

# Optimizing the tie-breaker regression discontinuity design

Art B. Owen\*                      Hal Varian  
Stanford University                Google Inc.

August 2018

## Abstract

Motivated by customer loyalty plans, we study tie-breaker designs which are hybrids of randomized controlled trials (RCTs) and regression discontinuity designs (RDDs). We quantify the statistical efficiency of a tie-breaker design in which a proportion  $\Delta$  of observed customers are in the RCT. In a two line regression, statistical efficiency increases monotonically with  $\Delta$ , so efficiency is maximized by an RCT. That same regression model quantifies the short term value of the treatment allocation and this comparison favors smaller  $\Delta$  with the RDD being best. We solve for the optimal tradeoff between these exploration and exploitation goals. The usual tie-breaker design experiments on the middle  $\Delta$  subjects as ranked by the running variable. We quantify the efficiency of other designs such as experimenting only in the second decile from the top. We also consider more general models such as quadratic regressions.

## 1 Introduction

Airlines, hotels and other companies may offer incentives such as free upgrades to their most loyal customers in the expectation that those customers will respond favorably with future business. The companies wish to measure the impact of those incentives while also trying to get the greatest benefit from them. An e-commerce company might want to offer some analytic tools to the customers most likely to benefit from them, while also measuring the impact of offering those tools.

These companies can rank their customers, offer the incentive to the highest ranked ones, and then measure impact with a regression discontinuity design (RDD). Or they can run a randomized controlled experiment (RCT) and measure impact by comparing results from customers with and without the incentive. The RDD is expected to have the greatest immediate payoff while the RCT is known to be more statistically efficient.

This tradeoff is naturally handled in a tie-breaker design. For a running variable  $x$ , subjects in a tie-breaker design are allocated to a control condition if  $x \leq A$ , to a test

---

\*Art Owen is a professor at Stanford University. The work reported here was done for Google and is not part of his Stanford responsibilities

condition if  $x \geq B$  and their treatment (test or control) is randomized if  $A < x < B$ . If  $A = B$  then no subjects are randomized and the data follow an RDD as introduced by Thistlethwaite and Campbell (1960). The treatment effect is estimated as the extent to which the regression has a jump discontinuity where  $x = A = B$ . At the other extreme, if all the  $x$  values are above  $A$  and below  $B$ , then the design is an RCT as described in texts on causal inference (Imbens and Rubin, 2015) or on experimental design (Box et al., 1978; Wu and Hamada, 2011). Tie-breaker designs are also called cutoff designs (Cappelleri and Trochim, 2003) and the running variable is also called an assignment variable or a forcing variable. Sometimes we refer to subjects getting the treatment or not, in place of getting test and control levels of the treatment.

Angrist et al. (2014) use a tie-breaker design to evaluate the effects of post secondary aid in Nebraska. In that setting,  $x$  was a student ranking. Students were triaged into top, middle and bottom groups. The top students received aid, the bottom ones did not, and those in the middle group were randomized to receive aid or not. Aiken et al. (1998) report on a study about allocation of students to remedial English classes where the running variable is a measure of students' reading ability before they matriculate.

Our interest is in optimizing the size of the RCT within a tie-breaker experiment. The RCT is well known to be more statistically efficient than the RDD. See for instance Jacob et al. (2012a, Section 6). However the positive impact from the test condition is ordinarily going to be better in an RDD. Companies may have more to gain by increasing business from their best customers. Similarly, merit-based scholarships are used when one wants to get academically stronger students into a class. There is thus an exploration-exploitation tradeoff here; the RCT is better for measuring impact while the RDD is expected to have more positive impact on the subjects under study.

It is possible to study this tradeoff via extensive Monte Carlo simulations or similar numerical exploration. While that approach can be used with very detailed assumptions about the distribution of  $x$  and flexible models for the response of interest, it does not provide much insight into the general nature of the tradeoff. We consider a special case where the running variable has been rescaled to have a symmetric distribution centered at  $x = 0$ , and the experimental range is from  $A = -\Delta$  to  $B = \Delta$ . We will use a linear regression model for the response with a separate slope and intercept for test and control. In this setting, Jacob et al. (2012a, Section 6) found the RCT to be 4 times as efficient as the RDD when  $x$  is uniformly distributed.

Figure 1 illustrates tie-breaker designs for four values of  $\Delta$ . The assignment variable there has a Gaussian distribution, that we assume has been centered and scaled. The outcome variable is simulated from a linear model with a constant treatment effect. For instance, in the third panel, the top 1/6 of customers get the treatment, the bottom 1/6 do not and a fraction  $\Delta = 2/3$  of the data in the middle have randomized allocation. For a Gaussian allocation variable, the experimental region in the middle of the data is where the data are most densely packed, which will typically be desirable.

