

## RESEARCH ARTICLE

# Regression discontinuity threshold optimization

Ioana Marinescu<sup>1\*</sup>, Sofia Triantafillou<sup>2</sup>, Konrad Kording<sup>3</sup><sup>1</sup> School of Social Policy & Practice, University of Pennsylvania, Philadelphia, PA, United States of America,<sup>2</sup> Department of Mathematics and Applied Mathematics, University of Crete, Heraklion, Crete, Greece,<sup>3</sup> Departments of Bioengineering and Neuroscience, University of Pennsylvania, Philadelphia, PA, United States of America\* [ioma@upenn.edu](mailto:ioma@upenn.edu)

## Abstract

Treatments often come with thresholds, e.g. we are given statins if our cholesterol is above a certain threshold. But which statin administration threshold maximizes our quality of life adjusted years? More generally, which threshold would optimize the average expected outcome? Regression discontinuity approaches are used to measure the local average treatment effect (LATE) and more recently also the Marginal Threshold Treatment Effect (MTTE), which shows how marginal changes in the threshold can affect the LATE. We extend this idea to define the problem of optimizing a policy threshold, i.e. selecting a threshold that optimizes the cumulative effect of the treatment on the treated. We present an estimator of the optimal threshold based on a constrained optimization framework. We show how to use machine learning (Gaussian process regression) for non-linear estimation. We also extend the estimation to a conservative threshold that is unlikely to produce harm, and we show how to include policy cost constraints. We apply these results to estimate an optimal tip-maximizing threshold for tip suggestions in taxi cabs Haggag (2014).

## OPEN ACCESS

**Citation:** Marinescu I, Triantafillou S, Kording K (2022) Regression discontinuity threshold optimization. PLoS ONE 17(11): e0276755. <https://doi.org/10.1371/journal.pone.0276755>

**Editor:** Baogui Xin, Shandong University of Science and Technology, CHINA

**Received:** April 29, 2022

**Accepted:** October 12, 2022

**Published:** November 16, 2022

**Copyright:** © 2022 Marinescu et al. This is an open access article distributed under the terms of the [Creative Commons Attribution License](https://creativecommons.org/licenses/by/4.0/), which permits unrestricted use, distribution, and reproduction in any medium, provided the original author and source are credited.

**Data Availability Statement:** The data underlying the results presented in the study are available from <https://www.aeaweb.org/articles?id=10.1257/app.6.3.1> and have been used in many previous papers.

**Funding:** There was no explicit funding for this work.

**Competing interests:** The authors have declared that no competing interests exist.

## 1 Introduction

Many real world optimization problems come in the form of a set of thresholds that need to be chosen. We give scholarships and access to better schools to students whose grades exceed certain thresholds [1], give drugs to patients whose lab tests exceed given values [2–4], give loans to people whose credit scores exceed some value [5], and show ads to unsuspecting web users whose likelihood to click on our ad exceeds a certain value [6]. In such cases, the decision maker faces an optimization problem: which level of the threshold will maximize the relevant welfare criterion? Threshold optimization has received little attention in the econometrics literature. We focus on Regression discontinuity (RD) designs where a threshold sharply defines treatment assignment.

RD designs are quasi-experimental designs that allow for local effect estimation [1]. In an RD design, an intervention is applied based on a cut-off on a running variable: units with a running variable value below the cutoff are assigned to the control condition, and units with a running variable value above the cutoff are assigned to the treatment condition (or vice versa).

Assuming that the average properties of units just above and just below the threshold are similar, this allows estimation of the effect of the intervention at the threshold (Local Average Treatment Effect, LATE). These approaches are often implemented using local polynomial regression [7] and rely on continuity assumptions near the threshold (see [8] for a review of assumptions and methods). Alternative approaches are based on viewing RD design as a local randomization procedure [9, 10]. Regardless of the approach, the validity of this estimation weakens as we move to units away from the threshold, and the treatment effect estimation is only valid for a very specific subpopulation with running variable values very close to the threshold.

Recent works have gone beyond LATE estimation, trying to extend the validity of the effect estimation away from the threshold. [11] use difference-in-difference type assumptions to estimate the effect away from the threshold. [12] assume that the running variable is one of multiple noisy measures of a latent factor, and show that this allows for non-parametric estimation of the average effect away from the threshold. Angrist and Rokkanen [13] rely on the existence of a set of covariates that break the dependence of the policy on the running variable, and allow for extrapolation away from the threshold. [14] develop external validity tests based on the independence between types of compliance and potential outcomes. [15, 16] use RD policies with multiple cut-offs to estimate causal effects for different levels of the score variable.

Works mentioned so far try to extend the validity of a LATE estimate beyond the threshold of the RD design. [17] also try to estimate how the LATE changes for small changes of the threshold; hence they evaluate the effect of marginally moving the policy threshold. This is called the Marginal Threshold Treatment effect (MTTE), and its estimation relies on some smoothness assumptions on the threshold. Building on the idea of estimating the effect of changing the threshold, we propose *optimizing* the threshold to maximize a welfare criterion (e.g. expected earnings over the population). For optimization, we need to make assumptions about the behavior of potential outcomes away from the threshold: first, we assume that linear extrapolation of potential outcomes is valid in a neighborhood around the threshold. Violations of this assumption introduce biases, and a highly relevant literature discusses related issues [18–20]. Our assumption is much stronger than continuity assumptions near the threshold; however, similar assumptions are often already necessary for the RD design to have statistical power. Second, we also assume that changing the threshold does not affect the treatment effect conditional on the running variable. Under these assumptions, we can extrapolate potential outcomes away from the threshold, and estimate the optimal threshold. While policy optimization has been explored for observational causal inference (e.g., [21, 22]), to the best of our knowledge, threshold optimization for RD designs has not been previously explored.

Our work contributes to the literature in the following ways: we introduce the welfare function in RD designs as an objective function for threshold optimization. We want to emphasize here that we can design meaningful welfare functions, primarily by assuming that we care about every individual being treated equally. We effectively want to maximize the population density weighted individual welfare. We derive estimates for optimal thresholds using local linear regression and Gaussian process regression. We obtain strategies to improve the threshold, including strategies that limit the risk of a negative local average treatment effect at the new threshold. These techniques could be a crucial piece for policy improvement. Finally, we apply the methods to estimate the tip-maximizing threshold for tip suggestions in taxi cabs [23].

Importantly, our calculations allow us to conceptualize RDD in new ways. We show that finding a zero treatment effect (on the relevant outcome that we want to maximize) in an RDD does not imply that the policy is ineffective but can indicate that the threshold is set optimally if the MTTE also changes at the threshold: in the canonical case of a positive treatment effect above the threshold, the LATE is zero for an optimally set threshold, and the treatment effect

derivative is larger to the right of the optimal threshold than to the left of the optimal threshold. Many real-world thresholds are set by experts at levels that are likely to be quite adequate, so a zero treatment effect is a plausible outcome.

The rest of the paper is organized as follows: Section 2 briefly discusses the framework of sharp regression discontinuity designs, and Section 3 discusses recent results in the literature on the derivatives of LATEs. Our contribution begins with section 4, which defines the problem of optimizing policy thresholds in sharp and fuzzy RDDs and proposes methods of threshold optimization with local linear regression (Section 4.1) and Gaussian process regression (Section 4.2). Section 5 addresses risk aversion by discussing a modification of the threshold optimization problem so that the new threshold is unlikely to yield a negative LATE. In Section 7, threshold optimization is used to propose a new threshold for default tip suggestion in cab fares, that would maximize tip percentage. In Section 8 we summarize results and discuss pitfalls and future work.

