

**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; December 7, 2017



Centre for Human Capital and Productivity (CHCP)

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*

December 7, 2017

Abstract

Regression-discontinuity (RD) designs estimate treatment effects only around 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. 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. Bounds are typically looser if the capacity constraint binds. Testable implications of cutoff optimality are derived. The results are demonstrated using previous RD studies.

*Department of Economics, University of Western Ontario. Email: nirav.mehta@uwo.ca. I thank Tim Conley, Nail Kashaev, Salvador Navarro, Shalini Roy, Paul Sullivan, and Todd Stinebrickner for helpful discussions. Mark Hoekstra provided me with assistance in applying this paper's results. I also thank George Orlov and Enrique Martin Luccioni for research assistance. This research was supported by the Social Sciences and Humanities Research Council.

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 average 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 units 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, the price to pay is that without more structure, the addition of which may be appropriate in some contexts, they are completely uninformative about treatment effects elsewhere. 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 treatment effect parameters 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 average treatment effect at the cutoff with a simple model of cutoff choice yields many insights. The most basic inference 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: we can reject cutoff optimality if the average treatment effect at the cutoff is negative. If the marginal cost is nonincreasing and the cutoff does not reflect a binding capacity constraint, optimality implies that the average treatment effect is nondecreasing at the cutoff (when approaching the cutoff from the untreated side) and several additional results relate the average 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

¹See 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.

the gain from treatment is very large if, as is commonly assumed by practitioners, average 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, obviating estimating the cost of treatment. Additionally, we can rule out constant treatment effects, 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 population average treatment effect (ATE). Bounds on the ATE are looser if the marginal cost of treatment is strictly increasing; they are also typically looser (and never tighter) if 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 should be lower than 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, *ceteris paribus*, 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)).

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.² 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

²Applications 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 average treatment effect is nonnegative and nondecreasing (when approaching from the untreated side) at the cutoff.

administrators.

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 paper simply shows 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 average treatment effects and her cost function. These informational assumptions simplify this paper's analysis at the cost of being stronger than those invoked in the Roy model, where individuals are only assumed to know their potential outcomes. Therefore, in addition to testing necessary conditions of optimality, I also examine how the results are affected by violations of these assumptions.

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 (1989), Manski (1990), Manski (1997), and Manski and Pepper (2000), which examine how boundedness and monotonicity assumptions can provide useful information about treatment effect parameters. 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. Though it uses stronger informational assumptions, some of the basic 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.

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 average 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.

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. 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 on average for students, given their characteristics, 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.³ 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 i . Student i has running variable w_i , which has continuous CDF F_w ; for example, the running variable could be student SAT scores. Let the random variable $Y_{\tau i}$ denote student i ’s potential outcome under treatment status τ , where $\tau = 1$ denotes the treatment and $\tau = 0$ the control. For example, this may be the wage earned as a function of being enrolled in the training program. The treatment effect for student i is $\Delta_i \equiv Y_{1i} - Y_{0i}$ and the average treatment effect for students with running variable w is $E[Y_{1i} - Y_{0i} | w_i = w]$. Note that there may be heterogeneous treatment effects among students with running variable w . It will be convenient to define the average treatment effect among students at an index denoting the x th quantile of w using the function $\Delta(x) \equiv E[Y_{1i} - Y_{0i} | w_i = F_w^{-1}(x)]$. Note that x is uniformly distributed over $[0, 1]$ (Embrechts and Hofert, 2013, Proposition 3.1). 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 administrator is constrained to choose a cutoff rule where $\tau(x) = 1$ if and only if $x \geq \kappa$ for some $\kappa \in (0, 1)$.⁴ To simplify exposition, I assume a “sharp” RD design, which

³Allowing 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.

⁴Note 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 expected gain from serving each quantile of students was at least as large as the marginal cost of serving them, point-wise. This assumption would imply that expected 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 expect to gain from treatment: $\int_0^1 1\{\Delta(x) > 0\}dx \geq \mu$. This

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).⁵ Let κ^* denote the treatment cutoff chosen by the administrator. The measure of students receiving treatment is $\mu^* = \int_0^1 \tau(x)dx = 1 - \kappa^*$; if there is a non-unit measure of students then let μ denote the share of students treated. 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 E). If $c(\mu)$ cannot capture such heterogeneity, then treatment effects 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 problem.

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 that κ^* , the treatment cutoff, was chosen by the administrator and the value of $\Delta(\kappa^*)$, from a RD design? Though I find that we can not say much about $\Delta(\cdot)$ for specific values of x that are not at the cutoff, we will be able to bound averages of $\Delta(\cdot)$ over different intervals; in particular, I focus on the average effect of treatment on the treated (ATT), untreated (ATUT), and population (ATE):

$$\begin{aligned} \text{ATT} &= \int_{\kappa^*}^1 \frac{\Delta(x)}{1-\kappa^*} dx \\ \text{ATUT} &= \int_0^{\kappa^*} \frac{\Delta(x)}{\kappa^*} dx \\ \text{ATE} &= \int_0^1 \Delta(x) dx, \end{aligned} \tag{1}$$

where superscripts are used to denote the lower and upper bounds of parameters, e.g., ATE^{LB} and ATE^{UB} respectively denote the lower and upper bounds on the population average treatment effect. The “local average treatment effect” (LATE) at the treatment cutoff is simply $\Delta(\kappa^*)$.

I briefly remark on notation before defining the administrator's problem. In a perhaps nonstandard manner, I define and study the average treatment effect at quantiles of the running variable w instead of at values of w . This is because the cost function is defined in terms of the measure (or share) of treated students $\mu \in [0, 1]$. Since treatment status depends only on which side of the cutoff a student lies, the ATT and ATUT (and population ATE) are unaffected by this choice.

would be relevant if, say, the population voted on whether to implement the treatment.

⁵Section 4 shows that the theoretical results are identical under a “fuzzy” design where, instead of perfect compliance, the probability of participation is constant other than for a discontinuous change at the treatment cutoff. Appendix A examines the case where probabilities are not constant on either side of the cutoff.

The administrator’s problem is to choose a cutoff to maximize the total treatment effect from treating the treated, net treatment cost, which is equivalent to maximizing the total average treatment effect from treating the treated, net treatment cost:⁶

$$\max_{\tilde{\kappa}} \beta \left(\int_{\tilde{\kappa}}^1 \Delta(x) dx \right) - c(1 - \tilde{\kappa}), \quad (2)$$

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 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.⁷

Assumption 1. *The following assumptions about costs and benefits of treatment are maintained until superseded:*

- (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) *The average treatment effect $\Delta(\cdot)$ is differentiable in x and known by the administrator.*
- (iii) *There exist finite lower and upper bounds on $\Delta(\cdot)$; 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, which assumes costs are separable across treated units. All the following results would obtain in the more general case where the marginal cost of treatment

⁶For an example of a slightly different objective function, see Heinrich et al. (2002), who study treatment decisions when administrators face performance standards.

⁷In the baseline case presented here, the administrator is a utilitarian who weighs gains for all students equally. Appendix D considers a case where gains are not weighed equally; the results derived in Section 3.1 are also obtained there.

was nonincreasing in μ , i.e., where the cost of treatment is weakly concave in μ .⁸ 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.⁹ 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 average treatment effects for students with index x ; so long as the administrator has an unbiased signal of $\Delta(x)$, uncertainty about average treatment effects does not affect the analysis, as the administrator’s objective is linear. When the bias can depend on x , however, bounds vary depending on the relationship between the bias and x . If the bias is positive for treated students and negative for untreated students then bounds are looser (i.e., wider); if instead the bias is negative for treated students and positive for untreated students then bounds are tighter.¹⁰ Assumption 1(iii) is an implication of potential outcomes $Y_{\tau i}$ having common and finite support for all i , which makes sense for outcomes such as wages, test scores, or probabilities.¹¹ 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 E, which studies variable marginal costs of treatment.

In what follows, all bounds presented in this paper are “sharp”, i.e., they are the tightest possible bounds given the model and data. This is because they are characterized by necessary and sufficient conditions. I refer to a bound as “uninformative” when it cannot be tightened relative to the bound on the parameter space, specified in Assumption 1(iii). For a parameter in eq. (1), an informative lower bound is a lower bound above $\underline{\Delta}$, an informative upper bound is an upper bound below $\overline{\Delta}$, and uninformative bounds are $[\underline{\Delta}, \overline{\Delta}]$.

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.

⁸Note 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 E for details).

⁹Note 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.

¹⁰The results for how uncertainty and bias affect bounds are derived in Appendix B.

¹¹Note that $Y_{\tau i}$ may contain negative values, which may capture a negative treatment effect or a positive opportunity cost of participating in an ineffective treatment.

3.1 Results without Capacity Constraints

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 unit interval) κ^* , which is not restrictive when analyzing results from an RD design.

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

(i) *Marginal Benefit = Marginal Cost:* $\Delta(\kappa^*) = c'(1 - \kappa^*) = \chi$

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

Proof. Differentiate the administrator's problem (2) 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). \square

Condition 1(i) will play a key role in bounding average 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 (2).*

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 eq. (2) 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.:*

$$\text{Participation: } \int_{\kappa^*}^1 \Delta(x) dx \geq c(1 - \kappa^*) = \chi(1 - \kappa^*). \quad (3)$$

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

$$\int_{\hat{\kappa}}^{\kappa^*} \Delta(x) dx < c(1 - \hat{\kappa}) - c(1 - \kappa^*) = \chi(\kappa^* - \hat{\kappa}). \quad (4)$$

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^* (recall that units above the cutoff are treated). 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 $\int_{\hat{\kappa}}^{\kappa^*} (\Delta(x) - \chi) dx < 0$. 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 $[\hat{\kappa}, \kappa^*]$, 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^* . \square

It is often said that there is no reason to believe average 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 a constant average treatment effect. 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 treatment effects for subsets of students.

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

Proof. Divide eq. (3) by the measure of treated students $(1 - \kappa^*)$ and combine this with Condition 1(i) to obtain

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

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.




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

Corollary 2. *The ATT is positive.*

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

Proposition 1 shows that the discontinuity-based estimate provides a lower bound for the average 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. 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 cost of treating them. Note that the level of the marginal cost does not need to be known by the researcher. Clearly, there is no informative (i.e., less than $\overline{\Delta}$) upper bound on the ATT.