This paper is organized as follows. Section 2 introduces a two-line regression relating an outcome to the assignment variable. The slope and intercept vary between treatment and control. The assignment variable will not always be Gaussian, but we can always rank order it, so that section is based on the ranks. Section 3 shows that the statistical efficiency of incorporating  $\Delta > 0$  experimentation versus the plain regression discontinuity design at

$\Delta = 0$  is  $1 + 3\Delta^2(2 - \Delta^2)$ . Thus, statistical efficiency is a monotone increasing function of the amount of experimentation. At the extreme, a pure RCT with  $\Delta = 1$  is 4 times as efficient as the RDD. We ordinarily expect that our outcome variable will show the greatest gains if we give the treatment to the highest ranked customers. Section 4 quantifies that cost in the two-line regression model and trades it off against statistical efficiency. The optimal  $\Delta$  is then dependent on the ratio between the value per customer of the short term return and the value of the information per customer that we get for a given  $\Delta$ . Although an experiment might be designed for a linear model, once the data are collected there may be nonlinearities that warrant a more flexible model. Section 5 repeats our analysis of the linear model for a pair of quadratic regression models. In this case, regression discontinuity design has a much higher variance than the experiment does. This is in line with recent findings of Gelman and Imbens (2017). Section 6 revisits the Gaussian case that we illustrate in Figure 1. It is similar to the uniform case. Here a full RCT is  $\pi/(\pi - 2) \doteq 2.75$  times as efficient as the RDD. It is qualitatively similar to the uniform case. Section 7 describes a numerical version of our approach that does not require a simplistic regression model. One can always use brute force optimization of a Monte Carlo simulation. We show how to replace the simulation inner loop by matrix algebra allowing faster and more thorough optimization. The tie-breaker literature has emphasized experiments in the middle range of the running variable  $x$ . Section 8 looks at off center experiments, such as experimenting in just the second decile from the top. In our motivating applications, the incentive might only be offered to a small fraction of customers. Section 9 contains a short discussion of how to use the findings.

We close this introduction with some additional references. Since Thistlethwaite and Campbell (1960), there have been many applications of regression discontinuity designs, particularly in economics and political science. Textbook treatments and surveys may be found in Angrist and Pischke (2009), Angrist and Pischke (2014), Jacob et al. (2012b), Imbens and Lemieux (2008), Jacob et al. (2012a), Klaauw (2008), and Lee and Lemieux (2010).

It is well known in the literature that experiments are more efficient than regression discontinuity designs. Section 6 of Jacob et al. (2012a) discusses this point in depth. They include the four-fold efficiency improvement we get for uniformly distributed running variables and a factor of 2.75 for normal running variables. The latter goes back to Goldberger (1972).

For an historical note, a tradeoff of this kind appeared in the Lanarkshire milk experiment, described by Student (1931). The goal was to measure the effect of a daily ration of milk on the health of school children. Among many complications was the fact that some of the schools chose to give the rations to the students that they thought needed it most. While that may have been the most beneficial way to allocate the school's milk, it was very damaging to the process of learning the causal impact of the milk rations. A tie-breaker experiment might have been a good compromise.




Figure 1: Illustrative data for tie-breaker designs with  $\Delta \in \{0, 1/3, 2/3, 1\}$ , and a standardized Gaussian assignment variable. The regression discontinuity design has  $\Delta = 0$ , the randomized controlled trial has  $\Delta = 1$ . Treated points are plotted in red, control in black. Allocation is deterministic for  $x$  outside the blue lines.

## 2 Setup

We begin with a simple setting where there are an even number  $N$  of customers  $i = 1, \dots, N$ , and exactly  $N/2$  of them will receive the treatment. There is an “assignment variable”  $x_i \in \mathbb{R}$  that measures the suitability of the customer for the program. The assignment variable might be the output of a statistical machine learning model based on multiple variables, or it could be based on a subjective judgment of one or more experts or stakeholders.

We will simplify the problem by transforming  $x_i$  to be equispaced in the interval  $[-1, 1]$ . That is, after sorting the customers in increasing order of  $x_i$ , we make a rank transformation

to  $x_i = (2i - N - 1)/N$ . If  $N = 6$ , the assignment variable is  $(-5, -3, -1, +1, +3, +5)/6$ . Let  $z_i$  indicate the treatment status; subjects that receive the treatment have  $z_i = +1$  and subjects that do not receive the treatment have  $z_i = -1$ .

We denote the experimental interval by  $(-\Delta, +\Delta)$  for  $\Delta$  in  $[0, 1]$ . In our hybrid design the treatment assignment takes the form:

$$z_i = \begin{cases} 1, & x_i \geq \Delta \\ \text{random}, & |x_i| < \Delta \\ -1, & x_i \leq -\Delta \end{cases} \quad (1)$$

If  $\Delta = 0$ , then we have a classic RDD with the discontinuity at  $x = 0$ . If  $\Delta = 1$ , then we have a classic RCT. If  $0 < \Delta < 1$ , then we have a tie-breaker design with  $\Delta$  measuring amount of the randomization.

The random allocation in equation (1) will make half of the  $z_i$  for  $|x_i| < \Delta$  equal 1 and the other half will be  $-1$ . One way to do this is to choose  $z_i = 1$  for a simple random sample of half of the elements in  $R = \{i \mid |x_i| < \Delta\}$ . Stratified schemes, setting  $z_i = 1$  for exactly one random member of each consecutive pair of indices in  $R$  are also easy to implement.

The impact of the treatment is measured by a scalar outcome  $Y$  where  $Y_i$  is a measure of the benefit derived from customer  $i$ . We suppose that the delay time between setting  $z_i$  and observing  $Y_i$  is long enough to make bandit methods (see for instance, Scott (2015)) unsuitable. We will instead compare experimental designs using the following two-line regression model:

$$Y_i = \beta_0 + \beta_1 x_i + \beta_2 z_i + \beta_3 x_i z_i + \varepsilon_i, \quad (2)$$

where  $\varepsilon_i$  are IID random variables with mean 0 and finite variance  $\sigma^2 > 0$ . Our analysis is based on the regression model (2) instead of the randomization because the treatment for subjects with  $x$  outside  $(-\Delta, \Delta)$  is not random.

