mu}$. Then we can apply the bound from Proposition $\overline{2a}$, which would be tighter because $\Delta(\overline{\kappa}^*)$ is smaller in Case $\overline{b}$.

There are two subcases based on a in case (ii), where $b \in (1 - \overline{\mu}, \overline{\kappa}^*]$:

Subcase (ii.i): If $a \geq 1 - \overline{\mu}$ the fact the administrator did not extend treatment from $\overline{\kappa}^*$ to a implies that $\int_a^b \Delta(x)dx + \int_b^{\overline{\kappa}^*} \Delta(x)dx < \chi(\overline{\kappa}^* - a)$, which implies that $\int_a^b \Delta(x)dx < \chi(\overline{\kappa}^* - a) - \underline{\Delta}(\overline{\kappa}^* - b)$, which, by Condition $\overline{1b}$(i), returns a lower bound of $\Delta(\overline{\kappa}^*)(\overline{\kappa}^* - a) - \underline{\Delta}(\overline{\kappa}^* - b)$.

[18] Assuming β was also known.

Subcase (ii.ii): If $a < 1-\overline{\mu}$ use the bound from subcase (ii.i), setting $a = 1-\overline{\mu}$, to bound $\int_{1-\overline{\mu}}^{b} \Delta(x)dx$. Note that $\int_{a}^{1-\overline{\mu}} \Delta(x)dx \leq \overline{\Delta}(1-\overline{\mu}-a)$. Sum to form the upper bound $\Delta(\overline{\kappa}^*)(\overline{\kappa}^* - 1 + \overline{\mu}) - \underline{\Delta}(\overline{\kappa}^* - b) + \overline{\Delta}(1-\overline{\mu}-a)$. □

Proposition $\overline{2a}$ applies to subsets of untreated students, and can be applied to two subsets of particular interest. Define $\text{ATUT}_a = (\int_0^{\kappa_{\overline{a}}^*} \Delta(x)dx)/\kappa_{\overline{a}}^*$, i.e., the unconstrained effect of treatment on the untreated from Definition 1.

Corollary $\overline{3a}$. *Suppose $\overline{\kappa}^* = 1-\overline{\mu}$ solves ($\overline{1}$). Upper bounds can be obtained for treatment effects for groups of untreated units:*

(i) $ATUT_a$ is bounded from above by the LATE at the treatment cutoff.

(ii) There is an informative upper bound for the $\overline{ATUT}$, if $\kappa_{\overline{a}}^*$ is known.

Proof. Part (i) follows by combining Condition $\overline{1a}$ with Corollary 3 and Condition 1(i).

To show part (ii), apply Proposition $\overline{2a}$, setting $a = 0$ and $b = \overline{\kappa}^*$, and divide by $\overline{\kappa}^*$, the measure of untreated units, obtaining $[\Delta(\overline{\kappa}^*)\kappa_{\overline{a}}^* + \overline{\Delta}(\overline{\kappa}^* - \kappa_{\overline{a}}^*)]/\overline{\kappa}^*$. □

Corollary $\overline{3a}$(i) is the analogue of Corollary 3, delivering an upper bound to the unconstrained ATUT_a. Corollary $\overline{3a}$(ii) shows that knowledge of how severe the capacity constraint is, i.e., $(\kappa_{\overline{a}}^*/\overline{\kappa}^*)$, tightens the upper bound on the $\overline{\text{ATUT}}$. The extent to which it can be tightened depends on $\kappa_{\overline{a}}^*$, but this bound is still looser than in the unconstrained case, because $\Delta(\overline{\kappa}^*) > \Delta(\kappa_{\overline{a}}^*)$. Intuitively, if the desired measure of treated students is not known then we would have to set $\kappa_{\overline{a}}^* = 0$ to maximize the upper bound, returning the uninformative upper bound of $\overline{\Delta}$.

Similarly, Proposition $\overline{2b}$ can also be applied to these subsets of untreated students.

Corollary $\overline{3b}$. *Suppose $\overline{\kappa}^* = \kappa_{\overline{b}}^*$ solves ($\overline{1}$). Upper bounds can be obtained for treatment effects for groups of untreated units:*

(i) $ATUT_a$ is bounded from above by the LATE at the treatment cutoff.

(ii) There is an informative upper bound for the $\overline{ATUT}$, which if tighter if $\kappa_{\overline{a}}^*$ is known.

Proof. Part (i) follows by combining Condition $\overline{1b}$ with Corollary 3 and Condition 1(i).

To show part (ii), note that ($\overline{3b}$) implies that $\int_{1-\overline{\mu}}^{\overline{\kappa}^*} \Delta(x)dx < \Delta(\overline{\kappa}^*)(\overline{\kappa}^* - (1-\overline{\mu}))$. By Assumption 1(iii), the treatment effect for any treated unit cannot exceed $\overline{\Delta}$, which if $\kappa_{\overline{a}}^*$ is not known is the upper bound for the average of $\Delta(\cdot)$ for $x \in [0, 1-\overline{\mu}]$. This results in an $\overline{\text{ATUT}}$ upper bound of $[\Delta(\overline{\kappa}^*)(\overline{\kappa}^* - (1-\overline{\mu})) + \overline{\Delta}(1-\overline{\mu})]/\overline{\kappa}^*$.

If $\kappa_{\overline{a}}^*$ is known, then $\int_0^{\kappa_{\overline{a}}^*} \Delta(x)dx < \Delta(\overline{\kappa}^*)\kappa_{\overline{a}}^*$ by (3), which results in a tighter $\overline{\text{ATUT}}$ upper bound of $[\Delta(\overline{\kappa}^*)(\overline{\kappa}^* - (1-\overline{\mu}) + \kappa_{\overline{a}}^*) + \overline{\Delta}(1-\overline{\mu}-\kappa_{\overline{a}}^*)]/\overline{\kappa}^*$. □

The bound from Corollary $\overline{3b}$(i) is tighter than that from Corollary $\overline{3a}$(i) because in Case $\bar{b}$, $\Delta(\overline{\kappa}^*) = \chi$, in contrast to Case $\overline{a}$, where $\Delta(\overline{\kappa}^*) > \chi$. Corollary $\overline{3b}$(ii) shows that, in Case $\bar{b}$, the upper bound on the $\overline{\text{ATUT}}$ gets tighter as capacity $\overline{\mu}$ increases. Note that the $\overline{\text{ATUT}}$ upper bound is tighter in Case $\bar{b}$ than in Case $\overline{a}$.

Finally, we can combine the previous results to bound the ATE, starting with Case $\overline{a}$.

Corollary $\overline{4a}$. *Suppose $\overline{\kappa}^* = 1 - \overline{\mu}$ solves ($\bar{1}$). The ATE has an informative lower bound when the capacity constraint binds, if $\underline{\Delta} < 0$. If $\kappa^*_{\overline{a}}$ is known, then the ATE also has an informative upper bound.*

Proof. By Corollary $\overline{2a}$, the lower bound on the ATE for units with $x \geq \overline{\kappa}^*$ is 0. This increases the ATE lower bound from $\underline{\Delta}$ to $\underline{\Delta}\overline{\kappa}^*$.

If $\kappa^*_{\overline{a}}$ is known, then by Corollary $\overline{3a}$(ii) the upper bound on the total gain for the untreated is $\Delta(\overline{\kappa}^*)\kappa^*_{\overline{a}} + \overline{\Delta}(\overline{\kappa}^* - \kappa^*_{\overline{a}})$, reducing the ATE upper bound from $\overline{\Delta}$ to $\Delta(\overline{\kappa}^*)\kappa^*_{\overline{a}} + \overline{\Delta}(1 - \kappa^*_{\overline{a}})$. □

There is no informative upper bound (i.e., less than $\overline{\Delta}$) for the ATE if $\kappa^*_{\overline{a}}$ is unknown. On the other hand, if χ and $\kappa^*_{\overline{a}}$ are both known then the ATE bounds in Case $\overline{a}$ would be the same as in Section 3.1.

Corollary $\overline{4b}$. *Suppose $\overline{\kappa}^* = \kappa^*_{\bar{b}}$ solves ($\bar{1}$). The ATE has informative lower and upper bounds. If $\kappa^*_{\overline{a}}$ is known, then the upper bound is tighter.*

Proof. First assume $\kappa^*_{\overline{a}}$ is not known. To form the lower bound for the ATE, note that measure κ^* units are untreated, and, by Assumption 1(iii), the treatment effect for each unit cannot be worse than $\underline{\Delta}$. Analogously, $1 - \kappa^*$ units are treated, and ($\overline{2b}$) implies the ATT is no smaller than $\Delta(\overline{\kappa}^*)$. Integrate and sum the two parts to form $\Delta^{LB} \equiv \underline{\Delta}\overline{\kappa}^* + \Delta(\overline{\kappa}^*)(1 - \overline{\kappa}^*)$.

To form the upper bound for the ATE, use the expression for the upper bound of $\overline{\text{ATUT}}$ from Corollary $\overline{3b}$(ii) and the fact that the upper bound on the treated students $x \in [\overline{\kappa}^*, 1]$ is $\overline{\Delta}$ (by Assumption 1(iii)), to integrate and sum to form $\Delta^{UB} \equiv \Delta(\overline{\kappa}^*)(\overline{\kappa}^* - (1 - \overline{\mu})) + \overline{\Delta}(1 - (\overline{\kappa}^* - (1 - \overline{\mu})))$.

If $\kappa^*_{\overline{a}}$ is known, then, as in Corollary $\overline{3b}$(ii), we can also decrease the ATE upper bound by noting that $\text{ATUT}_a \leq \Delta(\overline{\kappa}^*)$, shifting the mass $\kappa^*_{\overline{a}}$ from having an upper bound of $\overline{\Delta}$ to $\Delta(\overline{\kappa}^*)$ and resulting in a tighter upper bound of $\Delta(\overline{\kappa}^*)(\overline{\kappa}^* - (1 - \overline{\mu}) + \kappa^*_{\overline{a}}) + \overline{\Delta}(1 - (\overline{\kappa}^* - (1 - \overline{\mu})) - \kappa^*_{\overline{a}})$. □

Intuitively, in contrast to Case $\overline{a}$, there are informative lower *and* upper bounds on the ATE in Case $\bar{b}$even if $\kappa^*_{\overline{a}}$ is unknown, as the decision to not treat students between $1 - \overline{\mu}$ and $\kappa^*_{\bar{b}}$ contains useful information. Knowledge of $\kappa^*_{\overline{a}}$, though not necessary to produce informative bounds in Case $\bar{b}$, would further tighten them.