## 2 Regression discontinuity designs

We first briefly discuss the framework of RD designs and how they can be used to obtain LATE estimates. For each unit  $i$  there are two potential outcomes,  $Y_i(1)$  and  $Y_i(0)$ , corresponding to treatment and control, respectively. Let  $T_i$  denote the assignment for unit  $i$ , i.e.  $T_i = 1$  if unit  $i$  is treated and 0 otherwise. In sharp RD designs, the assignment variable  $T$  is a deterministic piecewise function of a running variable  $X$ :  $T_i = I(X_i \geq c)$ . For each unit, only one of the outcomes is observed (they are either treated or control):

$$y_i = T_i Y_i(0) + (1 - T_i) Y_i(1).$$

Let  $\mu_1(x) := E[Y(1)|X = x]$  and  $\mu_0(x) := E[Y(0)|X = x]$  be the conditional expectation of the outcome for treatment and control, respectively. The following assumption is necessary for identifying the LATE:

**Assumption 1** (Continuity).  $\mu_{1,0}$  are continuous in a small neighborhood around  $c$ .

Under the monotonicity of the conditional Average Treatment Effects (ATEs) around  $c$ , the RD design allows the identification of  $\pi(c) = E[Y(1) - Y(0)|X = c]$ , as follows:

$$\pi(c) = \lim_{x \rightarrow c^+} E[Y_i|X_i = x] - \lim_{x \rightarrow c^-} E[Y_i|X_i = x] = \lim_{x \rightarrow c^+} \mu_1(c) - \lim_{x \rightarrow c^-} \mu_0(c) = \mu_1(c) - \mu_0(c).$$

Thus, the LATE  $\pi(c)$  is identifiable using a parametric or non-parametric estimator for  $\mu_{1,0}$ , as  $\mu_1(c) - \mu_0(c)$ . Treatment assignment arbitrarily close to the threshold is considered random enough so that treated units right above the threshold are valid counterfactuals for non-treated units just below the threshold.

Typically,  $\mu_{1,0}$  are estimated by applying a locally linear regression (LLR) model on data in a bandwidth  $h$  around  $c$ . The choice of the bandwidth is important, as an overly broad bandwidth will induce biases when the relationship between running variable and outcome are nonlinear. In practice, researchers may try different thresholds and test how the chosen bandwidth affects the obtained estimates [8]. For example, for LLR estimates, [24] propose a method for selecting an asymptotically optimal bandwidth in terms of the bias in the estimated LATE. Ideally, the chosen bandwidth must be narrow enough for the observations to be close enough to the cutoff to avoid large biases, and wide enough to provide adequate power.

Fuzzy RD designs (FRDD) handle settings where the threshold affects treatment assignment, but not deterministically. Let  $T_i$  be a binary variable indicating whether unit  $i$  has received treatment, and  $T^* = I(X \geq c)$ . As with the potential outcomes  $Y(t)$ , the potential treatment status  $T(t^*)$ , indicates what an individual's treatment status would be if  $T^* = t^*$ . Let  $f(x) := P(T_i|X_i = x)$  be the probability that unit  $i$  with running variable  $x$  is treated. The probability of

treatment changes discontinuously at the threshold, albeit not deterministically:  $\lim_{x \rightarrow c^-} f(x) \neq \lim_{x \rightarrow c^+} f(x)$ ,  $0 < f(x) < 1$ .

For every individual, we only observe one of the potential treatments  $T(0)$ ,  $T(1)$ . Individuals can be classified based on their behavior towards treatment as: (a) compliers:  $T(0) < T(1)$ , (b) defiers:  $T(0) > T(1)$  (c) always-takers:  $T(0) = T(1) = 1$ , and (d) never-takers:  $T(0) = T(1) = 0$ . LATE at the threshold  $c$  is the effect of treatment on compliers:  $\pi_f(c) = E[Y(1) - Y(0)|X = c, T(0) < T(1)]$ . LATE identification then requires that, moving the threshold from  $c$  to  $c + \epsilon$  will not create any defiers, and that there is at least one complier:

**Assumption 2** (Strong Monotonicity).  $T(0) \leq T(1)$ .  $P(T(0) > T(1)|X = x)$  is strictly positive at  $X = c$ .

[25] show that under the additional assumption of local independence, the running variable can be viewed as an instrumental variable, that affects the outcome only through treatment assignment at the threshold.

**Assumption 3** (Local Independence).  $Y(0) - Y(1)$  is independent of  $T$  in a neighborhood of  $X = c$ .

Under Assumptions 1, 2 and 3, the LATE at  $X = c$  can be identified as

$$\begin{aligned} \pi_f(c) &= \frac{\lim_{x \rightarrow c^+} E[Y_i|X_i = x] - \lim_{x \rightarrow c^-} E[Y_i|X_i = x]}{\lim_{x \rightarrow c^+} E[T_i|X_i = x] - \lim_{x \rightarrow c^-} E[T_i|X_i = x]} \\ &= \frac{\lim_{x \rightarrow c^+} \mu_1(x) - \lim_{x \rightarrow c^-} \mu_0(x)}{\lim_{x \rightarrow c^+} f(x) - \lim_{x \rightarrow c^-} f(x)} \\ &= \frac{\mu_1(c) - \mu_0(c)}{\lim_{x \rightarrow c^+} f(x) - \lim_{x \rightarrow c^-} f(x)} \end{aligned} \tag{2.1}$$

The local independence assumption implies that the treatment effect does not depend on the running variable directly or indirectly through confounding. This is often problematic in empirical applications, and does not allow the exploitation of derivative changes near the threshold employed in some recent work [18]. [26] shows that Eq (2.1) holds if we replace the local independence assumption with a continuity of selection into a type of individual:

**Assumption 4** (Local Smoothness). The conditional means  $E[Y(t)|T(0) < T(1), X = x]$  and  $E[Y(t)|T(0) = T(1) = t, X = x]$ , as well as the probabilities  $P(T(0) < T(1)|X = x)$ ,  $P(T(0) = T(1)|X = x)$ ,  $t = 0, 1$  are continuous in  $x$  in a neighborhood of  $x = c$ .

Assumption 4 can replace Assumptions 1 and 3 to obtain Eq (2.1). See [26] for a discussion on the behavioral assumptions that can lead to each of these assumptions. Importantly, removing the local independence assumption allows us to exploit changes of slope near the threshold.

### 3 Effect derivatives in RD designs

RD designs allow the unbiased estimation of the LATE at the predefined threshold under weak assumptions, and therefore have strong internal validity. Researchers are often interested in the external validity of the LATE and the effects of marginal threshold changes in the LATE. [17] show how to estimate the derivatives of the LATE under similarly weak conditions.

Specifically, let  $S(x, c) = E[Y(1) - Y(0)|X = x, C = c]$  be the average treatment effect of the sharp RD policy applied at threshold  $c$  on individuals with running variable  $x$ . Thus, the effect depends both on the value of the running variable  $X$  and on the implemented threshold  $C$ . For a constant threshold  $C = c$ ,  $\pi(X) = S(X, c)$  shows how the treatment effect changes for different values of the running variable  $X$ . Respectively,  $\tau(C) = S(C, C)$  shows how the LATE changes for different thresholds  $C$ . Notice that in the previous section,  $\pi(x)$  and  $\pi_f(x)$  implicitly condition

on  $C = c$ . In this section, we view treatment outcome as a function of both the value of the running variable and the value of the threshold.

[17] define the Treatment Effect Derivative (TED) as  $\pi'(c) = \frac{\partial S(X,c)}{\partial X} \Big|_{X=c}$ , and the Marginal Threshold Treatment Effect (MTTE) as  $\tau'(c) = \frac{\partial S(C,C)}{\partial C} \Big|_{C=c}$ . Thus, TED shows how the LATE varies for marginal changes in the running variable (and a given threshold), and the MTTE shows how the LATE varies for small changes in the threshold.

For fuzzy designs, let  $S_f(x, c) = E[Y(1) - Y(0)|X = c, C = c, T(0) < T(1)]$  be the average treatment effect of the fuzzy RD design policy applied at threshold  $c$  on complier individuals with running variable  $x$ . TED and MTTE are similarly defined as  $\pi'_f(c) = \frac{\partial S_f(X,c)}{\partial X} \Big|_{X=c}$ , and the Marginal Threshold Treatment Effect (MTTE) as  $\tau'_f(c) = \frac{\partial S_f(C,C)}{\partial C} \Big|_{C=c}$ .

The authors provide conditions under which the TED and MTTE are identifiable. Specifically, non-parametric identifiability of TED for sharp RD designs is possible under the following assumption:

**Assumption 5** (Sharp RD TED).  $E[Y(t)|X = x]$  is continuously differentiable in  $x$  in a neighborhood of  $x = c$  for  $t = 0, 1$ .