The effect of the treatment averaged over customers  $i = 1, \dots, N$  is  $2\beta_2$ . The factor of 2 comes from comparing  $z_i = 1$  to  $z_i = -1$ . We can also estimate whether the effect increases or decreases with  $x$ , through the coefficient  $\beta_3$ . The quantity  $2\beta_2$  is also the magnitude of the treatment effect on a (hypothetical) average customer with  $x = 0$ .

Under model (2), we can distinguish customers for whom the treatment is effective from those for whom it is not. Suppose that  $\tau$  is the incremental cost of offering the treatment to one customer. If  $\beta_3 > 0$ , then there is a cutpoint  $x_* = (\tau - \beta_2)/\beta_3$  with  $\mathbb{E}(Y \mid z = 1) - \mathbb{E}(Y \mid z = 0) \geq \tau$  for customers with  $x_i \geq x_*$ . If  $x_* \notin [-1, 1]$  then the treatment either pays off on average at all  $x$ , or pays off on average for no  $x$ . If  $\beta_3 < 0$ , then the treatment only pays off for customers with  $x_i \leq x_*$ . We discuss that case further in Section 4.

### 3 Efficiency in the two-line model

We will analyze the data  $(x_i, Y_i)$  for  $i = 1, \dots, N$  by fitting model (2) by least squares. The parameter of interest is  $\beta = (\beta_0, \beta_1, \beta_2, \beta_3)^\top$  and we assume that  $Y_i$  are independent random variables with  $\text{Var}(Y_i) = \sigma^2$ . The design matrix is  $\mathcal{X} \in \mathbb{R}^{N \times 4}$  with  $i$ 'th row  $(1, x_i, z_i, x_i z_i)$ ,

and  $\text{Var}(\hat{\beta}) = (\mathcal{X}^\top \mathcal{X})^{-1} \sigma^2$ . Because  $\sigma^2$  does not depend on  $\Delta$ , we can compare designs assuming that  $\sigma = 1$ .

Next, we look at how  $\mathcal{X}^\top \mathcal{X}$  depends on  $\Delta$ . For large  $N$  we can replace  $\sum_i x_i^2$  by  $N \int_{-1}^1 x^2 dx/2 = N/3$ . Similar integral approximations yield

$$\frac{1}{N} \mathcal{X}^\top \mathcal{X} \approx \begin{pmatrix} 1 & 0 & 0 & \phi(\Delta) \\ 0 & 1/3 & \phi(\Delta) & 0 \\ 0 & \phi(\Delta) & 1 & 0 \\ \phi(\Delta) & 0 & 0 & 1/3 \end{pmatrix}, \quad (3)$$

where where  $\phi(\Delta)$  is the average value of  $z \times x$  over the design. We let

$$z(x) = \begin{cases} -1, & x \leq -\Delta \\ 0, & |x| < \Delta \\ 1, & x \geq \Delta \end{cases}$$

and find

$$\phi(\Delta) = \frac{1}{2} \int_{-1}^1 xz(x) dx = \frac{1}{2} \int_{-1}^{-\Delta} (-x) dx + \frac{1}{2} \int_{\Delta}^1 x dx = \frac{1 - \Delta^2}{2}. \quad (4)$$

The approximation error in (3) is  $O_p(1/\sqrt{N})$  when the random  $z_i$  are assigned by simple random sampling and it is much smaller under stratified sampling. We will work with (3) as if it were exact.

We can reorder the rows and columns of (3) to make it block diagonal,

$$\begin{matrix} & 1 & zx & z & x \\ 1 & \left( \begin{matrix} 1 & \phi & 0 & 0 \\ \phi & 1/3 & 0 & 0 \\ 0 & 0 & 1 & \phi \\ 0 & 0 & \phi & 1/3 \end{matrix} \right) \end{matrix}$$

where the labels on the matrix above refer to the variables that the  $\beta_j$  multiply and  $\phi = \phi(\Delta)$ . It follows that

$$N \times \text{Var} \left( \begin{pmatrix} \hat{\beta}_0 \\ \hat{\beta}_3 \\ \hat{\beta}_2 \\ \hat{\beta}_1 \end{pmatrix} \right) = \frac{1}{1/3 - \phi^2} \begin{pmatrix} 1/3 & -\phi & 0 & 0 \\ -\phi & 1 & 0 & 0 \\ 0 & 0 & 1/3 & -\phi \\ 0 & 0 & -\phi & 1 \end{pmatrix}. \quad (5)$$

Thus the variances scale by  $(1/3 - \phi^2)^{-1}$ . The individual component variances are  $\text{Var}(\hat{\beta}_0) = \text{Var}(\hat{\beta}_2) = 1/(1 - 3\phi^2)$  and  $\text{Var}(\hat{\beta}_1) = \text{Var}(\hat{\beta}_3) = 3/(1 - 3\phi^2)$ . These variances are smallest for small values of  $\phi$ , corresponding to large values of  $\Delta$ . That is, the more randomized experimentation there is in the data, the less variance there is in the estimates. Therefore, the regression discontinuity design is worst and the randomized experiment is best. Larger values of  $\phi$  also induce stronger correlations among the  $\hat{\beta}_j$ .




Figure 2: The top curve is the limiting value of  $N\text{Var}(\hat{\beta}_3)$  plotted versus the fraction  $\Delta$  of experimental data in the hybrid. Here  $\beta_3$  is the regression coefficient of  $xz$ . The bottom curve corresponds to  $N\text{Var}(\hat{\beta}_2)$  where  $\beta_2$  is the coefficient of  $z$ .