However, although we have quite a bit of information about averages of treatment effects over some intervals of interest, we cannot make statements about $\Delta(x)$ for students with index $x \neq \kappa^*$. Figure 1 plots an example average treatment effect function $\Delta(x)$ (dashed black curve) and marginal cost of treatment (solid red horizontal line) against the index x , 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 an average 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 $0 \leq a < b \leq \kappa^*$ if $\underline{\Delta} > \chi - (\bar{\Delta} - \chi)(b - a)/(\kappa^* - b)$.*

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 eq. (4) into two parts at \hat{x} and rearrange terms:

$$\int_{\hat{x}-\hat{\mu}}^{\hat{x}} \Delta(x)dx < c(1-(\hat{x}-\hat{\mu})) - c(1-\kappa^*) - \int_{\hat{x}}^{\kappa^*} \Delta(x)dx \Rightarrow \int_{\hat{x}-\hat{\mu}}^{\hat{x}} \Delta(x)dx < \chi(\kappa^* - (\hat{x}-\hat{\mu})) - \underline{\Delta}(\kappa^* - \hat{x}), \quad (5)$$

□

where the implication follows from Assumption 1(iii).¹²

The right side of eq. (5) in Proposition 2 provides an upper bound for the gain from treating students with indices $x \in [\hat{x} - \hat{\mu}, \hat{x}]$. Because we do not know $\Delta(\cdot)$, by assuming the worst possible average 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 Condition 2. Intuitively, this upper bound grows the further below the cutoff we go. To gain more intuition for Proposition 2, rearrange eq. (5) and divide by the measure of students under consideration $\hat{\mu}$ to obtain

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

The left side of eq. (6) is the average treatment effect among students with indices $x \in [\hat{x} - \hat{\mu}, \hat{x}]$. First, consider the extreme scenario where we want an upper bound for the average treatment effect for a student with index \hat{x} , $\Delta(\hat{x})$. Take the limit of eq. (6) as the measure of additional treated students goes to zero:

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

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

¹²Note this bound will be informative for all but very low values of $\underline{\Delta}$, i.e., those satisfying $\underline{\Delta} > \chi - (\bar{\Delta} - \chi)\hat{\mu}/(\kappa^* - \hat{x})$. This condition is always satisfied when using Proposition 2 to bound the ATUT or ATE.

administrator is considering extending treatment to all students below \hat{x} :

$$\int_0^{\hat{x}} \frac{\Delta(x)}{\hat{x}} dx < (\chi - \underline{\Delta}) \left(\frac{\kappa^*}{\hat{x}} \right) + \underline{\Delta}. \quad (7)$$

Equation (7) says that there is an informative upper bound on the average treatment effect among students with indices $x \leq \hat{x}$ (the left side), which 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.

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

Proof. Let $\hat{x} = \hat{\mu} = \kappa^*$ in eq. (5) and divide through by κ^* to obtain the result:

$$\underbrace{\int_0^{\kappa^*} \frac{\Delta(x)}{\kappa^*} dx}_{\text{ATUT}} < \frac{\chi \kappa^*}{\kappa^*} = \chi = \Delta(\kappa^*), \quad (8) \quad \square$$

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 prior results can be used to bound the population average treatment effect (ATE).

Corollary 4. *The ATE has informative bounds, with a lower bound of $ATE^{LB} = \underline{\Delta}\kappa^* + \Delta(\kappa^*)(1 - \kappa^*)$ and an upper bound of $ATE^{UB} = \Delta(\kappa^*)\kappa^* + \bar{\Delta}(1 - \kappa^*)$.*

Proof. Lower bound: Measure κ^* units are untreated, and, by Assumption 1(iii), the average 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 $ATE^{LB} = \underline{\Delta}\kappa^* + \Delta(\kappa^*)(1 - \kappa^*)$.

Upper bound: Corollary 3 says that the ATUT is no larger than $\Delta(\kappa^*)$. By Assumption 1(iii), the average treatment effect for any treated unit cannot exceed $\bar{\Delta}$. Integrate and sum to form $ATE^{UB} = \Delta(\kappa^*)\kappa^* + \bar{\Delta}(1 - \kappa^*)$. \square

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 $\text{ATUT} < \Delta(\kappa^*) \leq \text{ATT}$. Moreover, we can contrast what can be said about treatment effects for students with indices $x > \kappa^*$ and those for students 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 with indices $x > \kappa^*$. But the fact that the administrator is not choosing to extend (i.e., decrease) the cutoff to students with index $\hat{x} < \kappa^*$ provides us information for how large the average treatment effect can possibly be for students with indices $x \in [\hat{x}, \kappa^*)$.

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 E. 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 $\text{ATUT} < \Delta(\kappa^*) \leq \text{ATT}$.

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 eq. (3) 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 eq. (4)) 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 next.

3.2 Results with Capacity Constraints

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

$$\begin{aligned} \max_{\tilde{\kappa}} \quad & \beta \left(\int_{\tilde{\kappa}}^1 \Delta(x) dx \right) - c(1 - \tilde{\kappa}) \\ \text{s.t.} \quad & 1 - \tilde{\kappa} \leq \bar{\mu}. \end{aligned} \quad (\bar{2})$$

If the desired (i.e., unconstrained) measure of treated students does not exceed capacity, i.e., $\mu^* (\equiv 1 - \kappa^*) \leq \bar{\mu}$, then the constraint does not bind and the optimal cutoff κ^* and resulting analysis are unaffected, meaning the results from Section 3.1 would also apply here. 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.




Figure 2: Types of binding capacity constraints

Let $\bar{\kappa}^*$ denote the binding solution to the constrained problem $(\bar{2})$.¹³ There are two potential cases corresponding to optimal cutoffs, depicted in Figure 2. The cutoff κ_a^* indicates

¹³Assumption 2 is augmented slightly for this section, to also assume that $\bar{\kappa}^*$ uniquely maximizes the administrator's constrained problem, $(\bar{2})$.

the optimal cutoff, were the administrator unconstrained. Figure 2a indicates the administrator’s constrained-optimal cutoff in Case \bar{a} , where $\bar{\mu}$ is such that the administrator would set $\bar{k}^* = 1 - \bar{\mu}$, i.e., the capacity constraint is locally binding. For example, Case \bar{a} would apply if $1 - \bar{\mu}$ was just above κ_a^* . However, if $\bar{\mu}$ decreased by enough, then the gains from extending treatment to units up until capacity may not be worth the cost. This is Case \bar{b} , depicted in Figure 2b, where a more severe capacity constraint would cause the administrator to set the constrained-optimal cutoff to $\bar{\kappa}^* = \kappa_b^* > 1 - \bar{\mu}$. In Case \bar{b} , the capacity constraint is binding, but not locally. Note that, although Figure 2 depicts only two possible solutions, κ_a^* and κ_b^* , the following analysis does not depend on this.¹⁴ Rather, the following results rely only on knowledge of $\bar{\mu}$, $\bar{\kappa}^*$, $\Delta(\bar{\kappa}^*)$, and, potentially, κ_a^* . I start by adapting Condition 1.

Condition $\bar{1a}$ (Necessity). *If $\bar{\kappa}^* = 1 - \bar{\mu}$ solves the administrator’s capacity-constrained problem ($\bar{2}$), then $\Delta(\bar{\kappa}^*) > \chi$.*

Proof. Ignoring the measure-zero case(s) where $\Delta(\bar{\kappa}^*) = \chi$, 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. \square

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

Condition $\bar{1b}$ (Necessity). *If $\bar{\kappa}^* = \kappa_b^*$ solves the administrator’s capacity-constrained problem ($\bar{2}$), then the following necessary conditions must hold:*

(i) *Marginal Benefit=Marginal Cost: $\Delta(\bar{\kappa}^*) = \chi$*

(ii) *Increasing Marginal Benefit: $\Delta'(\bar{\kappa}^*) \geq 0$.*

Proof. Identical to Condition 1. \square

Conditions $\bar{1a}$ and $\bar{1b}(i)$ 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 average treatment effect derivative is nondecreasing at the cutoff—applies in Case \bar{b} but need not apply in Case \bar{a} . However, an alternative testable implication of optimality can be deduced by combining Condition 1(i) with Condition $\bar{1a}$: $\Delta(\bar{\kappa}^*) > \Delta(\kappa_a^*)$. 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 \bar{a} -binding (allowing estimation

¹⁴For example, $\Delta(\cdot)$ might cross χ twice between κ_a^* and $\bar{\kappa}^*$ in Figure 2a.

of $\Delta(\bar{\kappa}^*)$) 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 $\Delta(\kappa_{\bar{a}}^*)$). 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 \bar{a} . Interestingly, these last two results are only implementable if the constraint Case- \bar{a} -binds in at least one period—otherwise the cutoff, and resulting LATE, would be the same for both periods. Perhaps counterintuitively, it is possible that we could actually learn more when the constraint binds in at least one year.

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

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

$$\text{Participation: } \int_{\bar{\kappa}^*}^1 \Delta(x)dx > \chi(1 - \bar{\kappa}^*), \quad (\bar{3a})$$

where the strict inequality follows from combining Condition $\bar{1a}$ with eq. (3) from Condition 2.

Note that Condition $\bar{2a}$ does not have an analogue to the second part of Condition 2, i.e., eq. (4), which is about optimality of non-treatment of the untreated. 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 \bar{b} the following condition characterizes the cutoff.

Condition $\bar{2b}$. *Suppose $\bar{\kappa}^* = \kappa_{\bar{b}}^*$ solves eq. (2). 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.:*

$$\text{Participation: } \int_{\bar{\kappa}^*}^1 \Delta(x)dx \geq \chi(1 - \bar{\kappa}^*). \quad (\bar{3b})$$

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

$$\int_{\hat{\kappa}}^{\bar{\kappa}^*} \Delta(x)dx < \chi(\bar{\kappa}^* - \hat{\kappa}). \quad (\bar{4b})$$

The first part of Condition $\bar{2b}$ is the same as the first part of Condition 2. The second part, eq. ($\bar{4b}$), differs from eq. (4) in that it only applies to candidate values $\hat{\kappa}$ no less than

$1 - \bar{\mu}$. Intuitively, this reflects the fact that the administrator chose to set the cutoff at $\kappa_{\bar{b}}^*$ instead of $1 - \bar{\mu}$.