Assumption 5 differs from Assumption 1 required for LATE identification in that it requires continuous differentiability at  $x = c$ , instead of continuity. However, most LATE estimators (parametric or non-parametric) require continuous differentiability at  $x = c$ . Under Assumption 1,

$$\pi'(c) = \mu'_1(x) - \mu'_0(x). \tag{3.1}$$

For fuzzy RD designs, the function  $\pi_f(x)$  additionally conditions on compliers, i.e.  $T(0) < T(1)$ . The corresponding assumption for TED identifiability is the following:

**Assumption 6** (Fuzzy RD TED). The conditional means  $E[Y(t)|T(0) < T(1), X = x]$  and  $E[Y(t)|T(0) = T(1) = t, X = x]$ , as well as the probabilities  $P(T(0) < T(1)|X = x)$ ,  $P(T(0) = T(1)|X = x)$ ,  $t = 0, 1$  are continuously differentiable in  $x$  in a neighborhood of  $x = c$ .

Again, this assumption requires continuous differentiability at  $x = c$ , instead of just continuity (Assumption 4), which is required for LATE estimation. However, while not explicitly required for identification, the majority of parametric and non-parametric models used for estimation are continuously differentiable. Let  $p_f(c)$  be the ratio of compliers at  $X = c$ . Under Assumptions 2 and 6, the TED for fuzzy RD designs is:

$$\pi'_f(c) = \frac{\mu'_1(x) - \mu'_0(x)}{p_f(c)} - \frac{p'_f(c)\pi_f(c)}{p_f(c)}. \tag{3.2}$$

Eq (3.1) can be viewed as a special case of Eq (3.2), where  $p(c) = 1$  and  $p'(c) = 0$ . Eqs (3.1) and (3.2) can be used to identify the treatment effect derivative, which indicates how the effect would change for individuals with a slightly different value of the running variable. A large TED brings into question the external validity of the LATE estimator, since it implies that a small change in the running variable could have a large impact on the treatment effect.

TED captures the sensitivity of the LATE estimator to small changes in the running variable for a fixed threshold. MTTE answers the question: How would the LATE change for a small change in the implemented threshold? Naturally, since all of our samples come from a single RD design with a fixed threshold, additional assumptions are required to estimate MTTE: One has to make an assumption about how changing the threshold would affect potential outcomes. One common assumption is the **policy invariance assumption** [27], which states that *an individual's outcome only depends on the treatment assigned to the agent, and not*

independently on the policy. We use a weaker version of the policy invariance assumption, assuming mean independence instead of independence:

**Assumption 7** (Policy Invariance). *The potential outcomes for both treated and untreated are mean-independent of the threshold given the running variable:*

$$E[Y_i(t)|X_i, C] = E[Y_i(t)|X_i] \quad t = 0, 1.$$

This assumption may be unrealistic in many policies, in particular if the threshold is moved a lot. For example, if the threshold test score for acceptance at an elite school becomes lower, this may affect the treatment effect of admission for higher test score students due to peer effects. Additional examples are discussed in [27]. However, there are domains where the policy invariance assumption is more plausible: for example, the treatment effect of receiving a statin for patients with high cholesterol is not likely to change if patients with lower cholesterol also take the same statin. Moreover, to allow MTTE estimation, the policy invariance assumption only needs to hold *locally* at the threshold  $c$ .

Under (the local version of) Assumption 7, the MTTE is equal to the TED for sharp RD designs. Specifically,  $\tau'(c) = \pi'(c) + \frac{\partial S(X,C)}{\partial C} \Big|_{X=c, C=c}$ . Local policy invariance implies  $\frac{\partial S(X,C)}{\partial C} \Big|_{X=c, C=c} = 0$ , hence,  $\tau'(c) = \pi'(c)$ . Similarly, for fuzzy designs  $\tau'_f(c) = \pi'_f(c) + \frac{\partial S_f(X,C)}{\partial C} \Big|_{X=c, C=c}$ , and  $\frac{\partial S_f(X,C)}{\partial C} \Big|_{X=c, C=c} = 0$ , thus  $\tau'_f(c) = \pi'_f(c)$ . Thus, under local policy invariance, we can estimate how marginal changes of the threshold would affect the LATE at the (new) threshold.

### 4 Threshold optimization

The sign of the MTTE shows whether a marginal increase (decrease) of the threshold is likely to increase or decrease of the corresponding LATE. In many applications, we are interested in *optimizing* this threshold, i.e., we want to identify the threshold that maximizes the benefit of the interventions on the treated population. Let  $p(c)$  be the density of observations (individuals) at  $x$ , and that  $x$  ranges from 0 to  $x_{max}$ . For sharp designs, we define the *welfare function* as the cumulative average treatment effect over all treated units:

$$W(t) = \int_t^{x_{max}} S(x, t)p(x)dx,$$

where  $t$  is the threshold at which the intervention is applied. Following the notation in Section 3,  $S(x, t)$  denotes the average treatment effect of individuals with a running variable of  $x$  when the policy is implemented at threshold  $t$ .

Without loss of generality, we assume the desired treatment effect is positive, hence we want to find the threshold that maximizes  $W(t)$ . Assuming there are no spillover effects from the treated to the untreated, maximizing the welfare function corresponds to finding the threshold  $c^*$  such that:

$$c^* = \arg \max_t \int_t^{x_{max}} S(x, t)p(x)dx. \tag{4.1}$$

Under the policy invariance assumption,  $S(x, t) = \pi(x) \forall t$ , therefore the problem in Eq (4.1) becomes:

$$c^* = \arg \max_t \int_t^{x_{max}} \pi(x)p(x)dx. \tag{4.2}$$

Assuming that  $\pi(x), p(x)$  are continuous, and that  $p(x)$  is strictly positive in  $x$ , the first derivative of the welfare function is  $W'(x) = -\pi(x)p(x)$ , and the threshold that maximizes the welfare



function satisfies  $\pi(x) = 0$ . Thus, *the optimal threshold occurs at a value of the running variable where the LATE is zero*. Finding the optimal threshold is equivalent to finding the roots of  $\pi(x)$ . The sign of the derivative left and right of the root indicates if the root corresponds to a local maximum or minimum. Eq (4.2) can then be solved analytically or numerically, given an estimate of the function  $\pi(x)$ .

RD designs allow the estimation of  $\pi(x)$  only at  $x = c$ , since the identification of LATE is based on the assumption that treatment assignment is “as good as random” in an infinitesimally small neighborhood around  $c$ . Thus, the continuity of functions  $\mu_{1,0}$  at the threshold suffices to allow estimation of  $\pi(x)$  locally at  $c$ , but not elsewhere. In practice, however, a functional form for  $\hat{\mu}_{1,0}$  is obtained using the data in a bandwidth  $c - h < X \leq c + h$ . Our approach is to use these estimates to extrapolate within this bandwidth. Notice that in this setting, we are using the bandwidth as a range where extrapolation of  $\pi(x)$  is valid. Thus, the local optimizer is the solution to the following bound constrained optimization problem, i.e.:

$$c^* = \arg \max_t \int_t^{c+h} \pi(x)p(x)dx, \quad c - h \leq t \leq c + h. \tag{4.3}$$

In addition to the policy invariance assumption in  $[c - h, c + h]$ , we assume the following:

**Assumption 8** (Sharp optimization).  $\pi(x), p(x)$  are continuous in  $[c - h, c + h]$ .  $p(x) > 0$  in  $[c - h, c + h]$ .

Under Assumptions 7 and 8, it is straightforward to show that  $c^*$  is a root of  $\pi(x)$ , or lies on the border of the bandwidth.

For fuzzy RD designs, the welfare function is equal to the cumulative treatment effect on compliers:

$$W_f(t) = \int_t^{x_{max}} S_f(x, t)p_f(x)p(x)dx \tag{4.4}$$

Similar to the sharp RDD, we are interested in finding the optimal threshold  $c^*$  in the interval  $[c - h, c + h]$ :

$$c^* = \arg \max_t \int_t^{c+h} S_f(x, t)p_f(x)p(x)dx, \quad c - h \leq t \leq c + h. \tag{4.5}$$

Under the policy invariance assumption, Eq (4.5) becomes

$$c^* = \arg \max_t \int_t^{c+h} \pi_f(x)p_f(x)p(x)dx, \quad c - h \leq t \leq c + h. \tag{4.6}$$

Similar to the sharp RD, we make the following additional assumptions:

**Assumption 9** (Fuzzy optimization).  $\pi_f(x), p(x), p_f(x)$  are continuous in  $[c - h, c + h]$ .  $p_f(x), p(x) > 0$  in  $[c - h, c + h]$ .