The estimated gain from the intervention for a customer with a given  $x$  is  $\hat{\mathbb{E}}(Y \mid x, z = 1) - \hat{\mathbb{E}}(Y \mid x, z = -1) = 2(\hat{\beta}_2 + x\hat{\beta}_3)$ . Next

$$\text{Var}(2(\hat{\beta}_2 + x\hat{\beta}_3)) = 4 \times \frac{1/3 + x^2}{1/3 - \phi^2} = \frac{16(1 + 3x^2)}{1 + 3\Delta^2(2 - \Delta^2)} \quad (6)$$

after some algebra. The relative efficiency of the experiment versus regression discontinuity is

$$\frac{\text{Var}(2(\hat{\beta}_2 + x\hat{\beta}_3); \Delta = 0)}{\text{Var}(2(\hat{\beta}_2 + x\hat{\beta}_3); \Delta = 1)} = \frac{1 + 3(2 - 1)}{1 + 3 \times 0} = 4 \quad (7)$$

for all  $x$ . That is, the randomized experiment with  $N/4$  observations is as informative as the regression discontinuity with  $N$  observations and this holds uniformly over all levels of the assignment variable  $x$ . This factor of 4 is given by Jacob et al. (2012a).

Figure 2 shows the variance of the treatment effect parameters as a function of  $\Delta$ . Some values from the plot are shown in Table 1. The regression discontinuity design has four times the variance of the experiment as we saw in equation (7). The slope coefficient for treatment always has three times the variance of the intercept coefficient as follows from (5). Figure 3 show the variance of the estimated impact versus  $x$  for several choices of  $\Delta$ .

| Method                   | $\Delta$ | $\text{Var}(\hat{\beta}_2)$ | $\text{Var}(\hat{\beta}_3)$ |
|--------------------------|----------|-----------------------------|-----------------------------|
| Regression discontinuity | 0        | $4/N$                       | $12/N$                      |
| Experiment               | 1        | $1/N$                       | $3/N$                       |

Table 1: Variance of  $\hat{\beta}_2$  (treatment effect intercept) and  $\hat{\beta}_3$  (treatment effect slope) under regression discontinuity ( $\Delta = 0$ ) and randomized experiment ( $\Delta = 1$ ). It assumes that  $\text{Var}(Y | x, z) = 1$ .




Figure 3: Variance of  $2(\hat{\beta}_2 + x\hat{\beta}_3)$  versus  $x$  in the two-line model (2), for  $\Delta$  between 0 and 1 in steps of 0.1. Note that the vertical axis is logarithmic.

## 4 Cost of experimentation

We ordinarily expect the value of the incentive to increase with the variable  $x$ . In that case the greatest return on the  $N$  customers in the experiment arises from the regression discontinuity design with  $\Delta = 0$ . The information gain from  $\Delta > 0$  comes at some cost in the present sample. This section quantifies that cost.

For a deterministic allocation of  $z = 1$  or  $z = -1$  we have  $\mathbb{E}(Y | x, z) = \beta_0 + \beta_1 x + \beta_2 z + \beta_3 z x$ . When  $z$  is chosen randomly with  $\Pr(z = 1) = \Pr(z = -1) = 1/2$ , then  $\mathbb{E}(Y | x) =$



$\beta_0 + \beta_1 x$ . It follows that the expected gain per customer in the hybrid design is

$$\begin{aligned} g(\Delta) &\equiv \frac{1}{2} \int_{-1}^{-\Delta} (\beta_0 + \beta_1 x - \beta_2 - \beta_3 x) dx + \frac{1}{2} \int_{-\Delta}^{\Delta} (\beta_0 + \beta_1 x) dx \\ &\quad + \frac{1}{2} \int_{\Delta}^1 (\beta_0 + \beta_1 x + \beta_2 + \beta_3 x) dx \\ &= \beta_0 + \beta_3(1 - \Delta^2)/2. \end{aligned}$$

Neither  $\beta_1$  nor  $\beta_2$  appear in this gain and the value of  $\beta_0$  does not affect our choice of  $\Delta$ . Only  $\beta_3$  which models how the payoff from the incentive varies with the assignment variable  $x$  makes a difference. Compared to the regression discontinuity design with  $\Delta = 0$ , the cost of incorporating experimentation is

$$N(g(0) - g(\Delta)) = N\beta_3\Delta^2/2,$$

which grows slowly as  $\Delta$  increases from zero and then rapidly as  $\Delta$  approaches one.

If  $\beta_3 > 0$ , then as expected, we gain the most from the regression discontinuity design and the least from the experiment. This is a classic exploration-exploitation tradeoff.

It is also possible that some settings have  $\beta_3 < 0$ . This might happen if the incentive is additional free tutoring in the educational context, or if it is advice on how to best use an e-commerce company's products in a context where higher performing customers already knew about the advice. In these cases the minimal cost is to give the incentive to the bottom  $N/2$  customers and not the top  $N/2$  customers. The analysis of this paper goes through by reversing the customer ranking, thereby replacing  $x$  by  $-x$  and also changing the sign of  $\beta_3$ .