In summary, comparing Case $\overline{a}$ with the unconstrained problem, bounds on $\Delta(\cdot)$ are looser when the capacity constraint binds. The lower bound on the $\overline{\text{ATT}}$ is lower—it is zero instead of the RD-estimate LATE. The RD estimate bounds the ATUT, but not the $\overline{\text{ATUT}}$, from above. Table 3 in Appendix E presents bounds on the average effect of treatment on the treated, untreated, and population in the unconstrained and capacity-constrained cases. Figure 6 shows these bounds for the example shown in Figure 2. Knowledge of how severe the capacity constraint is—which as shown in Section 3.3 could be obtained by comparing budget requests and realized allocations—tightens the upper bound on the $\overline{\text{ATUT}}$ and, consequently, the ATE. Additionally, though the derivative-based test of cutoff optimality does not apply when the constraint Case-$\overline{a}$-binds—it would still apply in a Case-$\overline{b}$-binding constraint—the nonnegative LATE testable implication still does apply and there is a new testable implication of optimality (which could be applied if one had access to data for binding and non-binding periods), and a new policy-relevant result that may help explain why it is difficult to "scale up" successful programs to larger populations (see, e.g., Elmore (1996) and Sternberg et al. (2006)).

3.3 How to Distinguish Unconstrained and Constrained Cases

Computing the relevant bounds requires knowing whether the administrator's chosen cutoff corresponds to the unconstrained case, Case $\overline{a}$, or Case $\overline{b}$. This could be accomplished by gathering data on a (i) *requested budget*, (ii) *approved/allocated budget*, and (iii) *realized budget*, i.e., the amount actually used by the administrator. If the requested and allocated budgets were the same we could surmise the administrator was not constrained and apply bounds from Section 3.1; otherwise, we could surmise the capacity constraint was binding. If the administrator used the entire allocated budget, we could surmise that $\overline{\kappa}^*$ satisfies Case $\overline{a}$, i.e., the capacity constraint was locally binding. If, instead, the administrator did not use the entire allocated budget we could surmise that $\overline{\kappa}^*$ satisfies Case $\overline{b}$. Such data could be obtained via analysis of, say, a state's budget process. For example, in New York state, agencies make budget requests, or proposals, which after being amended are included in an executive budget. Then, the following fiscal year there is an end-of-cycle report on the realized amounts.[19] Alternatively, one could find a situation where the cutoff had been chosen before capacity constraints were implemented, or even obtain data on $\kappa^*_{\overline{a}}$ from credible comparison groups.

Results from Section 3.2 show that, in addition to identifying the relevant case as discussed above, bounds can be tightened in both Case $\overline{a}$ and Case $\overline{b}$ if we know the uncon-

[19] `https://www.budget.ny.gov/guide/brm/item2.html`

strained optimal cutoff, i.e., $\kappa^*_{\overline{a}}$, even when the marginal cost of treatment χ is unknown. For example, in Case $\overline{a}$ one could compare an administrator's budget request with the actual amount expended and exploit the fact that the fraction, $(\kappa^*_{\overline{a}}/\overline{\kappa}^*)$, which is required to compute the bound, is a known function of the ratio of requested and realized budgets, because the unknown marginal cost of treatment cancels when computing the requested/realized budget ratio.[20]

4 Applications

This section shows how this paper's theoretical results can be used to extend findings from regression-discontinuity designs. There are two applications, which happen to be in the economics of education and examine contexts where it seems reasonable to expect that program administrators had information about the gains and costs of treatment.

Recall that the administrator's objective depends on the total gain from treatment in the baseline model presented in Section 2. This specification is a good fit for many applications of interest, in particular, the applications studied here, which study either wages directly, or measures of human capital, such as GPA or standardized test scores. This is because, given a rental rate for human capital, maximizing human capital, maximizing wages, and maximizing output may be viewed as equivalent, meaning the objective considered here corresponds to the efficient allocation.

To most fully illustrate the theoretical results, I first examine a context where it seems likely that the administrator's capacity constraint binds; this is followed by a context where the constraint is likely not binding. I check the model implication that the treatment effect at the cutoff is nonnegative for both applications. For the latter, unconstrained, application, I also conduct the test whether the treatment effect is increasing at the cutoff.[21] Reassuringly, we cannot reject that the cutoff was chosen optimally by an informed administrator in each of the (three) falsification tests.

Two of the studies employ fuzzy designs, so I first show how the earlier results pertaining to sharp designs generalize here.[22] Some new notation is necessary. Let $\omega(x)$ denote the administrator's intended treatment group for student with index x.[23] For example, if students

[20]Let the requested budget be $B^* = (1-\kappa^*)\chi$ and the realized budget (which, in Case $\overline{a}$ would also be the approved budget) be $\overline{B}^* = (1-\overline{\kappa}^*)\chi$. We can compute $B^*/\overline{B}^* = ((1-\kappa^*_{\overline{a}})\chi)/((1-\overline{\kappa}^*)\chi) = (1-\kappa^*_{\overline{a}})/(1-\overline{\kappa}^*)$, which permits the solution for $\kappa^*_{\overline{a}}$, even when marginal cost χ is unknown.

[21] Recall that this condition does not hold for a Case-$\overline{a}$-binding constraint.

[22]These results are derived for when the capacity constraint is not binding; analogous results obtain when the constraint binds.

[23]Recall that students are distributed uniformly over $[0, 1]$.

with indices x above cutoff κ are targeted for a program, then $\omega(x) = 1$ for $x \geq \kappa$ and $\omega(x) = 0$ for $x < \kappa$. The probability of being treated ($\tau = 1$) depends on ω according to $\rho_\omega = \Pr\{\tau = 1|\omega\}$.[24] In a fuzzy design, $0 \leq \rho_0 < \rho_1 \leq 1$, i.e., not all students targeted for treatment are necessarily treated and some students not targeted for treatment may be treated. The fuzzy design requires the probability of treatment to increase discontinuously at the cutoff (Hahn et al. (2001)). This notation can also capture a sharp design when $\rho_0 = 0$ and $\rho_1 = 1$. In a fuzzy design, the administrator chooses the treatment cutoff $\widetilde{\kappa}$ to maximize her expected objective:[25]

$$\max_{\widetilde{\kappa}} \left(\rho_0 \int_0^{\widetilde{\kappa}} (\Delta(x) - \chi)dx + \rho_1 \int_{\widetilde{\kappa}}^1 (\Delta(x) - \chi)dx) \right). \tag{8}$$

The optimal cutoff κ^* is characterized by $(\rho_1 - \rho_0)\Delta(\kappa^*) = (\rho_1 - \rho_0)\chi$, implying that $\Delta(\kappa^*) = \chi$. Note that this condition is identical to Condition 1(i) for the sharp design. Moreover, multiplying through by ρ_ω shows that the fuzzy design returns exactly the same bounds for the ATT and ATUT as does the sharp design when ρ_ω are constant within treatment status.[26] Appendix A shows that the mean effect of intending-to-treat among the treated (ITT) can be bounded when treatment probabilities ρ_ω depend on x.

The tests of model assumptions can be described using a sharp design without any loss of generality.[27] In particular, the derivative sign test implied by Condition 1(ii) is the same. Therefore, I use the sharp design to show how we can test the model assumptions. Suppose students with index $x \geq \kappa^*$ were treated. In this context, the assumptions that the administrator knows $\Delta(\cdot)$ and is acting optimally would be rejected if we found that either $\Delta'(\kappa^*) < 0$, because the administrator would gain by increasing the cutoff and avoid treating inframarginal students with gains lower than that for students at the cutoff, or $\Delta(\kappa^*) < 0$, because the marginal cost of treatment is positive, contradicting optimality of treating students at κ^*. As commonly assumed in regression-discontinuity designs, assume the expected outcome for a student with index x, $Y_{\tau(x)}(x)$, depends on treatment status $\tau(x)$

[24] That is, treatment probability only depends on x through $\omega(x)$. The treatment probabilities could be measured by computing average treatment rates on either side of the cutoff.

[25] As was the case in the theoretical model, β has been set to 1. Recall that the marginal costs of treatment are assumed to be constant in the model. Though estimates of cost functions are not widely available, I was able to find evidence supporting this assumption for the applications studying university outcomes. This evidence is presented on page 26.

[26] Note that if the administrator were allowed to choose (ρ_0, ρ_1), subject to the constraint $0 \leq \rho_0 < \rho_1 \leq 1$, she would choose the sharp design because ATUT $< \chi \leq$ ATT, which means the administrator would always want to shift treatment probability from units below the cutoff to those above. Therefore, interior values of (ρ_0, ρ_1) must reflect a technological constraint precluding perfect enforcement (i.e., $\rho_1 < 1$) or exclusion (i.e., $\rho_0 > 0$).

[27] Dong and Lewbel (2015) show that a similar result holds for fuzzy designs.

and the running variable $(x - \kappa^*)$ according to the following statistical relationship:[28]

$$Y_{\tau(x)}(x) = \alpha_0 + \alpha_1(x - \kappa^*) + \alpha_2\tau(x) + \alpha_3\tau(x)(x - \kappa^*),$$

and the observed outcome for student i, $\check{Y}_i$, measures $Y_{\tau(x_i)}(x_i)$ with an independent measurement error ϵ_i according to:

$$\check{Y}_i = Y_{\tau(x)}(x) + \epsilon_i = \alpha_0 + \alpha_1(x - \kappa^*) + \alpha_2\tau(x) + \alpha_3\tau(x)(x - \kappa^*) + \epsilon_i. \tag{9}$$

The estimate of the LATE at the treatment cutoff is $\widehat{\Delta}(\kappa^*) = \widehat{\alpha}_2$. Dong and Lewbel (2015) show that the estimate of the treatment effect derivative at the cutoff here would be $\widehat{\Delta'}(\kappa^*) = \widehat{\alpha}_3$. Therefore, the model has a testable implication, i.e., is falsifiable, because using $\widehat{\alpha}_3$ to test the null hypothesis $H_0 : \alpha_3 \geq 0$, versus the alternative hypothesis $H_1 : \alpha_3 < 0$, amounts to a test of the model assumptions. Evidence strong enough to reject the null that $\alpha_3 \geq 0$ would cast doubt on the validity of Assumption 1. Moreover, evidence strong enough to reject the null hypothesis that $\alpha_2 \geq 0$ would also cast doubt on the validity of the model assumptions.