Under assumptions 7, 9, the optimal threshold is a root of the treatment effect function,  $\pi_f(x) = 0$ . If  $\pi_f(x)$  does not have a root in  $[c - h, c + h]$ , then the local optimum lies on one of the borders of the interval.

The global optima of Eqs (4.2) and (4.4) could lie outside the bandwidth around the implemented threshold, and can be very different from the local optima in the within-bandwidth versions in Eqs (4.3) and (4.6). However, we cannot safely extrapolate  $\pi(x)$  away from the bandwidth, without losing the benefits of the RD design, or without additional assumptions. Moreover, we argue that policy thresholds are typically implemented based on some prior evidence on the benefits/cost of the policy. We therefore expect that the global optimum should lie close to the implemented policy thresholds. Furthermore, we view policy design and

optimization as an iterative process. In this process, the prediction of an improved outcome results in moving the threshold within a small neighborhood of the existing threshold (in analogy to trust region methods used in machine learning [28]); in turn, researchers can collect new data from the new policy, extending the data support for  $\pi(x)$  and allowing for further extrapolation.

Importantly, the threshold optimization approach presented here also allows a new interpretation to RDD estimates. Typically, a zero LATE is considered to imply an ineffective intervention. If LATE is zero for all values of the running variable, then the intervention is ineffective, and changing the threshold will not help (the MTTE is zero). However, a zero LATE can indicate an intervention for which utility is maximized: *zero LATE accompanied by a non-zero MTTE at the threshold is consistent with the threshold being set optimally to maximize outcomes.*

### 4.1 Estimation with local linear regression

In the previous section, we defined the problem of threshold optimization, and showed that the optimal threshold corresponds to a zero LATE ( $\pi(x) = 0$ ). To identify the optimal threshold, we need to estimate  $\pi(x) = \mu_1(x) - \mu_0(x)$  in a neighborhood of  $c$  and identify the threshold that leads to the optimal value of the welfare function within this neighborhood.

Typically,  $\mu_{1,0}$  are estimated using locally linear regression (LLR), which has been shown to have desirable bias properties at the threshold [7]. In the case of a rectangular kernel, this approach is equivalent to taking standard OLS estimates on both sides of the threshold. Different kernels could be used, but typically have little impact on the estimated effect [8]. Thus, we will focus on LLR with a rectangular kernel.

For a given bandwidth  $h$ , the regression model below the threshold is

$$Y(0) = \alpha_0 + \beta_0(X - c) + \varepsilon, \quad -h \leq X - c < 0$$

while the regression model above the threshold is

$$Y(1) = \alpha_1 + \beta_1(X - c) + \varepsilon, \quad 0 \leq X - c \leq h,$$

where  $\varepsilon$  follows a zero mean Gaussian distribution with standard deviation  $\sigma$ . Using OLS estimates for  $\alpha_{0,1}, \beta_{0,1}$ , the treatment effect is a linear function of the running variable:

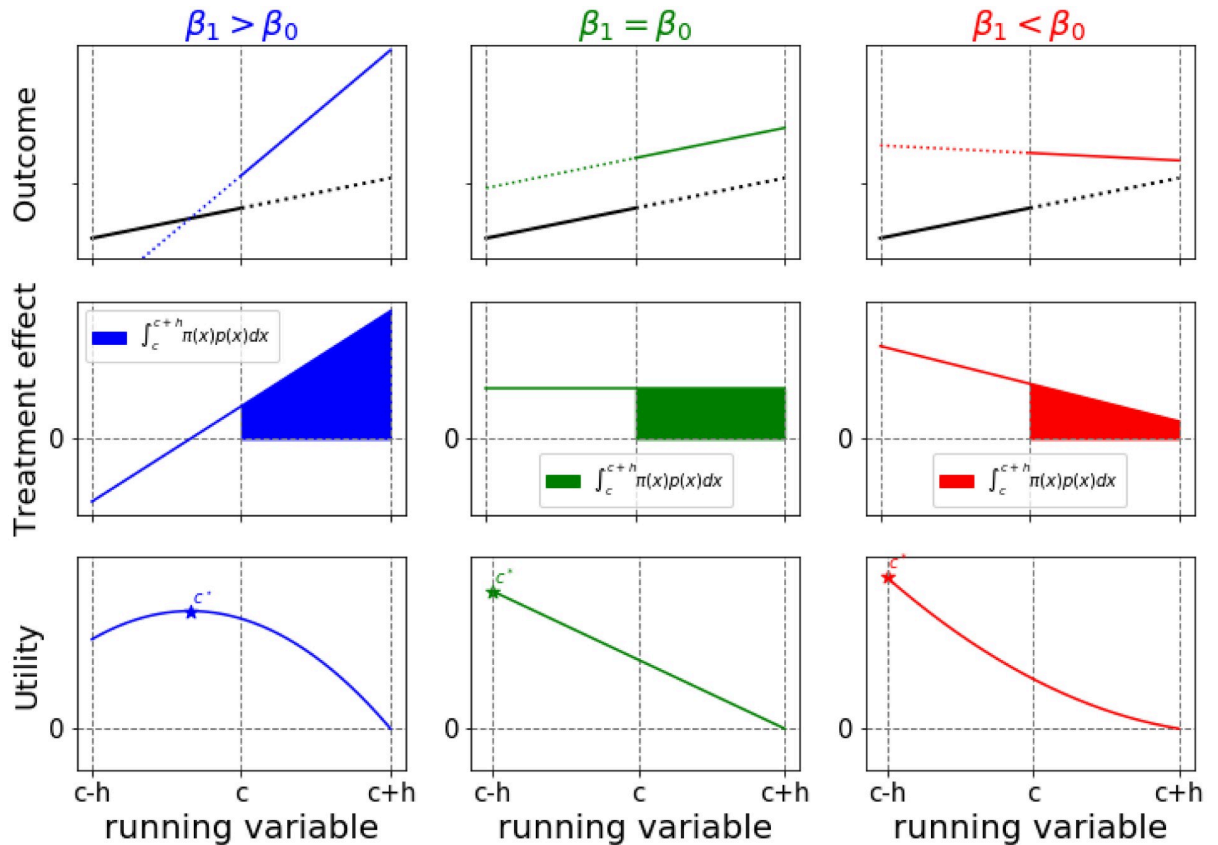
$$\hat{\pi}(x) = \hat{\mu}_1(x) - \hat{\mu}_0(x) = \hat{\alpha}_1 - \hat{\alpha}_0 + (\hat{\beta}_1 - \hat{\beta}_0)(x - c).$$

At the threshold  $c$ , we have the RDD treatment effect estimate which is equal to the difference of intercepts  $\hat{\pi}(c) = \hat{\alpha}_1 - \hat{\alpha}_0$ . The utility function has one global extremum at  $\hat{\pi}(x) = 0$ , thus

$$\hat{c}^* = c - \frac{\hat{\alpha}_1 - \hat{\alpha}_0}{\hat{\beta}_1 - \hat{\beta}_0}. \tag{4.7}$$

There are a range of different scenarios that matter for the threshold optimization problem in the context of LLR (Fig 1). Intuitively, we know that if the LATE is positive, we need to move the threshold to the left in order to treat more units, and if the LATE is negative we need to move it to the right. In Fig 1, we focus on cases where the LATE is positive, and vary the treatment effect derivative. The canonical case has a positive TED difference and a positive LATE (left panels in Fig 1). (Note that, if the TED were not higher right than left, the two lines would fail to cross.) However, there can be exceptions. For example, the LATE may be nonzero while the TED difference is zero. This may suggest that the treatment is always better (or worse) than the control and the threshold optimization problem becomes





**Fig 1. Optimizing the threshold for locally linear regression (LLR) models.** (top panel) Possible LLR models for outcome vs running variable. The discontinuity at  $c$  corresponds to the LATE  $\pi(c)$ . (middle panel) Treatment effect  $\pi(x)$  as a function of the running variable; the colored area corresponds to the utility. (bottom panel) Utility as a function of the chosen threshold, for uniform density of observations. The constrained optimization problem has one global optimal, either where  $\pi(x)$  crosses the x-axis (left column) or at the left boundary (middle, right column).