Now we turn to optimizing the choice of  $\Delta$  given some assumptions on the relative value of the information in the data for future decisions and the expected gain on the experiment. The precision (inverse variance) of our estimate of  $\hat{\beta}$  is a linear function of  $N$  and so is the expected gain. We can therefore trade off precision per customer with gain per customer. We think that  $\beta_3$  is the most important parameter so we take the precision gain per customer to be

$$p(\Delta) \equiv \frac{1}{N\text{Var}(\hat{\beta}_3)} = \frac{1}{3} - \phi^2 = \frac{1}{3} - \frac{(1 - \Delta^2)^2}{4}. \quad (8)$$

Alternatively, we could focus on  $2\beta_2$  which is both the average gain per customer and the gain for the customer at  $x = 0$ . The precision for  $2\beta_2$  turns out to be  $p(\Delta)/4$  so it perfectly aligned with precision on  $\beta_3$ . More generally the gain from the incentive at any specific  $x$  has a variance given by (6). Any weighted average of precision of  $2(\beta_2 + \beta_3 x)$  over points  $x \in [-1, 1]$  is a scalar multiple of  $p(\Delta)$  from (8).

We trade off gain per customer and precision per customer with the value function

$$v(\Delta) = g(\Delta) + \lambda p(\Delta) = \beta_0 + \beta_3 \frac{1 - \Delta^2}{2} + \lambda \left( \frac{1}{3} - \frac{(1 - \Delta^2)^2}{4} \right), \quad (9)$$

where  $\lambda > 0$  measures the value for future decisions of having greater precision on  $\beta_3$ .




Figure 4: The horizontal axis has  $\beta_3/\lambda$  where  $\beta_3$  affects immediate gain per customer of the treatment allocation and  $\lambda$  quantifies the value of precise slope estimation. The vertical axis gives the optimal  $\Delta$  in our hybrid experiment.

**Proposition 1.** *Let  $v(\Delta)$  be given by equation (9) with  $\lambda > 0$  and  $\beta_3 \geq 0$ . Then the maximum of  $v$  over  $\Delta \in [0, 1]$  occurs at*

$$\Delta_* = \begin{cases} 1, & \beta_3/\lambda \leq 0 \\ \sqrt{1 - \beta_3/\lambda}, & 0 \leq \beta_3/\lambda \leq 1 \\ 0, & 1 \leq \beta_3/\lambda. \end{cases} \quad (10)$$

*Proof.* Let  $\gamma = \Delta^2$ . We will first maximize  $v = c - \beta_3\gamma/2 - \lambda(1 - \gamma)^2/4$  over  $0 \leq \gamma \leq 1$ , where  $c$  does not depend on  $\gamma$ . Now  $v$  has a unique maximum over  $\gamma \in \mathbb{R}$  at  $\gamma_* = 1 - \beta_3/\lambda$ . The maximizing  $\gamma$  is  $\gamma_*$  when  $0 \leq \gamma_* \leq 1$ , it is 0 when  $\gamma_* < 0$  and it is 1 when  $\gamma_* > 1$ . Equation (10) translates these results back to the optimal  $\Delta$ .  $\square$

We see from equation (10) that the decision depends on the critical ratio  $\beta_3/\lambda$ . The numerator reflects the value of more efficient allocation and the denominator captures the value of improved information gathering. When  $\beta_3 \geq \lambda$  then the discontinuity design with  $\Delta = 0$  is optimal. The full experiment,  $\Delta = 1$ , is never optimal unless  $\beta_3 = 0$  or the value  $\lambda$  of information to be used in future decisions is infinite.

Figure 4 shows the value  $\Delta_*$  from equation (10) versus the ratio  $r = \beta_3/\lambda$  of the short term to long term value coefficients. The function is nearly equal to  $1 - r/2$  near the origin and has negative curvature on  $0 \leq r \leq 1$ . If future uses are important enough that  $r \leq 1/10$ , then one should use  $\Delta \geq 1 - 0.1/2 = 0.95$ . That is, when the future is very important the optimal hybrid is very close to an RCT.

## 5 Quadratic regression

A quadratic regression model

$$\mathbb{E}(Y) = \beta_0 + \beta_1 x + \beta_2 z + \beta_3 xz + \beta_4 x^2 + \beta_5 x^2 z \quad (11)$$

allows a richer exploration of the treatment effect. For instance, model (11) allows for the possibility that the treatment pays off if and only if  $x$  is in some interval. It also allows for a situation where the payoff only comes outside of some interval. This model has even (symmetric) predictors  $1, xz, x^2$  and odd (antisymmetric) predictors  $x, z, zx^2$ . As in the linear case, the even and odd predictors are orthogonal to each other.

Now  $(1/N)\mathcal{X}^\top \mathcal{X}$  is a  $6 \times 6$  block diagonal matrix. Some of the entries are

$$\phi_3 \equiv \phi_3(\Delta) = \frac{1}{2} \int_{-1}^1 zx^3 = \int_{\Delta}^1 x^3 dx = \frac{1 - \Delta^4}{4}$$

as well as  $\phi(\Delta)$  from Section 3 that we call  $\phi_1(\Delta)$  here. We find that

$$\frac{1}{N} \mathcal{X}^\top \mathcal{X} = \begin{matrix} & \begin{matrix} 1 & zx & x^2 & z & x & zx^2 \end{matrix} \\ \begin{matrix} 1 \\ zx \\ x^2 \\ z \\ x \\ zx^2 \end{matrix} & \begin{pmatrix} 1 & \phi_1 & 1/3 & 0 & 0 & 0 \\ \phi_1 & 1/3 & \phi_3 & 0 & 0 & 0 \\ 1/3 & \phi_3 & 1/5 & 0 & 0 & 0 \\ 0 & 0 & 0 & 1 & \phi_1 & 1/3 \\ 0 & 0 & 0 & \phi_1 & 1/3 & \phi_3 \\ 0 & 0 & 0 & 1/3 & \phi_3 & 1/5 \end{pmatrix} \end{matrix}. \quad (12)$$

Once again we get a block diagonal pattern with two identical blocks. This is a consequence of  $z^2 = 1$ , and it will happen for more general models with odd and even predictors.