4.1 Hoekstra (2009): "The Effect of Attending the Flagship State University on Earnings: A Discontinuity-based Approach"

This section applies this paper's results to Hoekstra (2009), who studies the effect of attending a flagship public university on subsequent mean wages for a sample of white males between the ages of 28 and 33. The objective considered in (1), where the administrator seeks to maximize the amount gained (i.e., increase in wages) net cost of treatment (i.e., having a student attend a high-quality public university) may be a good fit for this environment because a public university likely has the education of the state's denizens at heart, especially if these students become more productive and stay in the state upon graduation (70% of applicants to the flagship eventually earn wages in the same state).[29]

Hoekstra uses a fuzzy design in which treatment was targeted to students at or above

[28]This relationship only needs to be approximately true in a neighborhood around κ for the argument made here. However, if this were instead thought to be a reasonable approximation to the *global* behavior of $Y_{\tau(x)}(x)$, and, therefore, $\Delta(x)$, then Appendix B shows that inclusion of an additive independent error ϵ does not affect the choice of κ^* or theoretical results. Some studies also use polynomial functions of the running variable, which affects how to estimate $\Delta'(\kappa^*)$ but, does not affect the test results for these applications.

[29]Epple et al. (2006) find that a model where universities optimize student achievement can explain the data. Because students' achievement measures their human capital, which itself augments wages, one could therefore view universities as wanting to maximize future wages. Similarly, maximizing students' completing college or finding (or keeping) jobs would naturally be captured by having the administrator maximize wages, as these schooling and labor market outcomes are all positively related to wages. That is, though it admittedly abstracts from alternative dimensions universities may care about, modeling public universities as maximizing student wages may reasonably approximate their objectives.

a covariate-adjusted SAT score, i.e., $\omega(x) = 1 \Leftrightarrow x \geq \kappa^*$. The intended treated students ($\omega(x) = 1$) were offered admission to the flagship and, for the most part, attended it. The intended untreated students ($\omega(x) = 0$) represent a combination of students who do not pursue any higher education, students who attend some other institution of higher education, and a small number of students who attend the flagship university, though the author provides evidence that most likely attend another institution. A nonconstant $\Delta(x) \equiv Y_1(x) - Y_0(x)$ then represents the oft-studied heterogeneity in the returns to education, with respect to student characteristics (x) and applied to the case of selective public universities (Y_1) versus less-competitive institutions (Y_0). The inferential problem with extrapolating from the RD estimate is that the gain may vary between students.

Research by Epple et al. (2006) shows that universities admit students until their capacity constraints bind. Such locally binding capacity constraints mean it is likely that Hoekstra (2009) was implemented under a Case-$\overline{a}$-binding capacity constraint.[30] The marginal cost of treatment is assumed to be constant and, as in the model, is denoted by χ. This assumption is supported by Izadi et al. (2002), who estimate a CES cost function for universities. Based on parameter estimates provided in that paper, one cannot reject that university cost functions are linear in the number of students served.[31] This assumption is also supported by other work, such as Epple et al. (2006), who estimate a model of the higher education market for private colleges and do not find evidence that the cost of serving students is nonlinear.[32]

Recall that the fuzzy design returns exactly the same bounds as does the sharp design, when ρ_ω are constant within treatment status. Therefore, the lower bound for the average effect of treatment on the treated ($\overline{\text{ATT}}$) and upper bound for the average effect of treatment on the untreated ($\overline{\text{ATUT}}$) developed earlier also apply here. The main result reported in Hoekstra (2009) is that attending the flagship university increases log wages by 20% for students at the treatment cutoff, relative to attending a less-competitive institution. This positive estimate means this context is consistent with (constrained-) optimality of the cutoff, implied by Condition $\overline{1a}$; that is, we cannot falsify the model assumptions. Note that when we use the bounds for the $\overline{\text{ATT}}$ when the capacity constraint binds, we can only surmise that the effect of treatment on the treated is positive, i.e., $\overline{\text{ATT}} > 0$.

[30] The method proposed earlier to distinguish between Case-$\overline{a}$- and Case-$\overline{b}$-binding capacity constraints—compare proposed, approved, and enacted budgets—unfortunately cannot be implemented because the identity of the university is not publicly available. Therefore, assuming a Case-$\overline{a}$-binding constraint is reasonable, as it likely corresponds to more conservative bounds.

[31] Specifically, I test whether the second derivative of the cost of serving arts and science students is zero in the number of that type of student, and find that even an 80% confidence interval for the second derivative contains zero for both student types. Izadi et al. (2002) use data from the UK.

[32] See Table II on page 907 of Epple et al. (2006).

As noted by Hoekstra, this estimate seems fairly high. For example, Ashenfelter and Rouse (1998) estimate that an additional year of schooling increases log wages by 9%, while Behrman et al. (1996) find that an additional year of schooling increases log wages by 6-8% and that there is a 20% increase from attending a large public college versus only graduating from high school. In contrast, Hoekstra (2009) estimates a 20% increase from attending a flagship, versus mostly a less-competitive, institution—that is, among students pursuing higher educations. However, the relatively large estimated effect in Hoekstra (2009) is quite intuitive when viewed through the lens of Condition $\overline{1a}$—that the RD estimate exceeds the marginal cost of treatment when the capacity constraint binds.

4.2 Lindo et al. (2010): "Ability, Gender, and Performance Standards: Evidence from Academic Probation"

This section applies this paper's results to Lindo et al. (2010), which studies how being placed on academic probation affects subsequent outcomes for university students. They exploit a sharp discontinuity design, where students with GPAs below a chosen cutoff are assigned to academic probation, i.e., $\tau(x) = 1 \Leftrightarrow x \leq \kappa^*$, where x is the student's GPA last semester. Students on academic probation must keep their GPAs above a certain standard, else they will be placed on academic suspension. The estimation sample comprises students from three campuses of a public university in Canada.

As with Hoekstra (2009), the fact that the university is public means it is reasonable to expect that it would value student achievement. Therefore, I focus on effect of being placed on probation on subsequent GPA, which means that the treatment effect $\Delta(x)$ is the expected gain in subsequent GPA if student with prior GPA x were placed on academic probation. Lindo et al. (2010) use a simplified version of Bénabou and Tirole (2000) to motivate why there may be heterogeneity in the effect of probation on student outcomes; the takeaway being that students far above the cutoff naturally perform well in their classes, and, therefore, would gain little from being put on probation. The university faces a cost of placing students on probation, which captures the fact that students are offered additional counseling and support services to help them improve their achievement. Therefore, assigning all students to probation would mean incurring costs for students who have little expected gain. Because only a subset of students are placed on probation and the effects of probation likely depend on student ability, it is useful to think about how we can extrapolate away from the treatment cutoff. The university could have treated more students by sending out more probation letters and hiring the counselor/tutor for more hours, which means it is reasonable to assume the capacity constraint was not binding in this application.

I begin by conducting the falsification test on the treatment effect derivative, implied by Condition 1(ii). The sign of $\Delta'(\kappa^*)$, and therefore the rejection region for the falsification test, is reversed here because treatment is offered to students below κ^*, meaning that extending treatment to students above κ^* should not improve the administrator's objective. Using the information made available by the journal's replication policy, I ran regression (9), the results of which are presented in Table 1. I find that $\widehat{\alpha}_3 = 0.047$, with a standard error of 0.094, which means that there is not strong evidence that the treatment effect is increasing at the cutoff (p-value 0.31 that the treatment effect derivative is greater than zero).[33] Moreover, the positive RD estimate (see below) means the model passes the nonnegative LATE falsification test, implied by Condition 1(i). That is, there is not enough evidence to reject that the model assumptions hold here. Therefore, assuming that the model assumptions indeed hold, the first result is that we can rule out constant treatment effects, by Corollary 1(i).

Table 1: Results from main specification in Lindo et al. (2010)

Regressor:	*Dependent variable:* GPA next semester ($\check{Y}_i$)
Intercept	0.312
($\widehat{\alpha}_0$)	(0.019)
Running variable - cutoff	0.699
($\widehat{\alpha}_1$)	(0.053)
Treatment indicator	0.233
($\widehat{\alpha}_2$)	(0.031)
Treatment indicator × (running variable - cutoff)	0.047
($\widehat{\alpha}_3$)	(0.094)
Obs.	11,258
R^2	0.035

Note: Standard errors are in parentheses

The main result of Lindo et al. (2010) is that the estimated treatment effect of being placed on academic probation on the next term's grade performance for the full sample is $\widehat{\Delta}(\kappa^*) = 0.233$ higher GPA points.[34] By extending this finding using the results of this paper, we can bound the ATT and ATUT according to: ATT $\geq \widehat{\Delta}(\kappa^*) = 0.233 >$ ATUT. In other words, placing students below the treatment cutoff on academic probation would, on average, increase their GPA the next term by at least 0.233 points, while doing so for students

[33] The data are available at `https://www.aea-net.org/articles.php?doi=10.1257/app.2.2.95`; R code for replication can be found at Chi and Dow (2014); R Core Team (2016).

[34] This can be found in Table 5, in column (1) of panel A of Lindo et al. (2010).

above the treatment cutoff would increase their GPA next term by at most 0.233 points, on average. Intuitively, on average, academic probation may be more useful for students at the bottom of the grade distribution, by providing them with an external commitment to increase their performance above some minimal level.[35]

Lindo et al. (2010) affords us an opportunity to explore how the upper bound on the average gain from treating all students with indices above a prospective cutoff $\hat{x}$. We must first adapt equation (6) to take into account that the treatment in this example is assigned to students below the cutoff, resulting in an average effect of treating "the rest of the untreated" beyond $\hat{x}$, i.e., $\text{ATUT}(\hat{x}) \equiv (\int_{\hat{x}}^{1} \Delta(x)dx)/(1-\hat{x}) \leq [\chi(1-\kappa^*) - \underline{\Delta}(\hat{x}-\kappa^*)]/(1-\hat{x})$. Note this bound increases as we increase the prospective cutoff $\hat{x}$, i.e., decrease the size of the "rest of the untreated" group.

Thus, we need values for $(\chi, \kappa^*, \underline{\Delta})$ to solve for the upper bound on $\text{ATUT}(\hat{x})$. By Condition 1(i) we can use the estimated treatment effect at the cutoff, 0.233 GPA points, as an estimate of the (constant) marginal cost of treatment. Next, by noting that 25% of students received the treatment (Lindo et al. (2010), page 101) we can set $\kappa^* = 0.25$. Finally, we need to obtain the worst-case scenario from being assigned the treatment, which is mostly composed of a letter and some counseling and tutoring services. One option would be to assign a null effect, i.e., set $\underline{\Delta} = 0$. However, we could be more conservative by taking into account the opportunity cost of students' time, if we had an idea of how studying affected GPA and conservatively assumed that participating in the extra services (i) resulted in a complete crowd-out of study time and (ii) did not ceteris paribus increase grades. Assuming students on probation had one hour of time taken by these extra services per week, we can use the estimate from Stinebrickner and Stinebrickner (2008), that an hour of studying per day increases one's college GPA by 0.36 points, to roughly figure that the time-cost of this hourly meeting would be $\underline{\Delta} = -0.36/7 = -0.051$ GPA points.