<https://doi.org/10.1371/journal.pone.0276755.g001>

locally undefined (middle panels in Fig 1). Further, there may be some counterintuitive  $\pi$  where the LATE and the TED difference have opposite signs, indicating that the treatment stops being effective (or even becomes harmful) for large values of the running variable (right panels in Fig 1). In this paper we will generally analyze the canonical case in the left panels of Fig 1.

Within the bandwidth, if the treatment effect has a root  $c'$  on the left of our observed threshold  $c$  (Fig 1, left column), then this root is the global maximum of the utility function. This is the canonical case where the LATE is zero at the optimal threshold  $c^*$  and positive to the right of  $c^*$ , and we further have  $\beta_1 > \beta_0$ , i.e. the treatment effect derivative is greater above  $c^*$  than below  $c^*$ . This shows that a zero LATE does not demonstrate that a policy is ineffective: a zero LATE with an increase in the treatment effect derivative is the sign of an optimally set threshold.

If the treatment effect does not have a root  $c'$  on the left of threshold  $c$  (Fig 1, middle and right columns), the model suggests that the threshold should be moved to the left boundary (or even beyond). This yields a less trustworthy estimate, since we have less faith in the modeling assumptions as we move further and further away from the threshold. We later discuss more conservative estimates for threshold updating.

### 4.2 Non-linear estimation with Gaussian process regression

If we do not want to assume that the relationship between the running variable and the outcome is linear close to the threshold, we can use other strategies. For example, controlling for a higher-order local polynomial in the running variable has been popular. However, this method produces estimates that are particularly sensitive to the choice of polynomial degree [29]. Gaussian processes (GPs) refer to a class of approaches that offer an alternative way of setting up the nonlinear regression problem.

Recently, Branson et al. [30] proposed using Gaussian process regression (GPR) to model the conditional expectations  $E[Y(1)|X = x]$  and  $E[Y(0)|X = x]$  and estimate the LATE at the threshold  $c$  in RD designs. The authors show that GPR produces consistent estimates for LATE and can outperform state-of-the-art LLR methods, especially in terms of coverage.

GPR derives from the Bayesian idea that we should start with a prior belief over the functions we may encounter. Such a belief could for example be that smoother functions are more probable than less smooth ones. More specifically, we are interested in the function  $f$  that maps the running variable to the outcomes, so we need a prior belief  $p(f)$  about the possible functions  $y = f(x)$ , where  $y$  can denote either of the potential outcomes  $y_0$  or  $y_1$ . Then, given a data set  $D = \{(x_1, y_1), \dots, (x_N, y_N)\}$  of observations we can calculate the posterior probability over functions  $f$ :

$$p(f|D) = p(f) \frac{p(D|f)}{p(D)}$$

Thus, instead of inducing a probability distribution over parameters as in a linear regression, the dataset will induce a probability distribution over functions themselves.

We now need to specify a prior over functions  $f$  (note that the functions are random variables in this context) of the running variable  $X$ , and GPs derive from the use of Gaussian priors. Thus  $p(f)$  is a GP if for any subset of values of the running variable  $x_i, \dots, x_k$ , the joint distribution of the corresponding outcome functions  $f(x_i), \dots, f(x_k)$  is a multivariate Gaussian. GPs are parameterized by the mean of the GP  $\mu(x)$  and its covariance function or Kernel  $K(x, x')$ . For pairs of  $x_i, x_j$  we thus have

$$p(f(x_i, x_j)) = N(m, \Sigma)$$

where

$$m = \begin{bmatrix} m(x_i) \\ m(x_j) \end{bmatrix}$$

and

$$\Sigma = \begin{bmatrix} K(x_i, x_i) & K(x_i, x_j) \\ K(x_j, x_i) & K(x_j, x_j) \end{bmatrix}$$

For the canonical choice of  $m(x) = 0$  and a  $K$  that increases with distance, this induces a prior over functions  $f$  that have varying smoothness aspects [31].

A very popular choice is the squared exponential kernel:

$$K(x_i, x_j) := \sigma_f^2 \exp\left(-\frac{\|x_i - x_j\|^2}{2\ell^2}\right)$$

The parameters of the squared exponential kernel are the variance  $\sigma_f^2$ , which formalizes how

much the function varies around each mean, and the spatial scale  $\ell$ , which formalizes how quickly the function can change when  $x_i$  and  $x_j$  get further apart. The larger  $\ell$  the lower the probability of non-smooth functions. Once mean and kernel functions are defined, we can fully define the Bayesian model:

$$y_i = f(x_i) + \epsilon_i$$

where  $f \sim GP(0, K)$ ,  $\epsilon_i \sim N(0, \sigma_i^2)$

The likelihood of a GP is Gaussian, so the posterior is also a GP. We can, therefore use the properties of the multivariate Gaussian distribution to make predictions for the value of the outcome  $y_*$  at any covariate location  $\mathbf{x}_*$ :

$$p(y_* | \mathbf{x}_*, D) = \int p(y | \mathbf{x}_*, f, D) p(f | D) df$$

In the context of RD designs, GPs can be used to describe the potential outcomes as a function of the running variables [30]:

$$Y_i(0) = \mu_0(x_i) + \epsilon_{i0}, \epsilon_{i0} \sim N(0, \sigma_0^2)$$

$$Y_i(1) = \mu_1(x_i) + \epsilon_{i1}, \epsilon_{i1} \sim N(0, \sigma_1^2).$$

The conditional expectations  $\mu_0(x_i)$ ,  $\mu_1(x_i)$  are assumed to be independent. Thus, a separate model can be fit independently on each side of the threshold. Using the same kernel in both priors corresponds to assuming equal covariance matrices for treatment and control, and is similar in nature with fitting LLR with different slopes and intercepts but with the same bandwidth. The LATE is then the difference of the corresponding posteriors:

$$\hat{\pi}(c) = \mu_1(c) | \mathbf{x}, \mathbf{y} - \mu_0(c) | \mathbf{x}, \mathbf{y}.$$

To identify the optimal threshold for the utility function, we need to extrapolate the posterior functions to obtain an estimate of LATE at any covariate value  $x$ :

$$\hat{\pi}(x) = \mu_1(x) | \mathbf{x}, \mathbf{y} - \mu_0(x) | \mathbf{x}, \mathbf{y},$$


and then numerically solve the optimization problem (i.e. find a root for  $\hat{\pi}(x)$  near  $c$ ).

GP extrapolation is also called Kriging and we know much of the properties of this estimator [32, 33]. GPs implement many intuitions that we have about functions in many real world situations, such as smoothness of the function.

### 5 Risk averse threshold optimization

In many cases, we are risk averse. Moving the threshold to the value that maximizes the expected utility implies moving to an expected zero local average treatment effect with some uncertainty. However, we may not want to risk a negative LATE. More specifically, we may be unwilling to accept a new threshold if its LATE has considerable probability of being negative. After all, that would mean that there are a range of values of the running variable where the treatment is actually harmful. Such a risk-averse objective may be particularly relevant in domains in which the use of RDD approaches is not yet established, e.g. in medicine [34].

In this case, we want to move to a threshold that is into the right direction (i.e. towards maximizing the utility function) but guarantees, with some confidence, that it will not produce harm, i.e. it will not result in a negative LATE. A natural way of implementing this, would be to use a threshold  $c_\alpha$  which corresponds to the value for which the lower bound of the one-



**Fig 2. An example of conservative threshold estimation for a risk-averse policy, for 100 data points simulated from the model in Fig 1 (left).** Using the  $(100 - \alpha)\%$ -confidence intervals for  $\pi(x)$ , we can calculate conservative optimal thresholds  $c^*_\alpha$ , that improve the utility of the intervention, but guarantee with a probability of  $(100 - \alpha)\%$  that the LATE  $\pi(c^*_\alpha)$  will be non-negative.