**Proposition 2.** *For  $N > 0$ , let  $\mathcal{X}^\top \mathcal{X}$  be given by (12). Then*

$$(\mathcal{X}^\top \mathcal{X})^{-1} = \frac{1}{ND(\Delta)} \begin{pmatrix} M(\Delta) & 0 \\ 0 & M(\Delta) \end{pmatrix} \quad (13)$$

for a  $3 \times 3$  symmetric matrix

$$M(\Delta) = \begin{pmatrix} \frac{1}{15} - \phi_3^2 & \frac{\phi_3}{3} - \frac{\phi_1}{5} & \phi_3 \phi_1 - \frac{1}{9} \\ " & \frac{4}{45} & \frac{\phi_1}{3} - \phi_3 \\ " & " & \frac{1}{3} - \phi_1^2 \end{pmatrix},$$

and a determinant  $D(\Delta) = 4/135 - \phi_1^2/5 - \phi_3^2 + (2/3)\phi_1\phi_3$ .

*Proof.* Multiplying  $M(\Delta)$  above by the upper left  $3 \times 3$  submatrix in (12) yields  $I_3$  times  $D(\Delta)$ , after some lengthy manipulations.  $\square$




Figure 5: Variance of  $2(\hat{\beta}_2 + x\hat{\beta}_3 + x^2\hat{\beta}_5)$  versus  $x$  in the quadratic model (11), for  $\Delta$  between 0 and 1 in steps of 0.1.

Figure 5 show the variance of the estimated impact versus  $x$  for several choices of  $\Delta$ . Notice that the variance is given on a logarithmic scale there. The regression discontinuity design  $\Delta = 0$  in the top curve there, has extremely large variances especially where  $|x|$  is close to 1. The randomized design at the bottom has much smaller variance. Even the maximum variance in the experiment (at  $x = 1$ ) is smaller than the minimum variance in the regression discontinuity model (at  $x = 0$ ).

## 6 Gaussian case

The original assignment variable might have a nearly Gaussian distribution. Or we might believe that the two-line linear model fits better if we have transformed the assignment variable rank to normal scores  $x_i = \Phi^{-1}((i - 1/2)/N)$ , where  $\Phi(\cdot)$  is the cumulative distribution of the  $\mathcal{N}(0, 1)$  distribution.

We will experiment on the central data with  $|x_i| \leq \tau$  choosing  $\tau$  to get a fraction  $\Delta$  of data in the experiment. That leads to  $\tau = \Phi^{-1}((1 + \Delta)/2)$ . After reordering the variables we find in this case that

$$\frac{1}{N} \mathcal{X}^\top \mathcal{X} \approx \begin{matrix} & \begin{matrix} 1 & zx & z & x \end{matrix} \\ \begin{matrix} 1 \\ zx \\ z \\ x \end{matrix} & \begin{pmatrix} 1 & \phi_G & 0 & 0 \\ \phi_G & 1 & 0 & 0 \\ 0 & 0 & 1 & \phi_G \\ 0 & 0 & \phi_G & 1 \end{pmatrix} \end{matrix}.$$

Compared to the uniform scores case, the diagonal has changed from  $(1, 1/3, 1, 1/3)$  to  $(1, 1, 1, 1)$ . The value of  $\phi$  from the uniform case changes to

$$\begin{aligned}\phi_G &= \int_{-\infty}^{-\tau} (-x)\varphi(x) dx + \int_{\tau}^{\infty} x\varphi(x) dx = 2 \int_{\tau}^{\infty} x\varphi(x) dx \\ &= 2\varphi(\tau) = 2\varphi(\Phi^{-1}((1 + \Delta)/2)).\end{aligned}$$

Now

$$N \times \text{Var} \left( \begin{pmatrix} \hat{\beta}_0 \\ \hat{\beta}_3 \\ \hat{\beta}_2 \\ \hat{\beta}_1 \end{pmatrix} \right) = \frac{1}{1 - \phi_G^2} \begin{pmatrix} 1 & -\phi_G & 0 & 0 \\ -\phi_G & 1 & 0 & 0 \\ 0 & 0 & 1 & -\phi_G \\ 0 & 0 & -\phi_G & 1 \end{pmatrix}. \quad (14)$$

For this Gaussian case, all 4 estimated coefficients  $\hat{\beta}_j$  have the same variance, equal to  $1/(1 - \phi_G^2)$ . The variances for uniform assignment variables were not all the same. The difference stems from the points  $x_i$  having variance  $1/3$  in the uniform case instead of variance 1 here. As before as  $\Delta$  increases,  $\phi_G$  also increases and so  $\text{Var}(\hat{\beta}_j)$  decreases.

Now we work out the efficiency of the RCT compared to the RDD. For the RCT,  $\Delta = 1$  yields  $\tau = \infty$  and then  $\phi_G = 0$ . For the RDD,  $\Delta = 0$  yields  $\tau = 0$  and then  $\phi_G = 2\varphi(0)$ . Thus the efficiency of the RCT compared to the RDD is

$$\frac{1}{1 - [2\varphi(0)]^2} = \frac{\pi}{\pi - 2} \doteq 2.75$$

as reported by Goldberger (1972). This is somewhat less than the efficiency gain of 4 in the uniform case. The efficiency versus  $\Delta$  (not shown) has a qualitatively similar shape to the black curve for the coefficient of  $z$  in the uniform case (Figure 2).