Neither Corollary 1 nor Proposition 1 obtain in Case \bar{a} but they, or an analogue, can be obtained for Case \bar{b} .¹⁵

Corollary $\bar{1b}$. *Suppose $\bar{\kappa}^* = \kappa_{\bar{b}}^*$ solves eq. $(\bar{2})$. 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 $\bar{1b}$. *Suppose $\bar{\kappa}^* = \kappa_{\bar{b}}^*$ solves eq. $(\bar{2})$. The ATT is bounded below by the LATE at the treatment cutoff.*

Proof. Identical to Proposition 1. □

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

Proof. Because the marginal cost of treatment is positive (Assumption 1(i)), eq. $(\bar{3a})$ implies that $\int_{\bar{\kappa}^*}^1 \Delta(x) dx > \chi(1 - \bar{\kappa}^*) > 0$; divide through by $(1 - \bar{\kappa}^*)$ to obtain the result. □

Corollary $\bar{2a}$ shows that we can bound the ATT from below by zero. This lower bound is looser than that for the unconstrained ATT, obtained in Proposition 1, which was shown to be positive. This because, though we know that both $\Delta(\bar{\kappa}^*)$ and the ATT are nonnegative, we do not have enough information to order them. For example, consider $\Delta(\cdot)$ such that $\Delta'(x) < 0$ for all $x \geq \bar{\kappa}^*$; in this case the RD estimate would bound the ATT *from above*. However, we could use eq. $(\bar{3a})$ to tighten this bound if information about χ , the marginal cost of treatment, were available.¹⁶ In contrast, in Case \bar{b} we obtain the tighter bound obtained in the unconstrained problem.

Corollary $\bar{2b}$. *Suppose $\bar{\kappa}^* = \kappa_{\bar{b}}^*$ solves eq. $(\bar{2})$. The ATT is positive.*

Proof. Identical to Corollary 2. □

¹⁵Corollary 1(ii) would also obtain in Case \bar{a} if the administrator could choose which side of the cutoff to treat.

¹⁶Assuming β was also known.

As in the unconstrained scenario, we can bound the ATUT in both Case \bar{a} and Case \bar{b} . Due to their length and similarity to earlier results, the remaining proofs in this section are in Appendix C.

Corollary $\bar{3a}$. *Suppose $\bar{\kappa}^* = 1 - \bar{\mu}$ solves eq. ($\bar{2}$). The following ATUT bounds can be obtained:*

- (i) *There is an informative upper bound, of $ATUT^{UB} = [\Delta(\bar{\kappa}^*)\kappa_a^* + \bar{\Delta}(\bar{\kappa}^* - \kappa_a^*)]/\bar{\kappa}^*$, if and only if κ_a^* is known.*
- (ii) *There is an informative lower bound, of $ATUT^{LB} = \kappa_a^*\underline{\Delta}/\bar{\kappa}^*$, if and only if κ_a^* is known and $\underline{\Delta} < 0$.*

Proof. See Appendix C. □

Corollary $\bar{3a}(i)$ shows that knowledge of how severe the capacity constraint is, i.e., $(\kappa_a^*/\bar{\kappa}^*)$, tightens the upper bound on the ATUT. The extent to which it can be tightened depends on κ_a^* ; however, this bound is still looser than in the unconstrained case because $\Delta(\bar{\kappa}^*) > \Delta(\kappa_a^*)$. Intuitively, if the desired measure of treated students is not known, then we would have to set $\kappa_a^* = 0$ to maximize the upper bound, returning the uninformative upper bound of $\bar{\Delta}$. Similarly, Corollary $\bar{3a}(ii)$ shows that knowledge of $(\kappa_a^*/\bar{\kappa}^*)$ can also tighten the lower bound on the ATUT; if κ_a^* were not known then we would have to set $\kappa_a^* = \bar{\kappa}^*$, resulting in an uninformative lower bound of $\underline{\Delta}$. Note that, similar to the unconstrained case, the fact that the administrator would have desired to treat students between κ_a^* and $\bar{\kappa}^*$ does not provide information about the lower bound for strict subsets of students in this interval.

Corollary $\bar{3b}$. *Suppose $\bar{\kappa}^* = \kappa_b^*$ solves eq. ($\bar{2}$). The following ATUT bounds can be obtained:*

- (i) *There is an informative upper bound, which is tighter if κ_a^* is known. If κ_a^* is unknown the upper bound is $ATUT^{UB} = [\Delta(\bar{\kappa}^*)(\bar{\kappa}^* - (1 - \bar{\mu})) + \bar{\Delta}(1 - \bar{\mu})]/\bar{\kappa}^*$; if κ_a^* is known the upper bound is $ATUT^{UB} = [\Delta(\bar{\kappa}^*)(\bar{\kappa}^* - ((1 - \bar{\mu}) - \kappa_a^*)) + \bar{\Delta}((1 - \bar{\mu}) - \kappa_a^*)]/\bar{\kappa}^*$.*
- (ii) *There is an informative lower bound, which is tighter if κ_a^* is known. If κ_a^* is unknown the lower bound is $ATUT^{LB} = [\Delta(\bar{\kappa}^*)(\bar{\kappa}^* - (1 - \bar{\mu})) + \underline{\Delta}(1 - \bar{\mu})]/\bar{\kappa}^*$; if κ_a^* is known the lower bound is $ATUT^{LB} = [\Delta(\bar{\kappa}^*)(\bar{\kappa}^* - \kappa_a^*) + \underline{\Delta}\kappa_a^*]/\bar{\kappa}^*$.*

Proof. See Appendix C. □

Corollary $\overline{3b}$ (i) shows that, in Case \overline{b} , the upper bound on the ATUT gets tighter as capacity $\overline{\mu}$ increases. Note that the ATUT upper bound is tighter in Case \overline{b} than in Case \overline{a} . Corollary $\overline{3b}$ (ii) shows that, as in Case \overline{a} , information about $(\kappa_a^*/\overline{\kappa}^*)$ can also tighten the lower bound on the ATUT.

Finally, we can combine the previous results to bound the ATE.

Corollary $\overline{4a}$. *Suppose $\overline{\kappa}^* = 1 - \overline{\mu}$ solves eq. $(\overline{2})$. The following ATE bounds can be obtained.*

- (i) *If $\underline{\Delta} < 0$ and κ_a^* is unknown then there is an informative lower bound, of $ATE^{LB} = \underline{\Delta}\overline{\kappa}^*$, which is tightened to $ATE^{LB} = \underline{\Delta}\kappa_a^*$ if κ_a^* is known.*
- (ii) *There is an informative upper bound, of $ATE^{UB} = \Delta(\overline{\kappa}^*)\kappa_a^* + \overline{\Delta}(1 - \kappa_a^*)$, if and only if κ_a^* is known.*

Proof. See Appendix C. □

There is no informative upper bound for the ATE if κ_a^* is unknown. Of course, if χ and κ_a^* were both known then the ATE bounds in Case \overline{a} would be the same as in Section 3.1.

Corollary $\overline{4b}$. *Suppose $\overline{\kappa}^* = \kappa_b^*$ solves eq. $(\overline{2})$. The following ATE bounds can be obtained.*

- (i) *If κ_a^* is unknown there is an informative lower bound, of $ATE^{LB} = \Delta(\overline{\kappa}^*)(1 - (1 - \overline{\mu})) + \underline{\Delta}(1 - \overline{\mu})$, which is tightened to $ATE^{LB} = \Delta(\overline{\kappa}^*)(1 - \kappa_a^*) + \underline{\Delta}\kappa_a^*$, if κ_a^* is known.*
- (ii) *If κ_a^* is unknown there is an informative upper bound, of $ATE^{UB} = \Delta(\overline{\kappa}^*)(\overline{\kappa}^* - (1 - \overline{\mu})) + \overline{\Delta}(1 - (\overline{\kappa}^* - (1 - \overline{\mu})))$, which is tightened to $ATE^{UB} = \Delta(\overline{\kappa}^*)(\overline{\kappa}^* - ((1 - \overline{\mu}) - \kappa_a^*)) + \overline{\Delta}(1 - (\overline{\kappa}^* - ((1 - \overline{\mu}) - \kappa_a^*)))$ if κ_a^* is known.*

Proof. See Appendix C. □

Table 1 presents bounds on the average effect of treatment on the treated, untreated, and population in the unconstrained and capacity-constrained cases, assuming a constant marginal cost of treatment.¹⁷ Comparing Case \overline{a} with the unconstrained problem, the lower bound on the ATT is lower—it is zero instead of the LATE at the cutoff. In contrast with the unconstrained case, the RD estimate does not bound the ATUT from above. However, knowledge of κ_a^* would tighten bounds, in particular, it would provide an informative *lower* bound on the ATUT, which contains units the unconstrained administrator would have treated. In contrast to Case \overline{a} , in Case \overline{b} there are informative lower *and* upper bounds on

¹⁷Making 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 E.

Table 1: Comparison of Bounds for Unconstrained and Capacity-Constrained Problems

	Unconstrained	Binding Capacity Constraint			
		Case \bar{a}		Case \bar{b}	
		no κ_a^*	with κ_a^*	no κ_a^*	with κ_a^*
ATT ^{LB}	$\Delta(\kappa^*)$	0	0	$\Delta(\bar{\kappa}^*)$	$\Delta(\bar{\kappa}^*)$
ATT ^{UB}	$\bar{\Delta}$	$\bar{\Delta}$	$\bar{\Delta}$	$\bar{\Delta}$	$\bar{\Delta}$
ATUT ^{LB}	$\underline{\Delta}$	$\underline{\Delta}$	$\kappa_a^* \underline{\Delta} / \bar{\kappa}^*$	$[\Delta(\bar{\kappa}^*)(\bar{\kappa}^* - (1 - \bar{\mu})) + \underline{\Delta}(1 - \bar{\mu})] / \bar{\kappa}^*$	$[\Delta(\bar{\kappa}^*)(\bar{\kappa}^* - \kappa_a^*) + \underline{\Delta} \kappa_a^*] / \bar{\kappa}^*$
ATUT ^{UB}	$\Delta(\kappa^*)$	$\bar{\Delta}$	$[\Delta(\bar{\kappa}^*)\kappa_a^* + \bar{\Delta}(\bar{\kappa}^* - \kappa_a^*)] / \bar{\kappa}^*$	$[\Delta(\bar{\kappa}^*)(\bar{\kappa}^* - (1 - \bar{\mu})) + \bar{\Delta}(1 - \bar{\mu})] / \bar{\kappa}^*$	$[\Delta(\bar{\kappa}^*)(\bar{\kappa}^* - ((1 - \bar{\mu}) - \kappa_a^*)) + \bar{\Delta}((1 - \bar{\mu}) - \kappa_a^*)] / \bar{\kappa}^*$
ATE ^{LB}	$\underline{\Delta} \kappa^* + \Delta(\kappa^*)(1 - \kappa^*)$	$\underline{\Delta} \bar{\kappa}^*$	$\underline{\Delta} \kappa_a^*$	$\Delta(\bar{\kappa}^*)(1 - (1 - \bar{\mu})) + \underline{\Delta}(1 - \bar{\mu})$	$\Delta(\kappa^*)(1 - \kappa_a^*) + \underline{\Delta} \kappa_a^*$
ATE ^{UB}	$\frac{\Delta(\kappa^*)\kappa^* + \bar{\Delta}(1 - \kappa^*)}{\bar{\Delta}(1 - \kappa^*)}$	$\bar{\Delta}$	$\frac{\Delta(\bar{\kappa}^*)\kappa_a^* + \bar{\Delta}(1 - \kappa_a^*)}{\bar{\Delta}(1 - \kappa_a^*)}$	$\frac{\Delta(\bar{\kappa}^*)(\bar{\kappa}^* - (1 - \bar{\mu})) + \bar{\Delta}(1 - (\bar{\kappa}^* - (1 - \bar{\mu})))}{\bar{\Delta}(1 - (\bar{\kappa}^* - (1 - \bar{\mu})))}$	$\frac{\Delta(\bar{\kappa}^*)(\bar{\kappa}^* - ((1 - \bar{\mu}) - \kappa_a^*)) + \bar{\Delta}(1 - ((1 - \bar{\mu}) - \kappa_a^*))}{\bar{\Delta}(1 - ((1 - \bar{\mu}) - \kappa_a^*))}$