<https://doi.org/10.1371/journal.pone.0276755.g002>

sided  $(100 - \alpha)\%$ -confidence interval for  $\hat{\pi}(x)$  is equal to 0. If  $\sigma_{\hat{\pi}(x)}$  is the standard deviation of  $\hat{\pi}(x)$  as a function of the threshold, this threshold satisfies

$$\hat{\pi}(x) - Z_\alpha \sigma_{\hat{\pi}(x)} = 0, \tag{5.1}$$

where  $Z_\alpha$  is the critical value of the one tailed Z-test at level  $\alpha$ . Based on the estimation procedure for  $\hat{\pi}(x)$ , we can obtain standard errors and confidence or credible intervals. Note that one may also want to use robust estimates, see [18]. The threshold for which Eq (5.1) holds can then be solved numerically.

Being risk averse results in smaller threshold changes. Fig 2 shows an example. Fig 2(a) shows 100 data points simulated from the model shown in Fig 1(left column) and the corresponding fitted LLR models left and right. Coloured vertical lines indicate the predicted optimal threshold for various confidence levels.

For smaller sample sizes or large variances, it is possible that the threshold cannot be moved at all. However, if we have a LATE estimate that is significantly larger than zero at level  $\alpha$ , we can always lower the threshold to a value that results in a non-zero LATE with  $(100 - \alpha)\%$  probability and a larger utility value.

### 6 Threshold optimization with costs constraints

Most policies result in a cost per subject treated. Assume that treatment of each subject costs a fixed amount  $z$ . In the case of the linear regression (section 4.1), the optimal threshold is no longer for  $\hat{\pi}(x) = 0$ , but for  $\hat{\pi}(x) = z$ , reflecting the condition that marginal cost equals

marginal benefit. This yields

$$\hat{c}^* = c - \frac{(\hat{\alpha}_1 - z) - \hat{\alpha}_0}{\hat{\beta}_1 - \hat{\beta}_0}. \tag{6.1}$$

Typically, policy implementations are subject to additional constraints related to the budget for the intervention, so that the treatment cost for all treated subjects cannot exceed a given budget  $B$ . In that case, for an  $N$ -sized population (including treated and controls), finding the optimal threshold corresponds to the following constrained optimization problem:

$$\begin{aligned} \text{maximize}_t \quad & c^* = \arg \max_t \int_t^{c+h} \pi(x)p(x)dx, \\ \text{such that} \quad & N \int_{c^*}^{x_{max}} p(x)dx \leq B, \end{aligned} \tag{6.2}$$

Without additional assumptions about the population density  $p(x)$ , Problem (6.2) is a non-linear constrained optimization problem with non-linear constraints. The problem can be solved numerically, using any parametric or non-parametric density estimator for  $p(x)$ .

Other types of constraints are also possible. For example, a prescribed drug may have adverse effects that are also related to the running variable. Depending on the problem-specific cost constraints, the constrained threshold optimization problem can be solved analytically or numerically.

### 7 Case study

We use our approaches on an example from the economics literature, to show how we can estimate the optimal threshold for an intervention. We use the sharp RD design from [23], who study the effect of default options in consumer choices. The authors use data from all New York City yellow cab rides in 2009; summary statistics for the data can be found in Table 1. During this period, customers of a specific credit card machine company (vendor) were offered different default tip suggestions based on the the cab fare. Specifically, the default tip suggestions for fares under \$15 were \$2, \$3, and \$4. For fares over \$15, the tip suggestions were percentages of the fare amount: 20%, 25%, and 30%. At the \$15 discontinuity, this change of policy corresponds to an increase of the suggested amounts by respectively \$1, \$.75 and \$.5. Assuming that features that affect tips vary smoothly with the ride fare, the average difference of tips right and left of the margin can be interpreted as a result of policy change, i.e. higher default tip suggestions at the margin.

**Table 1. Summary statistics for the data used in the case study.** Standard deviations are in parenthesis. The sample is limited to rides without tolls, taxes, or surcharges (January 1,2009-January 31, 2009; 6 am-4 pm on Monday-Friday and 6 am-8 pm on Saturday and Sunday), as described in [23]. Data were downloaded from [https://www.aeaweb.org/aej/app/data/0603/2013-0098\\_data.zip](https://www.aeaweb.org/aej/app/data/0603/2013-0098_data.zip).

Summary statistics		
	T = 0 (Fare < \$15)	T = 1 (Fare > \$15)
Fare	8.727 (6.506)	18.419 (6.888)
Tip amount as percentage of fare	21.1853 (177.050)	17.054 (80.416)
Num. observations	453371	54269

<https://doi.org/10.1371/journal.pone.0276755.t001>

[23] used local linear regression with a bandwidth of \$10 around the discontinuity, and found that the local average treatment effect for cab fares of \$15 was an increase of approximately \$0.30 in tip amounts. This constitutes an increase in the average tip at the margin of more than 10%.

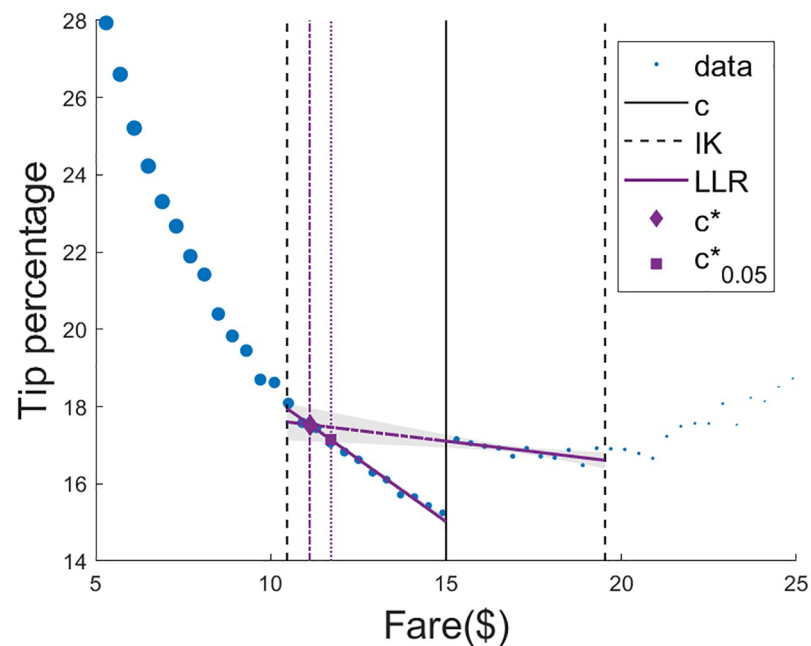
Optimizing the threshold in this case study can allow us to maximize the average tipping over all customers that receive the increased tipping suggestions. The positive effect at the discontinuity suggests that the policy threshold should be moved to the left, i.e. people with lower cab fares should also receive percentage tip suggestions. The original data set consists of 13,820,784 data points. To make the problem tractable, we used the data from all rides in January of 2009 with the specified vendor, resulting in 453,371 rides that cost \$5—\$15, and 54,269 rides that cost \$15—\$25. We used both local linear regression (LLR) and Gaussian Process regression (GPR) to model the conditional expectations  $\mu_1(x)$  and  $\mu_0(x)$  left and right of the threshold.

We used LLR with the bandwidth selection method of [24] a uniform kernel. The optimal bandwidth selected by this method was 4.54. LLR with a uniform kernel were estimated on both sides of the threshold. The optimal threshold  $c^*$  was calculated using Eq (4.1) at \$11.11 (see Fig 3).

The overall gain for drivers from switching to the optimal threshold corresponds to the expected gain in tip percentage:

$$\int_{c^*}^c (\hat{\mu}_1(x) - \hat{\mu}_0(x))p(x)dx.$$

Using the discretization bin of \$0.40 in the interval  $[c^*, c]$ , we calculate  $p(x_i)$ , the frequency



**Fig 3. Using local linear regression to predict the optimal threshold for changing default tip suggestions in taxi cabs.** Each dot is the average within a discrete fare amount (\$0.40 intervals). Dot size is proportional to the number of observations in that interval. Policy for default tip suggestions changes for \$15 rides. Imbens-Kalyanaraman (IK) bandwidth was estimated at \$4.54 around the threshold. LLR models were fit left and right of the threshold with in the IK bandwidth. The optimal threshold is at \$11.11, and the conservative threshold for significance level  $\alpha = 0.05$  is at \$11.72.