## 7 General numerical approach

The two line model for a running variable  $x$  with a symmetric distribution made it simple to study central experimental windows of the form  $(-\Delta, \Delta)$ . In that setting the means of  $x_i$  and  $z_i$  were both zero, and the variance of parameter estimates depended simply on just one quantity  $\Delta$ . We may want to use a more general regression model, allow experimental windows that are not centered around the middle value of  $x$ , have  $x$  values that are not uniform or Gaussian, and we might also want to use models other than two regression lines. There might even be more than one running variable as in Abdulkadiroglu et al. (2017). The price for this flexibility is high; users have to answer some hard questions about their goals, and then do numerical optimization over parameters with a potentially expensive Monte Carlo inner loop. In this section we show that the inner loop can be done algebraically.

We suppose that prior to treatment assignment, customer  $i$  has a known feature vector  $F_i \in \mathbb{R}^d$  which includes an intercept variable equal to 1, but not the treatment variable  $z_i$ . For instance in the linear and quadratic models, the features  $F_i$  are  $(1, x_i)^\top$  and  $(1, x_i, x_i^2)^\top$ , respectively. In the regression model

$$Y_i = F_i^\top \beta + z_i F_i^\top \gamma + \varepsilon_i,$$

we have  $\mathbb{E}(Y_i) = F_i^\top(\beta + \gamma)$  for the treated customers  $i$  and  $\mathbb{E}(Y_i) = F_i^\top(\beta - \gamma)$  for the others. Here  $\gamma \in \mathbb{R}^d$  models the effect of treatment.

The generalized tie-breaker study works with a vector  $\theta \in \mathbb{R}^d$  and sets

$$z_i = \begin{cases} 1, & \theta^\top F_i \geq \Delta \\ \text{random}, & |\theta^\top F_i| < \Delta \\ -1, & \theta^\top F_i \leq -\Delta. \end{cases}$$

In the random case, we suppose that  $z_i = 1$  with probability  $p$  and is  $-1$  with probability  $1 - p$  where  $p$  need not be  $1/2$ . Because  $F_i$  contains an intercept term, the experimental window  $|\theta^\top F_i| < \Delta$  need not be centered on a central value of  $\theta^\top F_i$ . The analyst must now choose  $\Delta \geq 0$ ,  $\theta \in \mathbb{R}^d$  and  $p \in (0, 1)$ .

The analogue of our previous approach is to find the matrix  $(\mathcal{X}^\top \mathcal{X})^{-1}$  where

$$\mathcal{X}^\top \mathcal{X} = \begin{pmatrix} A & B \\ B & A \end{pmatrix}, \quad A = \sum_i F_i F_i^\top, \quad B = \sum_i w_i F_i F_i^\top,$$

for

$$w_i = \mathbb{E}(z_i | F_i) = \begin{cases} 1, & \theta^\top F_i \geq \Delta, \\ 2p - 1, & |\theta^\top F_i| < \Delta, \\ -1, & \theta^\top F_i \leq -\Delta. \end{cases}$$

The lower right corner of  $\mathcal{X}^\top \mathcal{X}$  is  $A$  because it is using  $\mathbb{E}(z_i^2 | F_i) = 1$ . Averaging over the outcomes of  $z_i$  this way is statistically reasonable when  $n \gg d$ . If  $\varepsilon_i$  are independent with mean zero and variance  $\sigma^2$ , then

$$\text{Var} \left( \begin{pmatrix} \hat{\beta} \\ \hat{\gamma} \end{pmatrix} \right) = (\mathcal{X}^\top \mathcal{X})^{-1} \sigma^2.$$

This averages over the outcomes  $\varepsilon_i$  so that they do not have to be simulated.

One can now do brute force numerical search for good values of  $\theta$  and  $p$  and  $\Delta$ . A good choice would yield a favorably small  $\text{Var}(\hat{\gamma})$ . A bad choice will yield a larger variance covariance matrix. A very bad choice would lead to singular  $\mathcal{X}^\top \mathcal{X}$  and one would of course reject the corresponding triple  $(\theta, \Delta, p)$ . For instance, such a singularity would happen if  $\max_i \theta^\top F_i < -\Delta$  which is an obviously poor choice because then no customers would be in the treatment group.

Using a formula for the inverse of a block matrix we get

$$\text{Var}(\hat{\gamma}) = \text{Var}(\hat{\beta}) = (A - BA^{-1}B)^{-1}$$

and  $\text{Cov}(\hat{\beta}, \hat{\gamma}) = -A^{-1}B(A - BA^{-1}B)^{-1}$ . In an RCT with  $p = 1/2$  we have  $B = 0$ . For  $\Delta < 1$  certain components of  $B$  become nonzero (they were  $\phi$  values off the main diagonal in the two line regression) increasing  $BA^{-1}B$  and hence increasing  $\text{Var}(\hat{\gamma})$ .

## 8 Non-central experimental regions

Our treatment of the two line model assumed that the experimental region was in the center of the range of the running variable. For a loyalty program one might prefer instead to allocate the benefit in a different way. Perhaps the top 10% get the benefit, and the next 10% are randomized to receive the benefit or not, while the bottom 80% do not get the benefit. For a less expensive incentive, the company might want to offer it to the top 50% of customers and then randomize it to the bottom 50%. We can model these options by taking

$$z_i = \begin{cases} 1, & x_i \geq b \\ \text{random}, & a < x_i < b \\ -1, & x_i \leq a \end{cases} \quad (15)$$

for  $a \leq b$ .