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

the ATE even if κ_a^* is unknown, as the decision to not treat students between $1 - \bar{\mu}$ and κ_b^* contains useful information. Knowledge of κ_a^* would further tighten the ATE bounds in either Case \bar{a} or Case \bar{b} ; as shown in Section 3.3, this type of information could be obtained by comparing budget requests and realized allocations. Finally, though the derivative-based test of cutoff optimality does not apply when the constraint Case- \bar{a} -binds (though it would still apply in a Case- \bar{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 Distinguishing Between Unconstrained and Constrained Cases

Computing the relevant bounds requires knowing whether the administrator’s chosen cutoff corresponds to the unconstrained case, Case \bar{a} , or Case \bar{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 $\bar{\kappa}^*$ satisfies Case \bar{a} , i.e., the capacity constraint was locally binding. If instead, the administrator did not use the entire allocated budget we could surmise that $\bar{\kappa}^*$ satisfies Case \bar{b} . Such data could be obtained via analysis of, say, a state’s budget process. For example, in New York state, agencies make budget proposals, which after being amended are included in an executive budget; there is an end-of-cycle report on the realized amounts the next fiscal year.¹⁸ Alternatively, one could find a situation where the cutoff had been chosen before capacity constraints were implemented, or even obtain data on κ_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 \bar{a} and Case \bar{b} if we know the unconstrained optimal cutoff, i.e., κ_a^* , even when the marginal cost of treatment χ is unknown. For example, in Case \bar{a} one could compare an administrator’s budget request with the actual amount expended and exploit the fact that the fraction, $(\kappa_a^*/\bar{\kappa}^*)$, 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.¹⁹

3.4 Inference

Consider a partially-identified parameter $\theta = \text{ATT}, \text{ATUT}, \text{and ATE}$, where $\theta \in \Theta_{\text{ID}}$. The interval $\Theta_{\text{ID}} = [\theta^{\text{LB}}, \theta^{\text{UB}}] \subseteq \Theta \equiv [\underline{\Delta}, \bar{\Delta}]$ is the identified set (sometimes referred to as the “identification region”) for θ , i.e., the interval contained by the sharp bounds of θ at either end. For example, if θ were the ATT in the unconstrained case then $\Theta_{\text{ID}} = [\Delta(\kappa^*), \bar{\Delta}]$. Following Imbens and Manski (2004), I construct confidence intervals that asymptotically contain θ with a desired probability.²⁰

As shown in Table 1, bounds for the ATT, ATUT, and ATE in the unconstrained case only depend on $\kappa^*, \Delta(\kappa^*), \underline{\Delta}$, and $\bar{\Delta}$, while bounds in the capacity-constrained cases may additionally depend on κ_a^* and $\bar{\mu}$ (in addition to potentially depending on $\Delta(\bar{\kappa}^*)$ instead of $\Delta(\kappa^*)$). Treating the primitives of the environment as fixed within a particular period, κ^* (or, if applicable, $\bar{\kappa}^*$) can be treated as fixed, and the only source of variation in bounds estimates

¹⁸<https://www.budget.ny.gov/guide/brm/item2.html>

¹⁹Let the requested budget be $B^* = (1 - \kappa^*)\chi$ and the realized budget (which, in Case \bar{a} would also be the approved budget) be $\bar{B}^* = (1 - \bar{\kappa}^*)\chi$. We can compute $B^*/\bar{B}^* = ((1 - \kappa_a^*)\chi)/((1 - \bar{\kappa}^*)\chi) = (1 - \kappa_a^*)/(1 - \bar{\kappa}^*)$, which permits the solution for κ_a^* , even when marginal cost χ is unknown.

²⁰Tamer (2010) outlines two types of confidence intervals of potential interest in a partial-identification setting: the ones I consider here, which contain the partially identified parameter with a certain probability (Imbens and Manski (2004); Stoye (2009)), and ones that contain the identified set with a certain probability (Horowitz and Manski (2000)).

comes from sampling error in the estimate of $\Delta(\cdot)$ at the treatment cutoff.²¹ There are many alternatives to compute standard errors for the cutoff LATE, e.g., from a regression coefficient or via the bootstrap (e.g., Calonico et al. (2014)); this section takes the RD estimate and associated standard error as given.

Depending on θ and whether/how the capacity constraint is binding, different cases may apply. First, θ^{LB} and/or θ^{UB} may not depend on the estimated average treatment effect at the cutoff, in which case it/they would be treated as known.²² Second, if only θ^{LB} (or, mutatis mutandis, θ^{UB}) needed to be estimated then $\widehat{\Theta}_{\text{ID}} = [\widehat{\theta}^{\text{LB}}, \theta^{\text{UB}}]$, where $\widehat{\theta}^{\text{LB}}$ is identical to θ^{LB} except that $\widehat{\Delta}(\kappa^*)$ has been substituted for $\Delta(\kappa^*)$ everywhere. Under asymptotic normality, the $100(1 - \alpha)\%$ confidence interval for θ in this (one-sided) case would be $[\widehat{\theta}^{\text{LB}} - z_{1-\alpha}SE(\widehat{\theta}^{\text{LB}}), \theta^{\text{UB}}]$, where $z_{1-\alpha}$ is the $1 - \alpha$ quantile of the standard normal distribution and the standard error $SE(\widehat{\theta}^{\text{LB}})$ is calculated using the relevant bound in Table 1.²³ Third, if both θ^{LB} and θ^{UB} needed to be estimated then the $100(1 - \alpha)\%$ confidence interval for θ would be $[\widehat{\theta}^{\text{LB}} - c_n SE(\widehat{\theta}^{\text{LB}}), \widehat{\theta}^{\text{UB}} + c_n SE(\widehat{\theta}^{\text{UB}})]$, where c_n solves

$$\Phi \left(c_n + \frac{\widehat{\theta}^{\text{UB}} - \widehat{\theta}^{\text{LB}}}{\max\{SE(\widehat{\theta}^{\text{LB}}), SE(\widehat{\theta}^{\text{UB}})\}} \right) - \Phi(-c_n) = 1 - \alpha, \quad (9)$$

and where Φ denotes the standard normal CDF.²⁴ Table 2 summarizes identified sets and $100(1 - \alpha)\%$ confidence intervals for the main parameters in the scenario without capacity constraints. Confidence intervals for these parameters in other scenarios would be constructed by appropriately adjusting the relevant bounds in Table 1.

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 are in the economics of education and examine contexts where it seems reasonable to expect that program adminis-

²¹In practice, Θ may be unknown, making it necessary to use other information, e.g., estimates from another data sources, to estimate Θ . Since these estimates are likely very precise compared to the estimated average treatment effect at the cutoff, which typically relies on a small fraction of observations in the sample, I abstract from such uncertainty here.

²²For example, when there is a Case- \bar{a} -binding capacity constraint, the ATT has a lower bound of zero (and an uninformative upper bound, as in the unconstrained case).

²³Note that in the case where one bound is known, the corrections of Imbens and Manski (2004) and Stoye (2009), which account for the fact that θ can be up against only one of the bounds of Θ_{ID} when $\theta^{\text{UB}} - \theta^{\text{LB}} > 0$, are not necessary.

²⁴Note the expression for the critical value c_n is slightly different from that in Imbens and Manski (2004) because I use the standard error of the estimated bounds (both of which are computed using the same sample size n), canceling the \sqrt{n} in the numerator of their equation (7). If $\widehat{\theta}^{\text{UB}} - \widehat{\theta}^{\text{LB}}$ was not very large relative to the standard errors of the bounds estimates then one should use critical values from Stoye (2009).

Table 2: Identified Sets and Confidence Intervals for Parameters in Unconstrained Scenario

θ	Estimate of Identified Set for θ	100(1 - α)% Confidence Interval for θ
ATT	$[\widehat{\Delta}(\kappa^*), \overline{\Delta}]$	$[\widehat{\Delta}(\kappa^*) - z_{1-\alpha}SE(\widehat{\Delta}(\kappa^*)), \overline{\Delta}]$
ATUT	$[\underline{\Delta}, \widehat{\Delta}(\kappa^*)]$	$[\underline{\Delta}, \widehat{\Delta}(\kappa^*) - z_{1-\alpha}SE(\widehat{\Delta}(\kappa^*))]$
ATE	$[\underline{\Delta}\kappa^* + \widehat{\Delta}(\kappa^*)(1 - \kappa^*), \widehat{\Delta}(\kappa^*)\kappa^* + \overline{\Delta}(1 - \kappa^*)]$	$[\underline{\Delta}\kappa^* + \widehat{\Delta}(\kappa^*)(1 - \kappa^*) - c_nSE(\widehat{\Delta}(\kappa^*))(1 - \kappa^*), \widehat{\Delta}(\kappa^*)\kappa^* + \overline{\Delta}(1 - \kappa^*) + c_nSE(\widehat{\Delta}(\kappa^*))\kappa^*]$

Bounds are computed for unconstrained scenario (see Table 1). $SE(\widehat{\Delta}(\kappa^*))$ is the standard error of $\widehat{\Delta}(\kappa^*)$ and c_n is defined in eq. (9).

trators 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, e.g., GPA. 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 would result in 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 and then consider a context where the constraint is likely not binding. I check the model implication that the average treatment effect at the cutoff is nonnegative for both applications. For the latter, unconstrained, application, I also test whether the average treatment effect is increasing at the cutoff (approaching from the untreated side).²⁵ Reassuringly, I cannot reject that the cutoff was chosen optimally by an informed administrator in any of the falsification tests.