<https://doi.org/10.1371/journal.pone.0276755.g003>



of observations in each bin, where  $x_i$  are the discretized values of the running variable. The expected increase in tip percentage is then the finite sum  $\sum_i (\hat{\mu}_1(x) - \hat{\mu}_0(x))p(x_i) = 0.90$  percentage points. This corresponds to an expected gain of  $\sum_i (\hat{\mu}_1(x) - \hat{\mu}_0(x))p(x_i)x_i = \$0.12$  per ride. Thus, the method suggests that, by choosing the optimal threshold of \$11.11 for switching default tip suggestions, taxi drivers would have earned, in total, an additional \$11,011 over 89,805 rides with rates within [\$11.11, \$15] in the month of January.

We want to note though, that there may be a negative side effect of suggesting large tip amounts: disgruntled customers, who may end up tipping less than predicted by our extrapolation. This would push us to be a bit more conservative. The conservative optimal threshold  $c_{0.05}^*$  satisfying Eq (5.1) for  $\alpha = 0.05$  was estimated using MATLAB's default root finding function [35] at \$11.72. The expected increase in tip percentage is 1.17 percentage points, corresponding to an additional \$0.16 per ride. The gain per ride is higher for the conservative optimal threshold because we only include ranges of the running variable where the gain is high enough to avoid the risk of a negative LATE. On the other hand, the overall additional expected gain is lower with the conservative optimal threshold, at \$10,409 over 64,489 rides with rates within [\$11.72, \$15] in the month of January. (Fig 3). Table 2 shows bootstrap estimates and standard errors for  $c^*$  and  $c_{0.05}^*$ .

We also used GPR with the original bandwidth used by [23]. We used the double exponential kernel, with hyper-parameters optimized using cross-validation on each side of the threshold. The optimal threshold was estimated using MATLAB's default root finding function [35]. Optimization predicts that, to maximize average tipping across the bandwidth, the default tip suggestions should switch to percentages for rides above  $c^* = \$12.05$  (Fig 4). Due to the large sample size, we were unable to perform exact inference and obtain standard errors for the fit, and therefore do not include a conservative threshold prediction for Gaussian process regression. Moreover, due to the large sample size, we used subsampled bootstrap of 10,000 observations per repetition to compute bootstrap estimates and standard errors for LATE and  $c^*$ . In each sample, we again used the double exponential kernel, with hyper-parameters optimized using cross-validation on both sides of the threshold. Results are shown in Table 2.

Overall, both methods produce similar results. Moreover, the conservative threshold is really close to the unconstrained optimal threshold, due to the (unusually) large sample of the case study. Both methods suggest that, to maximize average tipping within the bandwidth, the default tip suggestions should switch to percentages for cab fares above \$11.11 – \$12.05 instead of \$15.


## 8 Discussion and conclusion

Here we have shown how RD design can be exploited not only to estimate the local average treatment effect, but also to optimize the threshold itself. We proposed threshold optimization

**Table 2. Optimal thresholds using LLR and GPR.** To estimate IK bandwidths and LLR coefficients, data were re-centered around the discontinuity  $c = 15$ . For both methods, bootstrap estimates were calculated over 100 bootstrap samples. Bootstrap standard errors are in parenthesis.

Optimal thresholds					
		IK bandwidth	LATE	$c^*$	$c_{0.05}^*$
LLR	All data	4.536	2.086	11.109	11.721
	Bootstrap	4.280 (0.576)	2.058 (0.112)	10.829 (0.777)	11.596 (0.502)
GPR	All data		2.017	12.048	
	Bootstrap		1.492 (0.783)	14.172 (2.458)	

<https://doi.org/10.1371/journal.pone.0276755.t002>



**Fig 4. Using Gaussian process regression to predict the optimal threshold for changing default tip suggestions in taxi cabs.** Each dot is the average within a discrete fare amount (\$0.40 intervals). Dot size is proportional to the number of observations in that interval. Policy for default tip suggestions changes for \$15 rides. Gaussian process regression was fit on the original bandwidth, using a double exponential kernel with parameters optimized using cross-validation. The optimal threshold is at \$12.05. Due to large sample size, we were unable to obtain standard errors and estimate the conservative optimal threshold.

<https://doi.org/10.1371/journal.pone.0276755.g004>

using local linear regression and Gaussian process regression, as well as conservative threshold optimization for situations where we are risk averse.

Most importantly, we show that RD designs with zero (or not statistically significant) estimates do not necessarily indicate ineffective policies. On the contrary, such interventions can correspond to policies implemented at the optimal threshold that maximizes welfare.

Optimizing the policy implementation threshold requires the assumption of policy invariance, i.e. that the LATE given the running variable does not depend on the threshold at which the policy is implemented. While there are cases where the assumption is violated, it is often satisfied in domains like medicine, where threshold optimization may be of great interest.

Given policy invariance, threshold optimization ultimately boils down to curve-fitting and extrapolation, since we are trying to guess the value of the LATE for unobserved regions of the running variable. Thus, any method is subject to modelling assumptions beyond the threshold.

Threshold optimization naturally implies that a meaningful welfare function exists and it is adequately measured. For example, if we want just want to maximize expected earnings without regard to any costs, our method can be applied straightforwardly as long as earnings data is available. However, in other cases, we have to provide as an outcome a summary measure that includes both costs and benefits. In many scenarios, e.g. educational policy [36], defining the utility function to be maximized is hard, since benefits and costs accrue to different entities. Consider the problem of student lifetime earnings versus cost of education to the

university: should a university persist in educating students with low GPAs? Even though an RDD estimation shows that the earnings effects of allowing a few more students to stay enrolled can be positive [37], should the university spend a dollar for every extra dollar these students make? In such cases, constrained threshold optimization that abides to cost constraints could be explored. On the other hand, there are many domains where there is a consensus on what outcome we should maximize. For example, maximizing quality of life adjusted years is a standard goal in health care. In these cases, exploiting RD structure to optimize policy thresholds can have a straightforward and significant impact on policy design. While defining a meaningful welfare function can be difficult in some cases, specific welfare functions are well accepted by the community of experts in many domains.

Interestingly, the optimization approach could be reversed. If we are, for example, interested in the motivations of a politician (or decision-making body), we might want to ask for which social welfare function the actually chosen threshold is optimal. This might give evidence, for example, about how much a politician weights the potential future earnings of disadvantaged students relative to the immediate cost of tax increases to fund public university education.

Direct threshold optimization problems abound. In medicine, countless thresholds guide physician behavior. These could all be optimized for outcomes. In finance, thresholds that affect access to financial products by customers could be optimized. For internet companies, the probability of someone clicking ads can be optimized. And, in public policy, the various program thresholds could also be optimized. There are many domains where optimizing thresholds in the context of an RDD seems highly promising.

## Author Contributions

**Conceptualization:** Ioana Marinescu, Sofia Triantafillou, Konrad Kording.

**Data curation:** Ioana Marinescu.

**Formal analysis:** Ioana Marinescu, Sofia Triantafillou, Konrad Kording.

**Funding acquisition:** Konrad Kording.

**Investigation:** Ioana Marinescu.

**Methodology:** Ioana Marinescu, Konrad Kording.

**Supervision:** Ioana Marinescu, Konrad Kording.

**Validation:** Ioana Marinescu, Sofia Triantafillou.

**Writing – original draft:** Ioana Marinescu, Sofia Triantafillou, Konrad Kording.

**Writing – review & editing:** Ioana Marinescu, Sofia Triantafillou, Konrad Kording.

## References

1. Thistlethwaite Donald L., and Campbell Donald T. 1960. "Regression-Discontinuity Analysis: An Alternative to the Ex-Post Facto Experiment." *The Journal of Educational Psychology*, 51(6): 309–317. <https://doi.org/10.1037/h0044319>
2. Pandya Ankur, Sy Stephen, Cho Sylvia, Weinstein Milton C, and Gaziano Thomas A. 2015. "Cost-effectiveness of 10-year risk thresholds for initiation of statin therapy for primary prevention of cardiovascular disease." *Jama*, 314(2): 142–150. <https://doi.org/10.1001/jama.2015.6822> PMID: 26172894
3. James Paul A, Oparil Suzanne, Carter Barry L, Cushman William C, Cheryl Dennison-Himmelfarb, Handler Joel, et al. 2014. "2014 evidence-based guideline for the management of high blood pressure in adults: report from the panel members appointed to the Eighth Joint National Committee (JNC 8)." *Jama*, 311(5): 507–520. <https://doi.org/10.1001/jama.2013.284427> PMID: 24352797