Figure 3 illustrates the results of the above calculations. We can see that the upper bound on $\text{ATUT}(\hat{x})$ starts at the upper bound on the ATUT—i.e., $\Delta(\kappa^*)$—when $\hat{x} = \kappa^*$, at the left side (blue dotted line) and then increases as the administrator increases the prospective cutoff above which all students would receive treatment. For example, the mean GPA increase from treating the top 70% of students in terms of prior GPA (i.e., $\hat{x} = 0.3$) would be no larger than 0.253 and the mean increase from treating the top half of students (i.e., $\hat{x} = 0.5$) would be no larger than 0.375 GPA points. Although these bounds increase in $\hat{x}$, they may be small enough to be of potential use to policymakers.

[35] This is a conjecture. Proposition 2 shows that we can obtain bounds for subsets of untreated students. However, no such bounds can be obtained for subsets of treated students, such as those with the lowest GPAs.

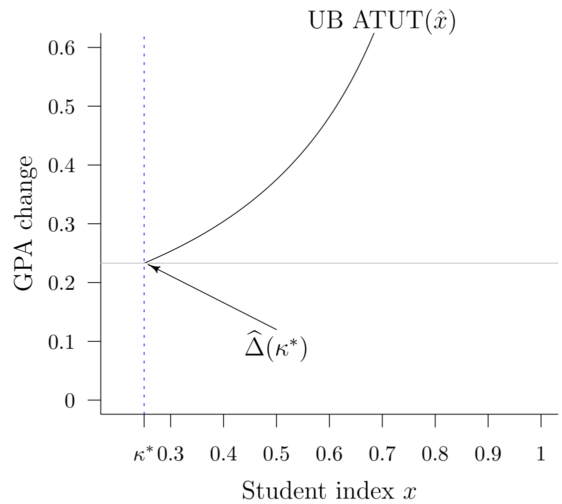

Figure 3: Upper Bound on Treating the "Rest of the Untreated"in Lindo et al. (2010)

5 Discussion

This paper represents a first step towards showing how one can use plausibly available information and a simple economic model to generalize findings from RD designs.[36] We can exploit information revealed by the optimizing behavior of the administrator to extrapolate from the LATE at the treatment cutoff, which is often available but can apply only to measure zero of the population, to obtain bounds for the treated, untreated, and the entire population. Perhaps the most intuitive findings relate to the case where the capacity constraint does not bind: i) if treating students is costly and the treatment cutoff has been chosen optimally, the ATT must be positive and treatment effects cannot be constant; ii) RD-based estimates provide a lower bound, or understate, the ATT; and iii) RD-based estimates provide an upper bound for the ATUT. Notably for applying these results, the model generates testable implications: if the treatment effect at the cutoff is negative or treatment effects are decreasing in the direction of treatment at the cutoff, then we can reject that the cutoff was chosen optimally by an administrator informed about the treatment effects. If the capacity constraint does bind then the treatment-effect sign test still allows one to falsify the model and bounds are generally looser. The treatment-effect-derivative test no longer applies, but there emerges a new testable implication of cutoff optimality, as well as an intuitive explanation for why program "scale-up" can be difficult in real-life applications. The theoretical results were then demonstrated using two applications. We cannot

[36]This paper also relates to work non-parametrically estimating treatment effects.

reject that the cutoff was chosen optimally by an informed administrator in every one of the falsification tests conducted, which may increase confidence that the contexts studied here constitute reasonable applications of the theoretical results.

The findings in this paper have several implications for the use of RD results in policy. Perhaps most novel is that we may incorrectly surmise that some programs are ineffective and eliminate them, even though in reality they are quite effective for the treated population. Strikingly, *such a mistake would be more likely for a program with a very low marginal cost, holding constant the ATT.* This point is illustrated in Figure 4, which plots treatment effects associated with hypothetical programs at two sites, A and B, with respective conditional treatment effect functions $\Delta_A(\cdot)$ (solid black line) and $\Delta_B(\cdot)$ (dashed black line). The programs have different marginal costs of treatment, where $\chi_B > \chi_A$, and happen to have the same cutoff κ^* and the same ATT. The difference in marginal costs means that optimization by respective site administrators implies that $\Delta_B(\kappa^*) > \Delta_A(\kappa^*)$. If only based on these RD estimates, a policymaker would likely fund B over A because it has a higher LATE, even though A provides the same gain on the treated, *at a lower cost.* In a sense, this paper provides an illustration of the importance of taking into account the costs, not just the benefits, of treatment. Additional policy-relevant results obtain if we can relate the policymaker's objective with that of the administrator. First, if a policymaker knew that their valuation of treatment gain in terms of treatment cost (i.e., β) was at least as high as the administrator's then he should definitely treat those units treated by the administrator. Second, the upper bound on the ATUT for subsets of the untreated, which increases in distance from the cutoff, can help rule out whether it would be worthwhile to extend treatment to subsets of units below the cutoff.

Though, in this setting, estimates of the LATE at the treatment cutoff must be positive if treating students is costly, we cannot compare them with the ATE, in the manner of LaLonde (1986), Dehejia and Wahba (1999), or Smith and Todd (2005), without further information. There is some work comparing findings from RD and experimental designs (Buddelmeyer and Skoufias (2004), Black et al. (2007), Cook and Wong (2008), Gleason et al. (2012), Barrera-Osorio et al. (2014)), but unfortunately, none consider the case of a program where the treatment cutoff seems to have been chosen by an administrator with institutional knowledge of the environment (e.g., in Barrera-Osorio et al. (2014) the evaluators were external and choose a poverty index as the threshold for treatment). However, the results here do suggest that RD estimates may be higher when cutoffs are chosen by external evaluators without institutional knowledge, hence less information about treatment effects. Related to this point, an unconstrained optimizing administrator would not choose to place the cutoff where they know the gain from treatment is quite large. Because RD estimates may understate the

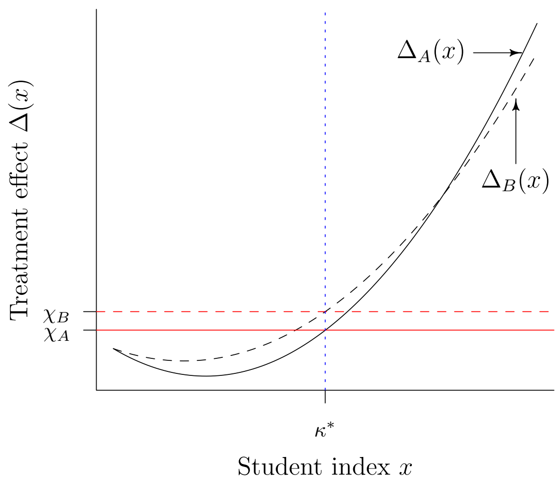

Figure 4: Example of when policy based only on the RD estimate could cause termination of the wrong program

ATT, there may be RD studies of useful programs that are simply not published because they lack statistically significant findings.

One practical variation of the environment considered here would introduce a more substantial form of uncertainty, for example, featuring learning about treatment effects, into the administrator's problem. Such uncertainty would pervade to the bounds obtained here, perhaps motivating a Bayesian approach. A more formal approach could also combine bounds for a particular treatment that had been implemented across multiple sites, to build up a picture of the population-level (as opposed to site-specific) heterogeneity in treatment effects. Another variation would investigate what could be learned if the administrator only knew certain features of treatment effects, say, the ATT. Such variations could be worthwhile ways to build on the basic point made in this paper: revealed preferences can provide quite a bit of useful information about treatment effects away from the cutoff in regression-discontinuity designs.

References

Angrist, J. D. and M. Rokkanen, "Wanna Get Away? Regression Discontinuity Estimation of Exam School Effects Away from the Cutoff," *Journal of the American Statistical Association*, 110(512):1331–1344, 2015.

Anwar, S. and H. Fang, "An Alternative Test of Racial Prejudice in Motor Vehicle Searches: Theory and Evidence," *American Economic Review*, 96(1):127–151, 2006.

Ashenfelter, O. C. and C. E. Rouse, "Income, Schooling, and Ability: Evidence From a New Sample of Identical Twins," *Quarterly Journal of Economics*, 113(1):253–284, 1998.

Barrera-Osorio, F., D. Filmer and J. McIntyre, "Randomized Controlled Trials and Regression-Discontinuity Estimations: An Empirical Comparison," 2014.

Behrman, J. R., M. R. Rosenzweig and P. Taubman, "College Choice and Wages: Estimates Using Data On Female Twins," *The Review of Economics and Statistics*, pp. 672–685, 1996.

Bénabou, R. and J. Tirole, "Self-Confidence and Social Interactions," *NBER Working Paper*, 2000.

Björklund, A. and R. Moffitt, "The Estimation of Wage Gains and Welfare Gains in Self-Selection Models," *The Review of Economics and Statistics*, pp. 42–49, 1987.

Black, D. A., J. Galdo and J. A. Smith, "Evaluating the Worker Profiling and Reemployment Services System Using a Regression Discontinuity Approach," *The American Economic Review*, pp. 104–107, 2007.

Brock, W., J. Cooley, S. Durlauf and S. Navarro, "On the Observational Implications of Taste-Based Discrimination in Racial Profiling," *Journal of Econometrics*, 2011.

Buddelmeyer, H. and E. Skoufias, *An Evaluation of the Performance of Regression Discontinuity Design On PROGRESA*, vol. 827, World Bank Publications, 2004.

Chi, O. and A. Dow, "Replication Data For: Improving Balance in Regression Discontinuity Through Matching: Estimating the Effect of Academic Probation After the First Year of College," Tech. rep., Harvard Dataverse, 2014.

Cook, T. D. and V. C. Wong, "Empirical Tests of the Validity of the Regression Discontinuity Design," *Annales d'Economie et de Statistique*, pp. 127–150, 2008.

Dehejia, R. H. and S. Wahba, "Causal Effects in Nonexperimental Studies: Reevaluating the Evaluation of Training Programs," *Journal of the American statistical Association*, 94(448):1053–1062, 1999.

DiNardo, J. and D. S. Lee, "Program Evaluation and Research Designs," *Handbook of Labor Economics*, 4:463–536, 2011.

Dong, Y. and A. Lewbel, "Identifying the Effect of Changing the Policy Threshold in Regression Discontinuity Models," *Review of Economics and Statistics*, 97(5):1081–1092, 2015.

Elmore, R., "Getting to Scale with Good Educational Practice," *Harvard Educational Review*, 66(1):1–27, 1996.

Epple, D., R. Romano and H. Sieg, "Admission, Tuition, and Financial Aid Policies in the Market for Higher Education," *Econometrica*, 74(4):885–928, 2006.

Ferrall, C. and B. Shearer, "Incentives and Transactions Costs Within the Firm: Estimating an Agency Model Using Payroll Records," *Review of Economic Studies*, 66(2):309–338, 1999.