One of the studies employs a fuzzy design, so I first show how the earlier results pertaining to sharp designs generalize.²⁶ Some new notation is necessary. Let $\omega(x)$ denote the administrator’s intended treatment group for student with index x .²⁷ For example, if students with indices x above κ are targeted for a program, then $\omega(x) = 1$ for $x \geq \kappa$ and $\omega(x) = 0$ for $x < \kappa$. For simplicity, here I consider the case where the probability of being treated ($\tau = 1$) depends on ω according to $\rho_\omega = \Pr\{\tau = 1|\omega\}$;²⁸ Appendix A shows that the mean effect of

²⁵ Recall that this condition does not hold for a Case- \bar{a} -binding constraint.

²⁶These results are derived for when the capacity constraint is not binding; analogous results obtain when the constraint binds.

²⁷Recall that the index x is distributed uniformly over $[0, 1]$.

²⁸That 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.

intending to treat among the (intended) treated (ITT) and intending not to treat among the (intended) untreated (ITUT) can be bounded when treatment probabilities ρ_w also depend on x , as may be the case in some applications. 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; however, the probability of treatment must increase discontinuously at the cutoff (Hahn et al. (2001)). In a fuzzy design, the administrator chooses the treatment cutoff $\tilde{\kappa}$ to maximize her expected objective:²⁹

$$\max_{\tilde{\kappa}} \left(\rho_0 \int_0^{\tilde{\kappa}} (\Delta(x) - \chi) dx + \rho_1 \int_{\tilde{\kappa}}^1 (\Delta(x) - \chi) dx \right). \quad (10)$$

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 (i.e., where $\rho_0 = 0$ and $\rho_1 = 1$). Moreover, multiplying through by ρ_w shows that the fuzzy design returns exactly the same bounds for the ATT and ATUT as does the sharp design when ρ_w are constant within treatment status.³⁰

The tests of model assumptions can be described using a sharp design without any loss of generality.³¹ 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; although x is defined as the quantile of the running variable w , the testable implication regarding optimality would be the same regarding the derivative of the cutoff LATE estimated using the running variable. 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 (recall that units above the cutoff are treated), or $\Delta(\kappa^*) < 0$, because the marginal cost of treatment is positive, contradicting optimality of treating students at κ^* . Assume the expected outcome for a student with index x_i , $E[Y_{\tau(x_i),i}]$, depends on treatment status $\tau(x_i)$ and the (transformed) running variable $(x_i - \kappa^*)$ according to the following statistical

²⁹As 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 first application.

³⁰Note 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., $\rho_0 > 0$).

³¹Dong and Lewbel (2015) show that a similar result holds for fuzzy designs.

relationship:³²

$$E [Y_{\tau(x_i),i}] = \alpha_0 + \alpha_1(x_i - \kappa^*) + \alpha_2\tau(x_i) + \alpha_3\tau(x_i)(x_i - \kappa^*),$$

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

$$\check{Y}_i = E [Y_{\tau(x_i),i}] + \epsilon_i = \alpha_0 + \alpha_1(x_i - \kappa^*) + \alpha_2\tau(x_i) + \alpha_3\tau(x_i)(x_i - \kappa^*) + \epsilon_i. \quad (11)$$

The estimate of the LATE at the treatment cutoff is $\hat{\Delta}(\kappa^*) = \hat{\alpha}_2$. Dong and Lewbel (2015) show that the estimate of the average treatment effect derivative at the cutoff here would be $\hat{\Delta}'(\kappa^*) = \hat{\alpha}_3$. Therefore, the model has a testable implication, i.e., is falsifiable, by using $\hat{\alpha}_3$ to test the null hypothesis $H_0 : \alpha_3 \geq 0$, versus the alternative hypothesis $H_1 : \alpha_3 < 0$. 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), which studies the effect of attending a flagship public university on subsequent mean earnings for a sample of white males between the ages of 28 and 33. The objective considered in eq. (2), 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 an interest in the education of the state’s denizens, 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).³³

Hoekstra uses a fuzzy design in which treatment was targeted to students at or above a covariate-adjusted SAT score, i.e., $\omega(x) = 1 \Leftrightarrow x \geq \kappa^*$. The intended-treated students

³²This 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 $E [Y_{\tau(x_i),i}]$, 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.

³³Epple 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.

($\omega(x) = 1$) 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 $\Delta(x)$ then represents the oft-studied heterogeneity in average 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 average gain may vary between students with different preparation levels (i.e., SAT scores).

Research by Epple et al. (2006) shows that universities admit students until their capacity constraints bind, suggesting that Hoekstra (2009) reflects a Case- \bar{a} -binding capacity constraint.³⁴ The marginal cost of treatment is assumed to be constant. 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.³⁵ 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.³⁶

Recall that the fuzzy design returns exactly the same bounds as does the sharp design, when ρ_ω are constant within treatment status.³⁷ The main result reported in Hoekstra (2009) is that attending the flagship university increases earnings 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

³⁴ The method proposed earlier to distinguish between Case- \bar{a} - and Case- \bar{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- \bar{a} -binding constraint is reasonable, as it likely corresponds to more conservative bounds.

³⁵ 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.

³⁶ See Table II on page 907 of Epple et al. (2006).

³⁷ Although Hoekstra (2009) allowed for non-constant ρ_ω , the estimates of this relationship, $\rho_\omega(x)$, were unfortunately not available; this forces this application to treat the ρ_ω as constant within treatment status. Figure 1 of Hoekstra (2009) provides visual guidance suggesting that, though there may be appreciable variation in treatment probabilities with respect to x in the intended-untreated group, there is not very much variation in treatment probabilities with respect to x among students in the intended-treated group. If then, we could approximate ρ_1 as constant, this would imply the results for the lower bound on the ATT would be the same as for the sharp case. Regardless, although perhaps not ideal, assuming a constant ρ_ω does not substantially affect the results of this application, due to the loose bounds achieved in Case- \bar{a} -binding constrained case, summarized in Table 3.

Condition $\bar{1}a$; that is, we cannot falsify the model assumptions.³⁸

As noted by Hoekstra, this estimate seems fairly high. For example, Ashenfelter and Rouse (1998) estimate that an additional year of schooling increases earnings by 9%, while Behrman et al. (1996) find that an additional year of schooling increases earnings 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. The large estimated effect in Hoekstra (2009), though perhaps not quite expected in a context in which the treatment cutoff was chosen by a utility-maximizing administrator, is somewhat easier to rationalize when viewed through the lens of Condition $\bar{1}a$ —that the RD estimate exceeds the marginal cost of treatment when the capacity constraint binds. In contrast, it might be hard to believe that an unconstrained utility-maximizing administrator chose a cutoff with such a large estimated cutoff average treatment effect.

I now show how bounds on the ATT, ATUT, and ATE are tightened by this paper’s results. I first discuss how I calibrate values for $\underline{\Delta}$, $\bar{\Delta}$, and $\bar{\kappa}^*$, which are required to bound parameters of interest. The uninformative bounds $\underline{\Delta} = -\$7,200$ and $\bar{\Delta} = \$18,000$ are calibrated using minimum and maximum estimates of quantile treatment effects from attending flagship (versus non-flagship) public universities on earnings, as reported in Tables 5 and 6 of Andrews et al. (2012) (the working paper version of Andrews et al. (2016)); the minimum and maximum effects, -16% and 40%, respectively, are then multiplied by \$45,000, which is roughly the average baseline annual income for the untreated.³⁹ The share of intended-treated students $\bar{\kappa}^*$ is calibrated to 0.5 based on acceptance rates of 40% and 66% at AT-Austin and Texas A&M, respectively.⁴⁰

When we use the bounds for the ATT when the capacity constraint Case- \bar{a} -binds, we can only surmise that the effect of treatment on the treated is positive, i.e., $ATT > 0$. Table 3 summarizes the bounds. As we can see from Col. (1), although the lower bound on the ATT is tightened by \$7,200 (i.e., 29%, as shown in Col. (2), which presents the width of each identified set in terms of the potential (i.e., uninformative) width $\bar{\Delta} - \underline{\Delta}$) and the lower bound of the population ATE is tightened by \$3,600 (i.e., 14%, as shown in Col. (2)), the Case- \bar{a} -binding constraint results in fairly wide bounds for parameters of interest. For example, the interval for the population ATE contains zero.

³⁸ This result is from Table 1 of Hoekstra (2009), the first two rows of which report RD estimates of the enrollment effect (i.e., effect of treatment on the treated for students at the cutoff) of 0.223 (with a standard error of 0.079) and 0.216 (with a standard error of 0.081), which respectively did not and did control for additional variables (e.g., residual high school GPA).

³⁹This is based on correspondence with the author.

⁴⁰See <https://www.usnews.com/best-colleges/university-of-texas-3658> and <https://www.usnews.com/best-colleges/texas-am-college-station-10366>.

Table 3: Summary of Parameter Bounds in Hoekstra (2009), measured in \$

Parameter θ	Identified Set Θ_{ID} (1)	Relative Width of Identified Set Θ_{ID} (2)
ATT	[0, 18,000]	0.71
ATUT	[-7,200, 18,000]	1.00
ATE	[-3,600, 18,000]	0.86
$[\underline{\Delta}, \overline{\Delta}]$	[-7,200, 18,000]	

Bounds are computed assuming Case- \bar{a} -binding capacity constraint, where κ_a^* is unknown (see Table 1). $\underline{\Delta} = -7,200$ and $\overline{\Delta} = 18,000$ are calibrated using minimum and maximum estimates of quantile treatment effects from attending flagship (versus non-flagship) public universities on earnings, as reported in Tables 5 and 6 of Andrews et al. (2012); the minimum and maximum effects, -16% and 40%, respectively, are then multiplied by \$45,000. $\bar{\kappa}^*$ is calibrated to 0.5 based on acceptance rates of 40% and 66% at AT-Austin and Texas A&M, respectively. See <https://www.usnews.com/best-colleges/university-of-texas-3658> and <https://www.usnews.com/best-colleges/texas-am-college-station-10366>.