4. Geneletti Sara, Ricciardi Federico, O’Keeffe Aidan G, and Baio Gianluca. 2019. “Bayesian modelling for binary outcomes in the regression discontinuity design.” *Journal of the Royal Statistical Society: Series A (Statistics in Society)*, 182(3): 983–1002. <https://doi.org/10.1111/rssa.12440>
5. Mester Loretta J, et al. 1997. “What’s the point of credit scoring?” *Business review*, 3(Sep/Oct): 3–16.
6. McMahan H. Brendan et al. 2013. “Ad Click Prediction: A View from the Trenches.” *KDD ’13*, 1222–1230. New York, NY, USA: ACM.
7. Fan Jianqing, and Gijbels Irene. 1996. *Local polynomial modelling and its applications: monographs on statistics and applied probability* 66. Vol. 66, CRC Press.
8. Lee David S, and Lemieux Thomas. 2010. “Regression discontinuity designs in economics.” *Journal of economic literature*, 48(2): 281–355. <https://doi.org/10.1257/jel.48.2.281>
9. Li Fan, Mattei Alessandra, and Mealli Fabrizia. 2015. “Evaluating the causal effect of university grants on student dropout: evidence from a regression discontinuity design using principal stratification.” *The Annals of Applied Statistics*, 9(4): 1906–1931. <https://doi.org/10.1214/15-AOAS881>
10. Mattei Alessandra, and Mealli Fabrizia. 2017. “Regression discontinuity designs as local randomized experiments.” *Observational Studies*, 3(2): 156–173. <https://doi.org/10.1353/obs.2017.0004>
11. Mealli Fabrizia, and Rampichini Carla. 2012. “Evaluating the effects of university grants by using regression discontinuity designs.” *Journal of the Royal Statistical Society: Series A (Statistics in Society)*, 175(3): 775–798. <https://doi.org/10.1111/j.1467-985X.2011.01022.x>
12. Rokkanen, Miikka AT. 2015. “Exam schools, ability, and the effects of affirmative action: Latent factor extrapolation in the regression discontinuity design.”
13. Angrist Joshua D, and Rokkanen Miikka. 2015. “Wanna get away? Regression discontinuity estimation of exam school effects away from the cutoff.” *Journal of the American Statistical Association*, 110(512): 1331–1344. <https://doi.org/10.1080/01621459.2015.1012259>
14. Bertanha Marinho, and Imbens Guido W. 2020. “External validity in fuzzy regression discontinuity designs.” *Journal of Business & Economic Statistics*, 38(3): 593–612. <https://doi.org/10.1080/07350015.2018.1546590>
15. Bertanha Marinho. 2020. “Regression discontinuity design with many thresholds.” *Journal of Econometrics*, 218(1): 216–241. <https://doi.org/10.1016/j.jeconom.2019.09.010>
16. Cattaneo Matias D, Keele Luke, Titiunik Rocío, and Gonzalo Vazquez-Bare. 2021. “Extrapolating treatment effects in multi-cutoff regression discontinuity designs.” *Journal of the American Statistical Association*, 116(536): 1941–1952. <https://doi.org/10.1080/01621459.2020.1751646>
17. Dong Yingying, and Lewbel Arthur. 2015. “Identifying the effect of changing the policy threshold in regression discontinuity models.” *Review of Economics and Statistics*, 97(5): 1081–1092. [https://doi.org/10.1162/REST\\_a\\_00510](https://doi.org/10.1162/REST_a_00510)
18. Calonico Sebastian, Cattaneo Matias D, and Titiunik Rocío. 2014. “Robust Nonparametric Confidence Intervals for Regression-Discontinuity Designs.” *Econometrica*, 82(6): 2295–2326. <https://doi.org/10.3982/ECTA11757>
19. Bartalotti Otávio. 2019. “Regression discontinuity and heteroskedasticity robust standard errors: evidence from a fixed-bandwidth approximation.” *Journal of Econometric Methods*, 8(1). <https://doi.org/10.1515/jem-2016-0007>
20. Armstrong Timothy B, and Michal Kolesár. 2020. “Simple and honest confidence intervals in nonparametric regression.” *Quantitative Economics*, 11(1): 1–39. <https://doi.org/10.3982/QE1199>
21. Athey Susan, and Wager Stefan. 2021. “Policy learning with observational data.” *Econometrica*, 89(1): 133–161. <https://doi.org/10.3982/ECTA15732>
22. Kitagawa Toru, and Tetenov Aleksey. 2018. “Who should be treated? empirical welfare maximization methods for treatment choice.” *Econometrica*, 86(2): 591–616. <https://doi.org/10.3982/ECTA13288>
23. Haggag Kareem, and Paci Giovanni. 2014. “Default Tips.” *American Economic Journal: Applied Economics*, 6(3): 1–19.
24. Imbens Guido, and Kalyanaraman Karthik. 2012. “Optimal Bandwidth Choice for the Regression Discontinuity Estimator.” *The Review of Economic Studies*, 79(3): 933–959. <https://doi.org/10.1093/restud/rdr043>
25. Hahn Jinyong, Todd Petra, and Klaauw Wilbert Van der. 2001. “Identification and estimation of treatment effects with a regression-discontinuity design.” *Econometrica*, 69(1): 201–209. <https://doi.org/10.1111/1468-0262.00183>
26. Dong Yingying. 2018. “Alternative Assumptions to Identify LATE in Fuzzy Regression Discontinuity Designs.” *Oxford Bulletin of Economics and Statistics*, 80(5): 1020–1027. <https://doi.org/10.1111/obes.12249>

27. Abbring Jaap H, and Heckman James J. 2007. "Econometric evaluation of social programs, part III: Distributional treatment effects, dynamic treatment effects, dynamic discrete choice, and general equilibrium policy evaluation." *Handbook of econometrics*, 6: 5145–5303. [https://doi.org/10.1016/S1573-4412\(07\)06072-2](https://doi.org/10.1016/S1573-4412(07)06072-2)
28. Conn, Andrew R, Nicholas IM Gould, and Philippe L Toint. 2000. Trust region methods. SIAM.
29. Gelman Andrew, and Imbens Guido. 2017. "Why high-order polynomials should not be used in regression discontinuity designs." *Journal of Business & Economic Statistics*, (just-accepted).
30. Branson, Zach, Maxime Rischard, Luke Bornn, and Luke Miratrix. 2017. "A Nonparametric Bayesian Methodology for Regression Discontinuity Designs." arXiv preprint arXiv:1704.04858.
31. Ghahramani, Zoubin. 2011. "A tutorial on Gaussian processes."
32. Cressie Noel. 1992. "Statistics for spatial data." *Terra Nova*, 4(5): 613–617. <https://doi.org/10.1111/j.1365-3121.1992.tb00605.x>
33. Wackernagel Hans. 1996. "Multivariate geostatistics: an introduction with applications." Vol. 33, 363A–363A, Springer.
34. Moscoe Ellen, Bor Jacob, and Till Bärnighausen. 2015. "Regression discontinuity designs are underutilized in medicine, epidemiology, and public health: a review of current and best practice." *Journal of clinical epidemiology*, 68(2): 132–143. <https://doi.org/10.1016/j.jclinepi.2014.06.021>
35. Brent Richard P. 2013. Algorithms for minimization without derivatives. Courier Corporation.
36. Jacob Brian A, and Lefgren Lars. 2004. "Remedial education and student achievement: A regression-discontinuity analysis." *Review of economics and statistics*, 86(1): 226–244. <https://doi.org/10.1162/003465304323023778>
37. Ost Ben, Pan Weixiang, and Webber Douglas. 2018. "The returns to college persistence for marginal students: regression discontinuity evidence from university dismissal policies." *Journal of Labor Economics*, 36(3): 779–805. <https://doi.org/10.1086/696204>