Gleason, P. M., A. M. Resch and J. A. Berk, "Replicating Experimental Impact Estimates Using a Regression-Discontinuity Approach," Tech. rep., Mathematica Policy Research, 2012.

Hahn, J., P. Todd and W. Van der Klaauw, "Identification and Estimation of Treatment Effects with a Regression-Discontinuity Design," *Econometrica*, 69(1):201–209, 2001.

Heckman, J., J. Smith and N. Clements, "Making the Most Out of Programme Evaluations and Social Experiments: Accounting for Heterogeneity in Programme Impacts," *The Review of Economic Studies*, 64(4):487–535, 1997.

Heckman, J. J., R. J. LaLonde and J. A. Smith, "The Economics and Econometrics of Active Labor Market Programs," *Handbook of Labor Economics*, 3:1865–2097, 1999.

Heckman, J. J. and J. A. Smith, "Evaluating the Welfare State," in "Econometrics and Economic Theory in the 20th Century: The Ragnar Frisch Centennial Symposium," 31, p. 241, Cambridge University Press, 1998.

Heckman, J. J. and S. Urzua, "Comparing IV with Structural Models: What Simple IV Can and Cannot Identify," *Journal of Econometrics*, 156(1):27–37, 2010.

Heckman, J. J. and E. J. Vytlacil, "Econometric Evaluation of Social Programs, Part II: Using the Marginal Treatment Effect to Organize Alternative Econometric Estimators to Evaluate Social Programs, and to Forecast Their Effects in New Environments," *Handbook of Econometrics*, 6:4875–5143, 2007.

Heinrich, C., J. J. Heckman and J. Smith, "The Performance of Performance Standards," *Journal of Human Resources*, 37(4):778–811, 2002.

Hoekstra, M., "The Effect of Attending the Flagship State University On Earnings: A Discontinuity-Based Approach," *The Review of Economics and Statistics*, 91(4):717–724, 2009.

Imbens, G. W., "Better LATE Than Nothing: Some Comments On Deaton (2009) and Heckman and Urzua (2009)," *Journal of Economic Literature*, pp. 399–423, 2010.

Imbens, G. W. and T. Lemieux, "Regression Discontinuity Designs: A Guide to Practice," *Journal of Econometrics*, 142(2):615–635, 2008.

Izadi, H., G. Johnes, R. Oskrochi and R. Crouchley, "Stochastic Frontier Estimation of a CES Cost Function: The Case of Higher Education in Britain," *Economics of Education Review*, 21(1):63–71, 2002.

Knowles, J., N. Persico and P. Todd, "Racial Bias in Motor Vehicle Searches: Theory and Evidence," *Journal of Political Economy*, 109(1), 2001.

LaLonde, R. J., "Evaluating the Econometric Evaluations of Training Programs with Experimental Data," *The American Economic Review*, pp. 604–620, 1986.

Lee, D. S. and D. Card, "Regression Discontinuity Inference with Specification Error," *Journal of Econometrics*, 142(2):655–674, 2008.

Lindo, J. M., N. J. Sanders and P. Oreopoulos, "Ability, Gender, and Performance Standards: Evidence From Academic Probation," *American Economic Journal: Applied Economics*, 2(2):95–117, 2010.

Manski, C. F., *Partial Identification of Probability Distributions*, Springer Science & Business Media, 2003.

—, "Statistical Treatment Rules for Heterogeneous Populations," *Econometrica*, 72(4):1221–1246, 2004.

—, "Choosing Treatment Policies Under Ambiguity," *Annual Review of Economics*, 3(1):25–49, 2011.

Manski, C. F. and J. V. Pepper, "Monotone Instrumental Variables: With an Application to the Returns to Schooling," *Econometrica*, 68(4):997–1010, 2000.

McCrary, J., "Manipulation of the Running Variable in the Regression Discontinuity Design: A Density Test," *Journal of Econometrics*, 142(2):698–714, 2008.

R Core Team, *R: A Language and Environment for Statistical Computing*, R Foundation for Statistical Computing, Vienna, Austria, 2016.

Roy, A., "Some Thoughts On the Distribution of Earnings," *Oxford Economic Papers*, 3(2):135, 1951.

Rubin, D. B., "Comment," *Journal of the American Statistical Association*, 75(371):591–593, 1980.

Smith, J. A. and P. E. Todd, "Does Matching Overcome LaLonde's Critique of Nonexperimental Estimators?" *Journal of Econometrics*, 125(1):305–353, 2005.

Sternberg, R. J., D. Birney, L. Jarvin, A. Kirlik, S. Stemler and E. L. Grigorenko, "Scaling Up Educational Interventions," *Translating educational theory and research into practice: Developments in content domains, large scale reform, and intellectual capacity*, 2006.

Stinebrickner, R. and T. R. Stinebrickner, "The Causal Effect of Studying On Academic Performance," *The BE Journal of Economic Analysis & Policy*, 8(1):1–55, 2008.

Van der Klaauw, W., "Regression–Discontinuity Analysis: A Survey of Recent Developments in Economics," *Labour*, 22(2):219–245, 2008.

Appendix

A Allow the Probability of Enrollment to Vary by x

Consider (8), where ρ_ω are no longer assumed to be uniform among treated students, but instead depends on x. The administrator's problem becomes:

$$\max_{\widetilde{\kappa}} \left(\int_0^{\widetilde{\kappa}} \rho_0(x)(\Delta(x) - \chi)dx \right) + \left(\int_{\widetilde{\kappa}}^1 \rho_1(x)(\Delta(x) - \chi)dx \right).$$

Note that a necessary condition for κ^* being optimal is still $(\rho_1(\kappa^*) - \rho_0(\kappa^*))\Delta(\kappa^*) = (\rho_1(\kappa^*) - \rho_0(\kappa^*))\chi \Rightarrow \Delta(\kappa^*) = \chi$, which is identical to Condition 1(i) for the sharp design. Optimality of κ^* further implies:

$$\int_{\widetilde{\kappa}}^1 \rho_1(x)\Delta(x)dx \geq \int_{\widetilde{\kappa}}^1 \rho_1(x)\chi dx = \overline{\rho_1}\chi(1-\kappa^*) \Leftrightarrow$$

$$\underbrace{\frac{\int_{\widetilde{\kappa}}^1 \rho_1(x)\Delta(x)dx}{1-\kappa^*}}_{\text{ITT}} \geq \underbrace{\overline{\rho_1}\chi}_{\text{expected cost of treating the treated}} = \overline{\rho_1}\widehat{\Delta}(\kappa^*), \tag{10}$$

where the second equality follows because $\Delta(\kappa^*) = \chi$. Equation (10) shows that if the average attendance probability among the treated ($\overline{\rho}_1$) were known then the RD estimate of the treatment effect can again be used to provide a lower bound for the mean effect of intending-to-treat among the treated (ITT).

B Treatment Effect Uncertainty

Suppose the administrator is uncertain about the treatment effect but has observed $\check{\Delta}(x)$, an unbiased signal of $\Delta(x)$. Let $\check{\Delta}(x) = \Delta(x) + \epsilon_i$, where ϵ is distributed independently from x, denote the administrator's noisy signal of the treatment effect for student i who has index x. Because the administrator has unbiased beliefs about $\Delta(x)$ at every point x, it must be the case that $\mathrm{E}[\epsilon_i] = 0$. The administrator chooses a cutoff to maximize her expected objective:

$$\max_{\widetilde{\kappa}} \mathrm{E}\left[\beta\left(\int_{\widetilde{\kappa}}^{1} \check{\Delta}(x)dx\right) - c\,(1-\widetilde{\kappa})\right] \Leftrightarrow \max_{\widetilde{\kappa}} \beta\,\mathrm{E}\left[\left(\int_{\widetilde{\kappa}}^{1} (\Delta(x)+\epsilon)dx\right)\right] - c\,(1-\widetilde{\kappa})$$
$$\Leftrightarrow \max_{\widetilde{\kappa}} \beta\left(\int_{\widetilde{\kappa}}^{1} \Delta(x)dx\right) + \beta\,\mathrm{E}[\epsilon] - c\,(1-\widetilde{\kappa})$$
$$\Leftrightarrow \max_{\widetilde{\kappa}} \beta\left(\int_{\widetilde{\kappa}}^{1} \Delta(x)dx\right) - c\,(1-\widetilde{\kappa})\,. \tag{11}$$

The first equivalence follows from the fact that the measure of students treated $(1-\widetilde{\kappa})$ is known because it is chosen by the administrator. The second follows from the independence assumption and the third from unbiasedness. The last expression is the administrator's original problem, (1). Therefore, the analysis for this case is identical. Intuitively, uncertainty does not affect the administrator's problem because it is linear in the amount gained.

We can also use this setup to examine what would happen if the administrator instead only had access to a biased measure of $\Delta(\cdot)$. Define $\delta(x) \equiv \mathrm{E}[\epsilon_i|x]$, i.e., the conditional expectation of ϵ given x. In the case of an unbiased $\check{\Delta}(x)$ we have $\delta(x) = 0$ for all x. I consider two types of biased beliefs.

Constant Bias First suppose $\delta(x) = \delta \neq 0$, i.e., ϵ is biased, but mean independent of x. In this case, Condition 1 would not be affected, as the optimal cutoff κ^* would not change from the unbiased case. Intuitively, if $\delta(\cdot)$ does not depend on x the bias does not affect the administrator's objective at the intensive margin.

Condition 2, however would be affected. Consider first the augmented participation condition:

$$\int_{\kappa^*}^{1} \Delta(x)dx \geq (\chi - \delta)(1-\kappa^*),$$

which implies the ATT lower bound would be shifted downwards by the constant amount δ. Similarly, the non-extension condition would become

$$\int_{\hat{\kappa}}^{\kappa^*} \Delta(x)dx < \chi(\kappa^* - \hat{\kappa}) = (\chi - \delta)(1-\kappa^*),$$

i.e., the upper bound on the ATUT would also be shifted down by the constant δ. These changes to Condition 2 would propagate to the other bounds results.

Differential Bias in x Now let $\delta(x)$ be variable in x. The augmented participation condition becomes

$$\int_{\kappa^*}^{1} \Delta(x)dx \geq \chi(1-\kappa^*) - \int_{\kappa^*}^{1} \delta(x)dx$$

and the augmented non-extension condition becomes

$$\int_{\hat{\kappa}}^{\kappa^*} \Delta(x)dx < \chi(\kappa^* - \hat{\kappa}) - \int_{\hat{\kappa}}^{\kappa^*} \delta(x)dx.$$

Consider the following two cases: (i) $\mathrm{E}\left[\delta(x)|x \in [\hat{\kappa}, \kappa^*)\right] < 0 < \mathrm{E}\left[\delta(x)|x \geq \kappa^*\right]$, $\forall \hat{\kappa} \in [0, \kappa^*)$ and (ii) $\mathrm{E}\left[\delta(x)|x \geq \kappa^*\right] < 0 < \mathrm{E}\left[\delta(x)|x \in [\hat{\kappa}, \kappa^*)\right]$, $\forall \hat{\kappa} \in [0, \kappa^*)$. In case (i) the augmented participation and non-extension conditions would reduce the lower bound on the ATT and increase the upper bound on the ATUT. That is, all bounds would be looser. In case (ii) the opposite would happen, i.e., bounds would tighten.