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 contains 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 the effect of being placed on probation on subsequent GPA, which means that the average treatment effect $\Delta(x)$ is the expected gain in subsequent GPA if a student with prior GPA quantile 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 average 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 information made available by the journal’s replication policy, I ran regression eq. (11), the results of which are presented in Table 4. I find that $\hat{\alpha}_3 = 0.047$, with a standard error of 0.094, which means that there is not strong evidence that the average treatment effect is increasing at the cutoff (p-value 0.31 that the average treatment effect derivative is greater than zero).⁴¹ 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 4: Results from main specification in Lindo et al. (2010)

<i>Regressor:</i>	<i>Dependent variable:</i> GPA next semester (\check{Y}_i)
Intercept ($\hat{\alpha}_0$)	0.312 (0.019)
Running variable - cutoff ($\hat{\alpha}_1$)	0.699 (0.053)
Treatment indicator ($\hat{\alpha}_2$)	0.233 (0.031)
Treatment indicator \times (running variable - cutoff) ($\hat{\alpha}_3$)	0.047 (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

⁴¹The 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).

placed on academic probation on the next term’s grade performance for the full sample is $\widehat{\Delta}(\kappa^*) = 0.233$ higher GPA points.⁴² 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 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.⁴³

Lindo et al. (2010) affords us an opportunity to explore the upper bound on the average gain from treating all students with indices above a prospective cutoff \hat{x} . We must first adapt equation eq. (7) 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., $ATUT(\hat{x}) \equiv (\int_{\hat{x}}^1 \Delta(x)dx)/(1 - \hat{x}) \leq [\chi(1 - \kappa^*) - \underline{\Delta}(\hat{x} - \kappa^*)]/(1 - \hat{x})$. 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 $ATUT(\hat{x})$. By Condition 1(i) we can use the estimated average 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, p. 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 calibrate the time-cost of this hourly meeting to $\underline{\Delta} = -0.36/7 = -0.051$ GPA points.

Figure 3 illustrates the results. We can see that the upper bound on $ATUT(\hat{x})$ 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

⁴²This can be found in Table 5, in column (1) of panel A of Lindo et al. (2010).

⁴³ 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.

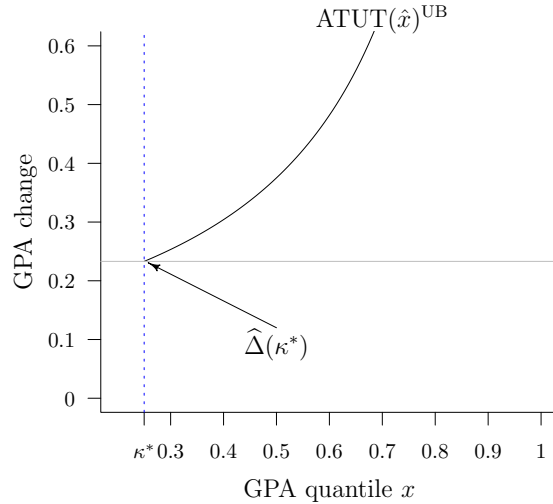


Figure 3: Upper Bound on Treating the “Rest of the Untreated” in Lindo et al. (2010)

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.

To show how much information is contained in this paper’s bounds, I report confidence intervals for the ATT, ATUT, and ATE, which depend on $(\underline{\Delta}, \overline{\Delta}, \kappa^*)$, in addition to the estimate of $\Delta(\kappa^*)$. I calibrate the remaining parameter $\overline{\Delta}$ by observing that Lindo et al. (2010) note that the GPA cutoff for receiving the treatment ranged from 1.5 to 1.6 between campuses (page 98); since the maximum GPA is 4.0, this results in a calibrated value of the maximum gain from treatment of $\overline{\Delta} = 2.5$ GPA points.

Table 5 shows bounds for the identified sets and confidence intervals for the ATT, ATUT, and ATE. Both the identified sets (Col. (1)) and confidence intervals (Col. (2)) for the ATT and ATUT are tighter than their possible ranges. This is perhaps more clearly shown in Col. (3), which presents the width of each confidence interval in terms of the potential (i.e., uninformative) width $\overline{\Delta} - \underline{\Delta}$. The 95% confidence interval for the population ATE is [0.007, 0.838] GPA points, which is one-third of the size of the potential width of the set containing this parameter, and strictly positive.⁴⁴

⁴⁴Following the formula for c_n in eq. (9), the lower bound is $\underline{\Delta}(1 - \kappa^*) + \widehat{\Delta}(\kappa^*)(\kappa^*) - 1.645SE(\widehat{\Delta}(\kappa^*))(\kappa^*)$ and the upper bound is $\widehat{\Delta}(\kappa^*)(1 - \kappa^*) + \overline{\Delta}(\kappa^*) + 1.645SE(\widehat{\Delta}(\kappa^*))(1 - \kappa^*)$, i.e., $c_n \approx z_{0.95}$.

Table 5: Summary of Parameter Bounds in Lindo et al. (2010), measured in GPA points

Parameter θ	Estimate of Identified Set $\widehat{\Theta}_{ID}$	95% Confidence Interval for θ	Relative Width of CI
	(1)	(2)	(3)
ATT	[0.233, 2.500]	[0.182, 2.500]	0.91
ATUT	[-0.051, 0.233]	[-0.051, 0.284]	0.13
ATE	[0.020, 0.800]	[0.007, 0.838]	0.33
$[\underline{\Delta}, \overline{\Delta}]$	[-0.051, 2.500]	–	

Bounds are computed assuming non-binding capacity constraint (see Table 1), adjusted for treatment below the cutoff. $\underline{\Delta} = -0.051$ GPA points is calibrated using opportunity cost of an hour of counseling per week ((Stinebrickner and Stinebrickner, 2008, Table 7a)) and $\overline{\Delta} = 2.5$ GPA points is calibrated using the largest difference between treatment cutoff (GPA of 1.5) and maximum GPA of 4.0 ((Lindo et al., 2010, p. 98)). κ^* is 0.25, per Lindo et al. (2010), page 101.

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.⁴⁵ We can exploit information revealed by the optimizing behavior of the administrator to extrapolate from the LATE at the treatment cutoff, which without further structure applies to only measure zero of the population, to obtain bounds for the effect of treatment on 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 for 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 average treatment effect at the cutoff is negative or decreasing (when approaching the cutoff from the untreated side), then we can reject that the cutoff was chosen optimally by an administrator informed about average 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. I cannot reject that the cutoff was chosen optimally by an informed administrator in any of the falsification tests conducted, which may increase confidence that the contexts studied

⁴⁵This paper also relates to work non-parametrically estimating treatment effects.

here constitute reasonable applications of the theoretical results.

The findings in this paper have several implications for the use of RD results in policy. 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 average 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 they should definitely treat those units treated by the administrator. Second, the upper bound on the average treatment effect 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 paper’s 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 only provide a lower bound for the ATT, there may be RD studies of useful—i.e., effective for treated units—programs that are simply not published because they lack statistically




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

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 larger-population-level (as opposed to site-specific) heterogeneity. Another variation would investigate what could be learned if the administrator only knew certain treatment effect parameters, 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

- Andrews, R. J., J. Li and M. Lovenheim, “Quantile Treatment Effects of College Quality On Earnings,” *Working Paper*, 2012.
- Andrews, R. J., J. Li and M. F. Lovenheim, “Quantile Treatment Effects of College Quality on Earnings,” *Journal of Human Resources*, 51(1):200–238, 2016.
- 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," *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," *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," *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.
- Calonico, S., M. D. Cattaneo and R. Titiunik, "Robust Nonparametric Confidence Intervals for Regression-Discontinuity Designs," *Econometrica*, 82(6):2295–2326, 2014.
- 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.
- Embrechts, P. and M. Hofert, "A Note on Generalized Inverses," *Mathematical Methods of Operations Research*, 77(3):423–432, 2013.
- 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,” *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.
- Hocking, T. D., *directlabels: Direct Labels for Multicolor Plots*, 2017, R package version 2017.03.31.
- Hoekstra, M., “The Effect of Attending the Flagship State University On Earnings: A Discontinuity-Based Approach,” *Review of Economics and Statistics*, 91(4):717–724, 2009.
- Horowitz, J. L. and C. F. Manski, “Nonparametric Analysis of Randomized Experiments with Missing Covariate and Outcome Data,” *Journal of the American statistical Association*, 95(449):77–84, 2000.
- 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 C. F. Manski, “Confidence Intervals for Partially Identified Parameters,” *Econometrica*, 72(6):1845–1857, 2004.
- 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,” *American Economic Review*, pp. 604–620, 1986.

- 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., “Anatomy of the Selection Problem,” *Journal of Human Resources*, pp. 343–360, 1989.
- , “Nonparametric Bounds on Treatment Effects,” *American Economic Review*, 80(2):319–323, 1990.
- , “Monotone Treatment Response,” *Econometrica*, pp. 1311–1334, 1997.
- , *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.
- 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.
- Sharpsteen, C. and C. Bracken, *tikzDevice: R Graphics Output in LaTeX Format*, 2016, r package version 0.10-1.
- 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,” *BE Journal of Economic Analysis & Policy*, 8(1):1–55, 2008.
- Stoye, J., “More on Confidence Intervals for Partially Identified Parameters,” *Econometrica*, 77(4):1299–1315, 2009.
- Tamer, E., “Partial Identification in Econometrics,” *Annual Review of Economics*, 2(1):167–195, 2010.
- 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

If the ρ_ω are no longer assumed to be uniform within intended treatment group, we adapt eq. (10) to obtain the administrator's problem:

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

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

$$\int_{\kappa^*}^1 \rho_1(x)\Delta(x)dx \geq \int_{\kappa^*}^1 \rho_1(x)\chi dx = \bar{\rho}_1\chi(1 - \kappa^*) \Rightarrow \underbrace{\frac{\int_{\kappa^*}^1 \rho_1(x)\Delta(x)dx}{1 - \kappa^*}}_{\text{ITT}} \geq \bar{\rho}_1\chi = \bar{\rho}_1\Delta(\kappa^*), \quad (12)$$

where $\bar{\rho}_1 \equiv \frac{\int_{\kappa^*}^1 \rho_1(x)dx}{1 - \kappa^*}$ is the average attendance probability among the intended treated and the second equality follows because $\chi = \Delta(\kappa^*)$. Equation eq. (12) shows that if the average attendance probability among the intended treated ($\bar{\rho}_1$) were known then the RD estimate of the treatment effect could be used to provide a lower bound for the mean effect of intending to treat among the (intended) treated (ITT).