C Weighted Objective

The administrator's original problem (1) was utilitarian, i.e., it weighed gains for all students equally. The most natural alternative to the unweighted objective would be a redistributive policy, which assigned people with lower running variable indices larger weights. For example, if x measured incoming human capital, then putting more weight on gains for students with lower indices allows the administrator to place additional value on students' becoming proficient. In this case, we can adapt equation (1) to allow the administrator to weigh gains for students depending on their index x by using weights $\phi(x)$, where $\phi' \leq 0$:

$$\max_{\widetilde{\kappa}} \left(\int_{\widetilde{\kappa}}^{1} \phi(x)\Delta(x)dx \right) - \chi\left(1 - \widetilde{\kappa}\right), \qquad (\hat{1})$$

and proceed with the analysis.

Condition $\hat{1}$ (Necessity)**.** *For problem ($\hat{1}$), the following necessary conditions must hold for κ^*:*

(i) MB=MC: $\phi(\kappa^*)\Delta(\kappa^*) = \chi$

(ii) Increasing MB: $\Delta'(\kappa^*) \geq 0$.

Proof. Differentiate the administrator's problem ($\hat{1}$) with respect to $\tilde{\kappa}$ to obtain (i). Note that if the derivative is negative at a candidate solution satisfying (i), the administrator would gain by not treating students just above κ^*, thereby obtaining (ii). The inequality is strict if $\phi' < 0$. □

Condition $\hat{2}$ (Sufficiency)**.** *The fact the program was implemented implies that the total gain from treating those units was at least as large as the total costs, i.e.:*

$$\textit{Participation: } \int_{\kappa^*}^{1} \phi(x)\Delta(x)dx \geq \chi(1-\kappa^*). \tag{$\hat{2}$}$$

The fact the program was not extended to $\hat{\kappa} < \kappa^*$ implies that treating these units would be sub-optimal, i.e.:

$$\int_{\hat{\kappa}}^{\kappa^*} \phi(x)\Delta(x)dx < \chi(\kappa^* - \hat{\kappa}). \tag{$\hat{3}$}$$

Proposition 1 remains true when $\phi' \leq 0$. To see this, divide ($\hat{2}$) by the measure of treated students and combine with Condition $\hat{1}$(i) to obtain $\left(\int_{\kappa^*}^{1} \phi(x)\Delta(x)dx\right) / (1-\kappa^*) \geq \phi(\kappa^*)\Delta(\kappa^*)$. Because $\phi' \leq 0$, this implies that $\left(\int_{\kappa^*}^{1} \Delta(x)dx\right) / (1-\kappa^*) \geq \Delta(\kappa^*)$, where the inequality is strict if $\phi' < 0$. Intuitively, the gains for treating the treated must be even larger than the LATE if the administrator values such gains less. Analogous reasoning applied to Corollary 3 shows that the ATUT is bounded above by the LATE when $\phi' \leq 0$, and that this bound is strict when $\phi' < 0$. Therefore, the corollaries, in particular Corollary 4 bounding the ATE, also still obtain with the weighted problem ($\hat{1}$). In summary, all of the bounds from the unweighted problem, including Corollary 4, which bounds the ATE, are also obtained for the weighted problem ($\hat{1}$).

D Variable Marginal Cost of Treatment

Begin by relaxing Assumption 1(i), replacing it with

Assumption 1′. *(i) The cost function $c(\cdot)$ is known and is non-negative, strictly increasing, and differentiable. The marginal cost function $c'(\cdot)$ is monotonic.*

Note that Assumption 1′(i) still implies that the marginal cost of providing treatment is strictly positive. The second part of Assumption 1′(i) relaxes the constant marginal cost assumption. Note that the cost can be variable in μ, but does not vary stochastically or *directly* with respect to x. However, it is possible to *indirectly* to pick up variation in costs with respect to x by using a reduced-form cost function $c_{\text{rf}}(\mu)$ in place of $c(\mu)$. Suppose the marginal cost was composed of two components and also took as an argument x: $c'_{\text{both}}(\mu, x) = c'(\mu) + c_x(x)$, where $c'(\mu)$ represented the marginal cost of the cost function in Assumption $1'$ and $c_x(\cdot)$ was monotonic in x. For example, suppose the first component was constant, i.e., $c'(\mu) = \chi$. Then, if $c_x(\cdot)$ is a constant χ_x the reduced-form marginal cost function $c'_{\text{rf}}(\mu) = \chi + \chi_x$ would also be constant. However, if $c_x(\cdot)$ is strictly increasing (decreasing) in x then the reduced-form total cost function would be $c_{\text{rf}}(\mu) = \int_{1-\mu}^{1} (c'(x) + c_x(x))\, dx$, which depends on the order in which students are treated. Then, the reduced form, $c'_{\text{rf}}(\mu)$, would be strictly decreasing (increasing), because students are added by extending the cutoff downward from 1. Indeed, an increasing $c_x(\cdot)$ could potentially transform an increasing marginal cost to a constant or even decreasing reduced-form marginal cost function, which would then be the one used in the analysis.

I first adapt the conditions characterizing κ^*, in terms of $\Delta(\cdot)$ and qualitative features of the (potentially reduced-form) cost function $c(\cdot)$. Specifically, I consider three cases for Assumption $1'$(i): where the marginal cost is constant, decreasing, and increasing; these correspond to linear, concave, and convex cost functions, respectively. I then provide results bounding treatment effects of interest.

Condition $1'$ (Necessity)**.** *The following necessary conditions must hold for κ^*:*

(i) MB=MC: $\Delta(\kappa^*) = c'(1-\kappa^*)$ for any cost function $c(\cdot)$ satisfying Assumption $1'$

(ii) Increasing MB: $\Delta'(\kappa^*) \geq 0$ if the marginal cost is constant or decreasing; this inequality is strict if the marginal cost is decreasing.

Proof. Differentiate the administrator's problem (1) with respect to $\tilde{\kappa}$ to obtain (i). Note that if the derivative is negative at a candidate solution satisfying (i) but the marginal cost is nonincreasing, the administrator would gain by not treating students just above κ^*, thereby obtaining (ii). □

Condition $1'$ is similar to Condition 1, except that Condition $1'$(ii) has a strict inequality if the marginal cost of treatment is decreasing. As before, to guarantee uniqueness, inspection of (1) implies two additional conditions sufficient for characterizing κ^*. These conditions are identical to those in Condition 2, the only difference being that χ no longer enters either expression.

Condition $2'$ (Sufficiency)**.** *The fact the program was implemented implies that the total gain from treating those units was at least as large as the total costs, i.e.:*

$$\textit{Participation:} \quad \int_{\kappa^*}^{1} \Delta(x)dx \geq c(1-\kappa^*). \tag{$2'$}$$

The fact the program was not extended to $\hat{\kappa} < \kappa^*$ implies that treating these units would be sub-optimal, i.e.:

$$\int_{\hat{\kappa}}^{\kappa^*} \Delta(x)dx < c(1-\hat{\kappa}) - c(1-\kappa^*). \tag{3'}$$

As before, a corollary immediately follows.

Corollary 1′. *The following are globally true about $\Delta(\cdot)$ if the marginal cost of treatment is nonincreasing:*

(i) $\Delta(\cdot)$ *cannot be constant.*

(ii) $\Delta(\cdot)$ *is not globally monotonically decreasing in x.*

Proof. Identical to proof of Corollary 1. □

As before, I next examine what can be deduced about averages of treatment effects for subsets of students.

Corollary 2′. *The ATT is positive for any cost function $c(\cdot)$ satisfying Assumption 1′.*

Proof. The left side of (2′) in Condition 2′ is the total effect of treatment on the treated, i.e., $(\int_{\kappa^*}^{1} \frac{\Delta(x)}{(1-\kappa^*)}dx)\,(1-\kappa^*)$. Because the marginal cost of treatment is positive (Assumption 1′(i)), (2′) implies that

$$\int_{\kappa^*}^{1} \Delta(x)dx \geq c(1-\kappa^*) > 0.$$

Divide through by $(1-\kappa^*)$ to obtain the result:

$$\underbrace{\int_{\kappa^*}^{1} \frac{\Delta(x)}{(1-\kappa^*)}dx}_{\text{ATT}} \geq \underbrace{\frac{c(1-\kappa^*)}{(1-\kappa^*)}}_{\text{avg. cost of treating treated}} > 0.$$ □

Although Corollary 2′ provides a lower bound for the average effect of treatment on the treated, there is no informative (i.e., lower than $\overline{\Delta}$) upper bound. Corollary 2′ makes no further assumptions about the shape of the cost function. However, if the marginal cost of treating students is nonincreasing, the lower bound on the average effect of treatment on the treated increases.

Proposition 1′. *If the marginal cost of treatment is nonincreasing, the ATT is bounded below by the LATE at the treatment cutoff.*

Proof. If the marginal cost of treatment is nonincreasing then $c'(1-\kappa^*) \leq \frac{c(1-\kappa^*)}{1-\kappa^*}$, i.e., the marginal cost of treatment for $1-\kappa^*$ is no greater than the average cost of providing treatment for treated students. Insert this inequality into (2′) and combine with this with Condition 1′(i) to obtain

$$\underbrace{\frac{\int_{\kappa^*}^{1} \Delta(x)dx}{1-\kappa^*}}_{\text{ATT}} \geq \frac{c(1-\kappa^*)}{1-\kappa^*} \geq c'(1-\kappa^*) = \underbrace{\Delta(\kappa^*)}_{\text{LATE at } \kappa^*} . \qquad \square$$

As with Proposition 1, Proposition 1′ shows that, if the marginal cost of treatment is nonincreasing, the discontinuity-based estimate provides a lower bound for the average effect of treatment on the treated. One should note that only qualitative information about the shape, not the level, of the marginal cost of treatment is all that is required for this result.

Although Corollary 1′(ii) rules out a treatment effect that is decreasing everywhere (if the marginal cost of treatment is nonincreasing), it could be the case that $\Delta(\cdot)$ is decreasing for some $x < \kappa^*$.[37] Therefore, as before, it is useful to bound averages of $\Delta(\cdot)$ itself for strict subsets of untreated students.

Proposition 2′. *There exists an informative upper bound for* $\int_a^b \Delta(x)dx$ *for* $0 \leq a < b \leq \kappa^*$.

Proof. Suppose we would like to characterize $\Delta(\cdot)$ for values less than $\hat{x} < \kappa^*$. Let $\hat{\mu}$ be the measure of students under consideration and split (3′) into two parts at $\hat{x}$ and rearrange terms:

$$\int_{\hat{x}-\hat{\mu}}^{\hat{x}} \Delta(x)dx < c(1-(\hat{x}-\hat{\mu}))-c(1-\kappa^*)-\int_{\hat{x}}^{\kappa^*} \Delta(x)dx \Rightarrow \int_{\hat{x}-\hat{\mu}}^{\hat{x}} \Delta(x)dx < c(1-(\hat{x}-\hat{\mu}))-c(1-\kappa^*)-\underline{\Delta}\left(\kappa^*-\hat{x}\right), \qquad (4')$$

where the implication follows from Assumption 1(iii). $\square$

Setting the measure of students to whom the treatment is extended equal to κ^* provides the following result about the ATUT.

Corollary 3′. *The ATUT has an informative upper bound. If the marginal cost of treatment is nonincreasing, this upper bound is the LATE at the treatment cutoff.*

[37]Note that Corollary 1′(ii) would also obtain when the marginal cost of treatment was strictly increasing, if the administrator could choose which side of the cutoff to treat.

Proof. Let $\hat{x} = \hat{\mu} = \kappa^*$ in (4′) and divide through by κ^* to obtain the first result:

$$\underbrace{\int_0^{\kappa^*} \frac{\Delta(x)}{\kappa^*} dx}_{\text{ATUT}} < \underbrace{\frac{c(1) - c(1-\kappa^*)}{\kappa^*}}_{>0,<\infty}, \tag{7′}$$

where the right hand side is positive from Assumption 1′(i). For the second result, note that a nonincreasing marginal cost implies

$$\frac{c(1) - c(1-\kappa^*)}{\kappa^*} < c'(1-\kappa^*) = \underbrace{\Delta(\kappa^*)}_{\text{LATE at } \kappa^*},$$

where the equality follows from Proposition 1′(i). □

Analogously to the upper bound for the ATT, although Corollary 3′ bounds the average of treatment effects for all untreated students, there is no informative (i.e., greater than $\underline{\Delta}$) lower bound.

To summarize, optimality of κ^* implies a lower bound on the ATT and an upper bound on the ATUT. If the marginal cost of treatment is constant or decreasing then it must be the case that ATUT $< \Delta(\kappa^*) \leq$ ATT. Though the ATT and ATUT are respectively bounded below and above by the cutoff LATE when marginal costs are nonincreasing, the LATE does not bound these moments when the marginal cost of treatment is strictly increasing.

D.1 Bounding the ATE

This section studies the interplay between qualitative features of the cost of treatment and inferences about treatment effects, by comparing three cases: constant, decreasing, and increasing marginal cost of treatment, where each marginal cost curve passes through the point $(\kappa^*, \Delta(\kappa^*))$. A decreasing marginal cost ($c'' < 0$) might result from economies of scale, while an increasing marginal cost ($c'' > 0$) might result from congestion effects, say if it becomes increasingly difficult to find a good fit for the program.

To begin, suppose the cost function is $c(\mu) = \mu\chi$. Then, as was shown in Section 3.1, the ATE lower bound is $\Delta^{LB} \equiv \underline{\Delta}\kappa^* + \chi(1-\kappa^*)$ and the ATE upper bound is $\Delta^{UB} \equiv \chi\kappa^* + \overline{\Delta}(1-\kappa^*)$, because $\chi = \Delta(\kappa^*)$ by Condition 1(i). Figure 5 builds on the example in Figure 1 to provide intuition for how the marginal cost of treatment bounds the ATE. Start with the solid red line representing a constant marginal cost of treatment, and rotate the cost function counterclockwise about the point $(\kappa^*, \Delta(\kappa^*))$ to represent a decreasing marginal cost of treatment (long-dashed red line).[38] This rotation implies the ATT must be higher than the case corresponding to the constant marginal cost in order to satisfy (2′). Analogously, the maximum ATUT must be lower when marginal costs are decreasing; were they the same as with constant marginal costs, the administrator might gain from extending treatment to untreated units given that they now have a lower cost of being treated, violating (3′). The opposite holds true for when we rotate the cost curve clockwise about the point $(\kappa^*, \Delta(\kappa^*))$, to reflect an increasing marginal cost of treatment (dot-dashed red line). Table 2 summarizes these results, showing that when the marginal cost of treatment is decreasing, bounds on the ATE are tighter than they would be with a constant marginal cost, while when marginal cost is increasing, bounds on the ATE are looser.

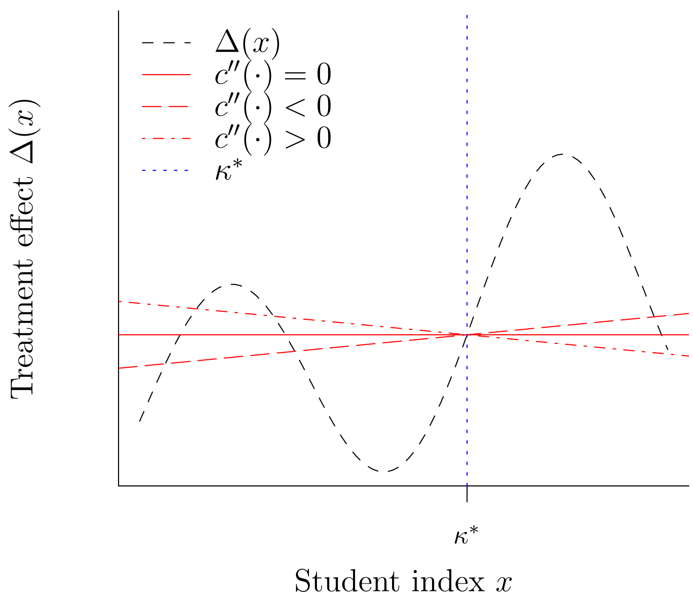

Figure 5: Example with different cost functions

	ATE bounds	
Marginal cost	Lower	Upper
Const. ($c'' = 0$)	$=\Delta^{LB}$	$= \Delta^{UB}$
Dec. ($c'' < 0$)	$> \Delta^{LB}$	$< \Delta^{UB}$
Inc. ($c'' > 0$)	$< \Delta^{LB}$	$> \Delta^{UB}$

Table 2: Summary of bounds on ATE

E Comparison of Bounds

[38] Recall that this line is decreasing in x because the treatment is being extended from $x = 1$ downwards.

Table 3: Comparison of bounds for unconstrained and capacity-constrained problems

	Unconstrained	Binding Capacity Constraint			
		Case $\overline{a}$		Case $\overline{b}$	
		$\kappa^*_{\overline{a}}$ unknown	$\kappa^*_{\overline{a}}$ known	$\kappa^*_{\overline{a}}$ unknown	$\kappa^*_{\overline{a}}$ known
ATT lower bound	$\Delta(\kappa^*)$	0	0	$\Delta(\overline{\kappa}^*)$	$\Delta(\overline{\kappa}^*)$
ATUT upper bound	$\Delta(\kappa^*)$	$\overline{\Delta}$	$[\Delta(\overline{\kappa}^*)\kappa^*_{\overline{a}} + \overline{\Delta}(\overline{\kappa}^* - \kappa^*_{\overline{a}})]/\overline{\kappa}^*$	$[\Delta(\overline{\kappa}^*)(\overline{\kappa}^* - (1-\overline{\mu})) + \overline{\Delta}(1-\overline{\mu})]/\overline{\kappa}^*$	$[\Delta(\overline{\kappa}^*)(\overline{\kappa}^* - (1-\overline{\mu}) + \kappa^*_{\overline{a}}) + \overline{\Delta}(1-\overline{\mu}-\kappa^*_{\overline{a}})]/\overline{\kappa}^*$
ATE lower bound	$\underline{\Delta}\kappa^* + \Delta(\kappa^*)(1-\kappa^*)$	$\underline{\Delta}\overline{\kappa}^*$	$\underline{\Delta}\overline{\kappa}^*$	$\underline{\Delta}\overline{\kappa}^* + \Delta(\overline{\kappa}^*)(1-\overline{\kappa}^*)$	$\underline{\Delta}\overline{\kappa}^* + \Delta(\overline{\kappa}^*)(1-\overline{\kappa}^*)$
ATE upper bound	$\Delta(\kappa^*)\kappa^* + \overline{\Delta}(1-\kappa^*)$	$\overline{\Delta}$	$\Delta(\overline{\kappa}^*)\kappa^*_{\overline{a}} + \overline{\Delta}(1-\kappa^*_{\overline{a}})$	$\Delta(\overline{\kappa}^*)(\overline{\kappa}^* - (1-\overline{\mu})) + \overline{\Delta}(1-(\overline{\kappa}^* - (1-\overline{\mu})))$	$\Delta(\overline{\kappa}^*)(\overline{\kappa}^* - (1-\overline{\mu}) + \kappa^*_{\overline{a}}) + \overline{\Delta}(1-(\overline{\kappa}^* - (1-\overline{\mu})) - \kappa^*_{\overline{a}})$

Note: Assuming $\underline{\Delta} < 0$; otherwise the ATT lower bound in uninformative in Case $\overline{a}$. These results assume a constant marginal cost of treatment. Appendix D examines nonconstant marginal costs of treatment.