We can also bound the mean effect of intending to not treat the untreated (ITUT) by considering the second part of the sufficient conditions for this case, i.e., adapting eq. (4):

$$\int_0^{\kappa^*} \rho_0(x)\Delta(x)dx < \int_0^{\kappa^*} \rho_0(x)\chi dx = \bar{\rho}_0\chi\kappa^* \Rightarrow \underbrace{\frac{\int_0^{\kappa^*} \rho_0(x)\Delta(x)dx}{\kappa^*}}_{\text{ITUT}} < \bar{\rho}_0\chi = \bar{\rho}_0\Delta(\kappa^*), \quad (13)$$

where $\bar{\rho}_0 \equiv \frac{\int_0^{\kappa^*} \rho_0(x)dx}{\kappa^*}$ is the average attendance probability among the intended untreated and, as with the ITT above, the second equality follows because $\chi = \Delta(\kappa^*)$. Analogous to the case for the ITT, eq. (13) shows that if the average attendance probability among the intended untreated is known then the RD estimate of the treatment effect could be used to provide an upper bound for the mean effect of intending to not treat among the (intended) untreated (ITUT).

B Treatment Effect Uncertainty

Suppose the administrator is uncertain about the average treatment effect $\Delta(x)$ 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 of x , denote the administrator's noisy signal of the average treatment effect for students with index x . Because the administrator has unbiased beliefs about $\Delta(x)$ at every point x , it must be the case that $E[\epsilon_i] = 0$. The administrator chooses a cutoff to maximize her expected objective:

$$\begin{aligned} \max_{\tilde{\kappa}} E \left[\beta \left(\int_{\tilde{\kappa}}^1 \check{\Delta}(x) dx \right) - c(1 - \tilde{\kappa}) \right] &\Leftrightarrow \max_{\tilde{\kappa}} \beta E \left[\left(\int_{\tilde{\kappa}}^1 (\Delta(x) + \epsilon) dx \right) \right] - c(1 - \tilde{\kappa}) \\ &\Leftrightarrow \max_{\tilde{\kappa}} \beta \left(\int_{\tilde{\kappa}}^1 \Delta(x) dx \right) + \beta E[\epsilon] - c(1 - \tilde{\kappa}) \\ &\Leftrightarrow \max_{\tilde{\kappa}} \beta \left(\int_{\tilde{\kappa}}^1 \Delta(x) dx \right) - c(1 - \tilde{\kappa}). \end{aligned} \quad (14)$$

The first equivalence follows from the fact that the measure of students treated $(1 - \tilde{\kappa})$ 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, eq. (2). Therefore, the analysis of 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 $\delta(x) \equiv E[\epsilon_i|x]$, 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) $E[\delta(x)|x \in [\hat{\kappa}, \kappa^*]] < 0 < E[\delta(x)|x \geq \kappa^*]$, $\forall \hat{\kappa} \in [0, \kappa^*]$ and (ii) $E[\delta(x)|x \geq \kappa^*] < 0 < E[\delta(x)|x \in [\hat{\kappa}, \kappa^*]]$, $\forall \hat{\kappa} \in [0, \kappa^*]$. 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 Omitted Proofs from Section 3.2

Corollary 3a. *Suppose $\bar{\kappa}^* = 1 - \bar{\mu}$ solves eq. (2). The following ATUT bounds can be obtained:*

(i) *There is an informative upper bound, of $ATUT^{UB} = [\Delta(\bar{\kappa}^*)\kappa_a^* + \bar{\Delta}(\bar{\kappa}^* - \kappa_a^*)]/\bar{\kappa}^*$, if and only if κ_a^* is known.*

(ii) *There is an informative lower bound, of $ATUT^{LB} = \kappa_a^* \underline{\Delta}/\bar{\kappa}^*$, if and only if κ_a^* is known and $\underline{\Delta} < 0$.*

Proof. (i) Upper bound: By Proposition 2 the upper bound for the total treatment effect over $[0, \kappa_a^*]$ is $\chi\kappa_a^*$; if χ is unknown then apply Condition 1a to form the upper bound $\Delta(\bar{\kappa}^*)\kappa_a^*$. The administrator would have chosen to treat students $x \in [\kappa_a^*, \bar{\kappa}^*]$, meaning there would be no informative upper bound, resulting in an upper bound of $\bar{\Delta}(\bar{\kappa}^* - \kappa_a^*)$ for the total

treatment effect. Sum these and divide by the measure of untreated students to obtain $ATUT^{UB} = [\Delta(\bar{\kappa}^*)\kappa_a^* + \bar{\Delta}(\bar{\kappa}^* - \kappa_a^*)]/\bar{\kappa}^*$.

(ii) Lower bound: As in the unconstrained problem, the interval below κ_a^* has an uninformative lower bound. Adapting the first part of Condition 2 for the interval $[\kappa_a^*, \bar{\kappa}^*]$, we obtain $\chi(\bar{\kappa}^* - \kappa_a^*) < \int_{\kappa_a^*}^{\bar{\kappa}^*} \Delta(x)dx$; if χ is unknown then by using the reasoning in Corollary $\bar{2}a$, $0 < \int_{\kappa_a^*}^{\bar{\kappa}^*} \Delta(x)dx$. Sum over these intervals and divide by the measure of untreated students to obtain $ATUT^{LB} = \kappa_a^*\underline{\Delta}/\bar{\kappa}^*$.

For either upper or lower bound, if κ_a^* was unknown then one would have to adopt the worst-case scenario corresponding to the loosest bounds, setting $\kappa_a^* = 0$ for the upper bound and $\kappa_a^* = \bar{\kappa}^*$ for the lower bound, resulting in uninformative bounds. \square

Corollary $\bar{3}b$. *Suppose $\bar{\kappa}^* = \kappa_b^*$ solves eq. $(\bar{2})$. The following ATUT bounds can be obtained:*

(i) *There is an informative upper bound, which is tighter if κ_a^* is known. If κ_a^* is unknown the upper bound is $ATUT^{UB} = [\Delta(\bar{\kappa}^*)(\bar{\kappa}^* - (1 - \bar{\mu})) + \bar{\Delta}(1 - \bar{\mu})]/\bar{\kappa}^*$; if κ_a^* is known the upper bound is $ATUT^{UB} = [\Delta(\bar{\kappa}^*)(\bar{\kappa}^* - ((1 - \bar{\mu}) - \kappa_a^*)) + \bar{\Delta}((1 - \bar{\mu}) - \kappa_a^*)]/\bar{\kappa}^*$.*

(ii) *There is an informative lower bound, which is tighter if κ_a^* is known. If κ_a^* is unknown the lower bound is $ATUT^{LB} = [\Delta(\bar{\kappa}^*)(\bar{\kappa}^* - (1 - \bar{\mu})) + \underline{\Delta}(1 - \bar{\mu})]/\bar{\kappa}^*$; if κ_a^* is known the lower bound is $ATUT^{LB} = [\Delta(\bar{\kappa}^*)(\bar{\kappa}^* - \kappa_a^*) + \underline{\Delta}\kappa_a^*]/\bar{\kappa}^*$.*

Proof. (i) Upper bound: Note that eq. $(\bar{4}b)$ implies that $\int_{1-\bar{\mu}}^{\bar{\kappa}^*} \Delta(x)dx < \Delta(\bar{\kappa}^*)(\bar{\kappa}^* - (1 - \bar{\mu}))$. By Assumption 1(iii), the average treatment effect for any treated unit cannot exceed $\bar{\Delta}$, which if κ_a^* is not known is the upper bound for the average treatment effect for $x \in [0, 1 - \bar{\mu}]$. This results in an upper bound of $ATUT^{UB} = [\Delta(\bar{\kappa}^*)(\bar{\kappa}^* - (1 - \bar{\mu})) + \bar{\Delta}(1 - \bar{\mu})]/\bar{\kappa}^*$. If κ_a^* is known, then $\int_0^{\kappa_a^*} \Delta(x)dx < \Delta(\bar{\kappa}^*)\kappa_a^*$ by eq. (4), which results in a tighter upper bound of $ATUT^{UB} = [\Delta(\bar{\kappa}^*)(\bar{\kappa}^* - ((1 - \bar{\mu}) - \kappa_a^*)) + \bar{\Delta}((1 - \bar{\mu}) - \kappa_a^*)]/\bar{\kappa}^*$.

(ii) Lower bound: Suppose that κ_a^* is known. Use the same argument as for the lower bound in Corollary $\bar{3}a$, but also include the interval $[1 - \bar{\mu}, \bar{\kappa}^*]$, which has a total lower bound of $\underline{\Delta}(\bar{\kappa}^* - (1 - \bar{\mu}))$, and note that in Case \bar{b} we have $\Delta(\bar{\kappa}^*) = \chi$. This would result in a ATUT lower bound $[\Delta(\bar{\kappa}^*)((1 - \bar{\mu}) - \kappa_a^*) + \underline{\Delta}(\kappa_a^* + (\bar{\kappa}^* - (1 - \bar{\mu})))]/\bar{\kappa}^*$. However, this expression does not exploit the fact that we are in a Case- \bar{b} -binding constrained scenario. By adapting Condition 2 to take as the upper limit $\bar{\kappa}^*$ (instead of 1), this latter piece of information implies that $\int_{\kappa_a^*}^{\bar{\kappa}^*} \Delta(x)dx \geq \chi(\bar{\kappa}^* - \kappa_a^*) = \Delta(\bar{\kappa}^*)(\bar{\kappa}^* - \kappa_a^*)$; intuitively, the unconstrained administrator would have treated $x \in [\kappa_a^*, \bar{\kappa}^*]$. Using (only) this information, the ATUT lower bound would be $[\Delta(\bar{\kappa}^*)(\bar{\kappa}^* - \kappa_a^*) + \underline{\Delta}\kappa_a^*]/\bar{\kappa}^*$. Taking the maximum (i.e., tightest lower bound) results in $ATUT^{LB} = [\Delta(\bar{\kappa}^*)(\bar{\kappa}^* - \kappa_a^*) + \underline{\Delta}\kappa_a^*]/\bar{\kappa}^*$. If κ_a^* is not known, then we must

choose κ_a^* to minimize the lower bound, i.e., $\kappa_a^* = (1 - \bar{\mu})$, which would result in a looser bound of $\text{ATUT}^{LB} = [\Delta(\bar{\kappa}^*)(\bar{\kappa}^* - (1 - \bar{\mu})) + \underline{\Delta}(1 - \bar{\mu})]/\bar{\kappa}^*$. \square

Corollary 4a. *Suppose $\bar{\kappa}^* = 1 - \bar{\mu}$ solves eq. (2). The following ATE bounds can be obtained.*

(i) *If $\underline{\Delta} < 0$ and κ_a^* is unknown then there is an informative lower bound, of $\text{ATE}^{LB} = \underline{\Delta}\bar{\kappa}^*$, which is tightened to $\text{ATE}^{LB} = \underline{\Delta}\kappa_a^*$ if κ_a^* is known.*

(ii) *There is an informative upper bound, of $\text{ATE}^{UB} = \Delta(\bar{\kappa}^*)\kappa_a^* + \bar{\Delta}(1 - \kappa_a^*)$, if and only if κ_a^* is known.*

Proof. (i) Lower bound: If κ_a^* is unknown, by Corollary 2a, the lower bound on the ATE for units with $x \geq \bar{\kappa}^*$ is 0, tightening the ATE lower bound from $\underline{\Delta}$ to $\underline{\Delta}\bar{\kappa}^*$. If κ_a^* is known, then by Corollary 3a(ii) the lower bound on for units $x \in [\kappa_a^*, \bar{\kappa}^*]$ is 0, further tightening the ATE lower bound to $\underline{\Delta}\kappa_a^*$.