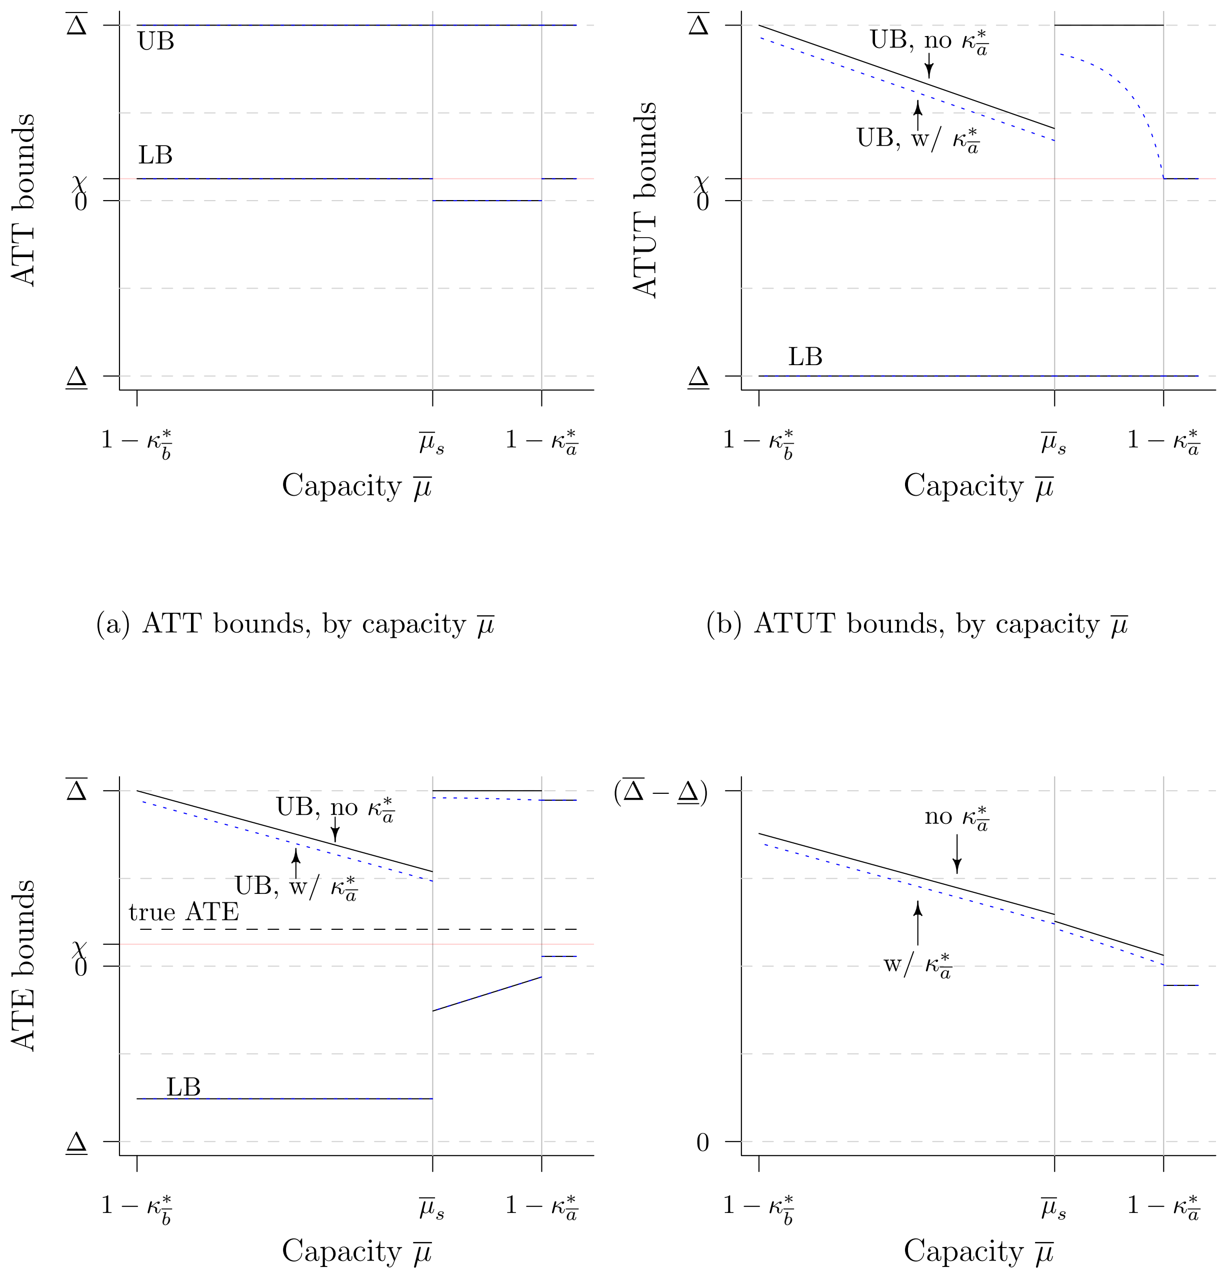

(a) ATT bounds, by capacity $\overline{\mu}$

(b) ATUT bounds, by capacity $\overline{\mu}$

(c) ATE bounds, by capacity $\overline{\mu}$

(d) Width of ATE bound, by capacity $\overline{\mu}$

Figure 6: Bounds on treatment effects, by capacity $\overline{\mu}$

Table 3 presents bounds in the unconstrained and capacity-constrained cases, assuming a constant marginal cost of treatment.[39] To provide intuition for how these bounds might look I computed these bounds using $(\Delta(\cdot), \chi)$ from the example shown in Figure 2, as the capacity $\overline{\mu}$ increases from $1 - \kappa^*_{\overline{b}}$—making $\overline{\kappa}^* = \kappa^*_{\overline{b}}$ feasible—to one—allowing treatment of the population; i.e., the measure of over-capacity units, $1 - \overline{\mu}$, would be zero. The value of $\overline{\mu}$ determines the relevant scenario. When $\overline{\mu} = 1 - \kappa^*_{\overline{b}}$ we are in Case $\overline{b}$. We remain in Case $\overline{b}$ as $\overline{\mu}$ increases, until we switch into Case $\overline{a}$ at $\overline{\mu} = \overline{\mu}_s$, the capacity above which the administrator would find it optimal to (also) treat units $x \in [1 - \overline{\mu}, \kappa^*_{\overline{b}})$. We remain in Case $\overline{a}$ for increasing capacities until $\overline{\mu} = 1 - \kappa^*_{\overline{a}}$, at which point the capacity constraint no longer binds.

Figure 6a presents bounds on the ATT as a function of $\overline{\mu}$. The values $\overline{\Delta}$ and $\underline{\Delta}$, indicated on the y-axis, respectively denote the uninformative upper and lower bounds of $\Delta(x)$. The ATT upper bound, indicated by the solid black line labeled "UB", cannot be tightened from $\overline{\Delta}$. At $\overline{\mu} = 1 - \kappa^*_{\overline{b}}$ (meaning we are in Case $\overline{b}$) the ATT lower bound, indicated by the solid black line labeled "LB", is $\Delta(\overline{\kappa}^*) = \chi$, tighter than the uninformative lower bound of $\underline{\Delta}$. As we increase $\overline{\mu}$, the measure of students treated and, hence, $\overline{\kappa}^*$ and the associated estimated treatment effect at the cutoff, remain constant until $\overline{\mu} = \overline{\mu}_s$, at which point we switch into Case $\overline{a}$. Thus, the black line representing the ATT lower bound is constant at the marginal cost of treatment χ when $\overline{\mu} < \overline{\mu}_s$. When the capacity locally binds in Case $\overline{a}$ the lower bound on the ATT is looser than it would be in Case $\overline{b}$, which is indicated by the drop in Figure 6a. However, when $\overline{\mu} \geq 1 - \kappa^*_{\overline{a}}$ the constraint no longer binds and the ATT lower bound jumps up again to the marginal cost of treatment. Note that knowledge of $\kappa^*_{\overline{a}}$ does not tighten the ATT lower bound.

Figure 6b presents bounds on the ATUT as a function of $\overline{\mu}$. The ATUT lower bound cannot be tightened from $\underline{\Delta}$, which is indicated by the solid black line labeled "LB". Starting in Case $\overline{b}$, in the case where we do not know $\kappa^*_{\overline{a}}$, indicated by the black line labeled "UB, no $\kappa^*_{\overline{a}}$", we can see the upper bound on the ATUT tightens as we increase the capacity $\overline{\mu}$. Intuitively, the knowledge that treating units between $1-\overline{\mu}$ and $\kappa^*_{\overline{b}}$ would be suboptimal leads to tighter upper bounds on the untreated as this measure of units increases. However, when $\overline{\mu} \geq \overline{\mu}_s$ in Case $\overline{a}$ then, without knowledge of $\kappa^*_{\overline{a}}$, the administrator's capacity constraint becomes locally binding, meaning we cannot rule out any values for $\Delta(x)$ for $x \leq 1 - \overline{\mu}_s$. When the constraint no longer binds, however, we recover the ATUT upper bound from the unconstrained case, χ. Knowledge of $\kappa^*_{\overline{a}}$ tightens the upper bound on the ATUT, as units $x < \kappa^*_{\overline{a}}$ would not be treated and, therefore, have an average $\Delta(x)$ less than χ (dashed blue line labeled "UB, w/ $\kappa^*_{\overline{a}}$"); this bound decreases, or becomes tighter, as the capacity

[39] Making an assumption about the extent to which marginal cost of treatment varied would afford computation of bounds when marginal costs were not constant. See Appendix D.

increases because the measure of students with an upper bound of $\overline{\Delta}$ decreases as $\overline{\mu}$ increases, while in Case $\overline{a}$ (additionally, in this example, the RD-estimated LATE would decrease as $\overline{\mu}$ increases, creating a concave dashed blue line).

Figure 6c combines the upper bound on the ATUT and lower bound on the ATT to present bounds on the ATT. The actual value of the ATE for this example is depicted in the dashed black line labeled "true ATE". As was the case with the ATT, starting in Case $\overline{b}$ at $\overline{\mu} = 1 - \kappa^*_{\overline{b}}$, the lower bound on the ATE is tighter than $\underline{\Delta}$ but does not vary with $\overline{\mu}$, because the measure of treated students would not change in Case $\overline{b}$ (nor can the ATT lower bound be tightened). The tighter ATUT upper bound results in a tighter ATE upper bound both when $\kappa^*_{\overline{a}}$ is unknown and known (solid black and dotted blue lines, respectively), which, as with the ATUT, tighten as capacity increases while in Case $\overline{b}$. Although the ATT lower bound is looser in Case $\overline{a}$, the discrete increase in the measure of treated students leads to a jump up in the ATE lower bound when we switch from Case $\overline{b}$ to Case $\overline{a}$. In contrast, the upper bound on the ATE is uninformative in Case $\overline{a}$ when we do not know $\kappa^*_{\overline{a}}$ (solid black line)—as was the case with the ATUT upper bound. It would be slightly lower when we did know $\kappa^*_{\overline{a}}$ (dotted blue line), and, as with the ATUT upper bound, would decrease as capacity increased while in Case $\overline{a}$. Finally, the ATE lower bound is the tightest when the constraint no longer binds, as the lower bound on the ATT is tightest and the ATT comprises a larger measure of units than in either binding case. Though the ATUT bound here would be tightest, the relatively small value of $\kappa^*_{\overline{a}}$ chosen in this example, on the other hand, leads to an ATE upper bound that is looser than in some of the constrained scenarios. As expected, the true value of the ATE is contained in the ATE bounds.

Figure 6d presents the width of the ATE bounds (i.e., the difference between upper and lower bounds). We can see that the ATE bound width is decreasing in $\overline{\mu}$, and would always be smaller when we know $\kappa^*_{\overline{a}}$.