(ii) Upper bound: If κ_a^* is known, then by Corollary 3a(i) the upper bound on the total gain for the untreated is $\Delta(\bar{\kappa}^*)\kappa_a^* + \bar{\Delta}(\bar{\kappa}^* - \kappa_a^*)$, reducing the ATE upper bound from $\bar{\Delta}$ to $\Delta(\bar{\kappa}^*)\kappa_a^* + \bar{\Delta}(1 - \kappa_a^*)$. \square

Corollary 4b. *Suppose $\bar{\kappa}^* = \kappa_b^*$ solves eq. (2). The following ATE bounds can be obtained.*

(i) *If κ_a^* is unknown there is an informative lower bound, of $\text{ATE}^{LB} = \Delta(\bar{\kappa}^*)(1 - (1 - \bar{\mu})) + \underline{\Delta}(1 - \bar{\mu})$, which is tightened to $\text{ATE}^{LB} = \Delta(\bar{\kappa}^*)(1 - \kappa_a^*) + \underline{\Delta}\kappa_a^*$, if κ_a^* is known.*

(ii) *If κ_a^* is unknown there is an informative upper bound, of $\text{ATE}^{UB} = \Delta(\bar{\kappa}^*)(\bar{\kappa}^* - (1 - \bar{\mu})) + \bar{\Delta}(1 - (\bar{\kappa}^* - (1 - \bar{\mu})))$, which is tightened to $\text{ATE}^{UB} = \Delta(\bar{\kappa}^*)(\bar{\kappa}^* - ((1 - \bar{\mu}) - \kappa_a^*)) + \bar{\Delta}(1 - (\bar{\kappa}^* - ((1 - \bar{\mu}) - \kappa_a^*)))$ if κ_a^* is known.*

Proof. For each bound, first assume κ_a^* is not known.

(i) Lower bound: Measure $\bar{\kappa}^*$ units are untreated, and Corollary 3b(ii) shows that the ATUT is no smaller than untreated of $[\Delta(\bar{\kappa}^*)(\bar{\kappa}^* - (1 - \bar{\mu})) + \underline{\Delta}(1 - \bar{\mu})]/\bar{\kappa}^*$. Analogously, $1 - \bar{\kappa}^*$ units are treated, and eq. (3b) implies that the ATT is no smaller than $\Delta(\bar{\kappa}^*)$. Integrate and sum the two parts to form $\text{ATE}^{LB} = \Delta(\bar{\kappa}^*)(1 - (1 - \bar{\mu})) + \underline{\Delta}(1 - \bar{\mu})$. If κ_a^* is known, we can use the tighter lower bound on the ATUT from Corollary 3b(ii), resulting in a tighter lower bound of $\text{ATE}^{LB} = \Delta(\bar{\kappa}^*)(1 - \kappa_a^*) + \underline{\Delta}\kappa_a^*$.

(ii) Upper bound: Use the expression for the upper bound of ATUT from Corollary 3b(i) and the fact that the upper bound on the treated students, with $x \in [\bar{\kappa}^*, 1]$, is $\bar{\Delta}$ (by Assumption 1(iii)), to integrate and sum to form $\text{ATE}^{UB} = \Delta(\bar{\kappa}^*)(\bar{\kappa}^* - (1 - \bar{\mu})) + \bar{\Delta}(1 - (\bar{\kappa}^* - (1 - \bar{\mu})))$. If κ_a^* is known, then, similar to Corollary 3b(i), we can shift the measure κ_a^* from having an upper bound of $\bar{\Delta}$ to $\Delta(\bar{\kappa}^*)$, resulting in a tighter upper bound of $\text{ATE}^{UB} = \Delta(\bar{\kappa}^*)(\bar{\kappa}^* - ((1 - \bar{\mu}) - \kappa_a^*)) + \bar{\Delta}(1 - (\bar{\kappa}^* - ((1 - \bar{\mu}) - \kappa_a^*)))$. \square

D Weighted Objective

The administrator's original problem (2) 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 quantiles of 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 eq. (2) to allow the administrator to weigh gains for students depending on their index x by using weights $\phi(x)$, where $\phi' \leq 0$:

$$\max_{\tilde{\kappa}} \left(\int_{\tilde{\kappa}}^1 \phi(x)\Delta(x)dx \right) - \chi(1 - \tilde{\kappa}), \quad (\hat{2})$$

and proceed with the analysis.

Condition $\hat{1}$ (Necessity). *For problem $(\hat{2})$, the following necessary conditions must hold for κ^* :*

(i) *Marginal Benefit=Marginal Cost:* $\phi(\kappa^*)\Delta(\kappa^*) = \chi$

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

Proof. Differentiate the administrator's problem $(\hat{2})$ 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 $\phi' < 0$. \square

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.:*

$$\text{Participation: } \int_{\kappa^*}^1 \phi(x)\Delta(x)dx \geq \chi(1 - \kappa^*). \quad (\hat{3})$$

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

$$\int_{\hat{\kappa}}^{\kappa^*} \phi(x)\Delta(x)dx < \chi(\kappa^* - \hat{\kappa}). \quad (\hat{4})$$

Proposition 1 remains true when $\phi' \leq 0$. To see this, divide eq. (3) by the measure of treated students and combine with Condition 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 eq. (2). In summary, all of the bounds from the unweighted problem, including Corollary 4, which bounds the ATE, are also obtained for the weighted problem (2).

E 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 effect parameters of interest.

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

- (i) *Marginal Benefit=Marginal Cost: $\Delta(\kappa^*) = c'(1 - \kappa^*)$ for any cost function $c(\cdot)$ satisfying Assumption 1'*
- (ii) *Increasing Marginal Benefit: $\Delta'(\kappa^*) \geq 0$ if the marginal cost is constant or decreasing; this inequality is strict if the marginal cost is decreasing.*

Proof. Differentiate the administrator's problem (2) 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 eq. (2) 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.:*

$$\text{Participation: } \int_{\kappa^*}^1 \Delta(x)dx \geq c(1 - \kappa^*). \quad (3')$$

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

$$\int_{\hat{\kappa}}^{\kappa^*} \Delta(x)dx < c(1 - \hat{\kappa}) - c(1 - \kappa^*). \quad (4')$$

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 eq. (3') 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)), eq. (3') implies that

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

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

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

Although Corollary 2' provides a lower bound for the average effect of treatment on the treated, there is no informative (i.e., lower than $\bar{\Delta}$) 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 eq. (3') and combine with this with Condition 1'(i) to obtain

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

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 an average 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 $x < \kappa^*$.⁴⁶ 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 $0 \leq a < b \leq \kappa^*$ if $\underline{\Delta} > \frac{c(1-a)-c(1-\kappa^*)-\bar{\Delta}(b-a)}{\kappa^*-b}$.*

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 eq. (4') into two parts at \hat{x} and rearrange terms:

$$\int_{\hat{x}-\hat{\mu}}^{\hat{x}} \Delta(x)dx < c(1-(\hat{x}-\hat{\mu}))-c(1-\kappa^*) - \int_{\hat{x}}^{\kappa^*} \Delta(x)dx \Rightarrow \int_{\hat{x}-\hat{\mu}}^{\hat{x}} \Delta(x)dx < c(1-(\hat{x}-\hat{\mu}))-c(1-\kappa^*) - \underline{\Delta}(\kappa^* - \hat{x}), \quad (5')$$

where the implication follows from Assumption 1(iii).⁴⁷ □

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.*

Proof. Let $\hat{x} = \hat{\mu} = \kappa^*$ in eq. (5') and divide through by κ^* to obtain the first result:

$$\underbrace{\int_0^{\kappa^*} \frac{\Delta(x)}{\kappa^*} dx}_{\text{ATUT}} < \underbrace{\frac{c(1) - c(1 - \kappa^*)}{\kappa^*}}_{>0, <\infty}, \quad (8')$$

⁴⁶Note 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.

⁴⁷As with Proposition 2, this bound will be informative for all but very low values of $\underline{\Delta}$. This condition is satisfied when using Proposition 2' to bound the ATUT or ATE if the average cost of treating the untreated is less than $\bar{\Delta}$.

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). \square

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 $\text{ATUT} < \Delta(\kappa^*) \leq \text{ATT}$. 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.


E.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 $\text{ATE}^{\text{LB}} = \underline{\Delta}\kappa^* + \chi(1 - \kappa^*)$ and the ATE upper bound is $\text{ATE}^{\text{UB}} = \chi\kappa^* + \bar{\Delta}(1 - \kappa^*)$, 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).⁴⁸ This rotation implies the ATT must be higher than the case corresponding to the constant marginal cost in order to satisfy eq. (3'). 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 eq. (4'). The opposite holds true for when we rotate the cost curve

⁴⁸Recall that this line is decreasing in x because the treatment is being extended from $x = 1$ downwards.

clockwise about the point $(\kappa^*, \Delta(\kappa^*))$, to reflect an increasing marginal cost of treatment (dot-dashed red line). Table 6 summarizes these results, showing that when the marginal cost of treatment is decreasing, bounds on the population ATE are tighter than they would be with a constant marginal cost, while when marginal cost is increasing, bounds on the population ATE are looser.



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

Table 6: Summary of bounds on population ATE

Figure 5: Example with different cost functions