Let the running variable  $x \in \mathbb{R}$  be random with  $\mathbb{E}(x^4) < \infty$ . Let  $z \in \mathbb{R}$  be random with a finite value of  $\mathbb{E}(x^4)$ . Let  $z = 1$  with probability  $p(x)$  and  $z = -1$  otherwise. Then letting  $\mathcal{X}$  be the design matrix in the two line regression, and noting that  $z^2 = 1$ , we have

$$\frac{1}{n} \mathcal{X}^\top \mathcal{X} = \begin{pmatrix} 1 & \mathbb{E}(x) & \mathbb{E}(z) & \mathbb{E}(xz) \\ \mathbb{E}(x) & \mathbb{E}(x^2) & \mathbb{E}(xz) & \mathbb{E}(x^2z) \\ \mathbb{E}(z) & \mathbb{E}(xz) & 1 & \mathbb{E}(x) \\ \mathbb{E}(xz) & \mathbb{E}(x^2z) & \mathbb{E}(x) & \mathbb{E}(x^2) \end{pmatrix} + O_p\left(\frac{1}{\sqrt{n}}\right)$$

under random sampling of  $x_i$  and  $z_i$  given  $x_i$  for  $i = 1, \dots, n$ . The  $O_p(n^{-1/2})$  error holds because  $\mathbb{E}(x^4) < \infty$ . The error could be less than  $O_p(1/\sqrt{n})$  if  $p(x)$  is a simple enough function to make stratification tractable.

We can center  $x$  so that  $\mathbb{E}(x) = 0$  and then

$$\text{Var}(\hat{\beta}) \doteq \frac{1}{n} \begin{pmatrix} D & C \\ C & D \end{pmatrix}^{-1}, \text{ for } C = \begin{pmatrix} \mathbb{E}(z) & \mathbb{E}(xz) \\ \mathbb{E}(xz) & \mathbb{E}(x^2z) \end{pmatrix} \text{ and } D = \begin{pmatrix} 1 & 0 \\ 0 & \mathbb{E}(x^2) \end{pmatrix}.$$

We can scale  $x$  to get  $\mathbb{E}(x^2) = 1$  so that  $D = I_2$ . We retain more general scaling because  $x \sim U[-1, 1]$  has  $\mathbb{E}(x^2) = 1/3$  and rescaling would require working with the less convenient distribution  $U[-\sqrt{3}, \sqrt{3}]$ .

We need the inverse of a block diagonal matrix containing just two unique square blocks. The following proposition specializes block matrix inversion to our case.

**Proposition 3.** *Let  $D$  be an invertible matrix and  $C$  be a square matrix with the same dimensions as  $D$ . If  $D - CD^{-1}C$  is invertible, then*

$$\begin{pmatrix} D & C \\ C & D \end{pmatrix}^{-1} = \begin{pmatrix} A & B \\ B & A \end{pmatrix}$$

for  $A = (D - CD^{-1}C)^{-1}$  and  $B = -ACD^{-1}$ .

*Proof.* Multiplying,

$$\begin{pmatrix} A & B \\ B & A \end{pmatrix} \begin{pmatrix} D & C \\ C & D \end{pmatrix} = \begin{pmatrix} AD + BC & AC + BD \\ BD + AC & BC + AD \end{pmatrix}.$$

Now  $AC + BD = AC - ACD^{-1}D = 0$  and  $AD + BC = A(D - CD^{-1}C) = I$ .  $\square$

| Method          | $a$   | $b$  | $\text{Var}(\hat{\beta}_3)$ |
|-----------------|-------|------|-----------------------------|
| Experiment      | -1.00 | 1.00 | 3.00/ $N$                   |
| RDD             | 0.00  | 0.00 | 12.00/ $N$                  |
| Bottom 50%      | -1.00 | 0.00 | 13.09/ $N$                  |
| Second 10%      | 0.60  | 0.80 | 137.56/ $N$                 |
| Middle 10%      | -0.10 | 0.10 | 11.32/ $N$                  |
| Tiny (median)   | -0.01 | 0.01 | 11.99/ $N$                  |
| Tiny (90th)     | 0.79  | 0.81 | 739.96/ $N$                 |
| Skew RDD (90th) | 0.80  | 0.80 | 751.03/ $N$                 |
| Tiny (80th)     | 0.59  | 0.61 | 94.86/ $N$                  |
| Skew RDD (80th) | 0.60  | 0.60 | 95.21/ $N$                  |

Table 2: Variance of  $\hat{\beta}_3$  (treatment effect slope) for some central and non-central experimental regions.

Using Proposition 3 we get

$$\text{Var}(\hat{\beta}) \doteq \frac{1}{n} \begin{pmatrix} (D - CD^{-1}C)^{-1} & -(D - CD^{-1}C)^{-1}CD^{-1} \\ -(D - CD^{-1}C)^{-1}CD^{-1} & (D - CD^{-1}C)^{-1} \end{pmatrix}.$$

Our primary interest is in  $\text{Var}(\hat{\beta}_3)$ , for the coefficient of  $xz$ . This is the lower right element of  $(D - CD^{-1}C)^{-1}$ . Now

$$\begin{aligned} D - CD^{-1}C &= \begin{pmatrix} 1 - \mathbb{E}(z)^2 - \mathbb{E}(xz)^2/\mathbb{E}(x^2) & -\mathbb{E}(xz)\mathbb{E}(z) - \mathbb{E}(x^2z)\mathbb{E}(xz)/\mathbb{E}(x^2) \\ -\mathbb{E}(xz)\mathbb{E}(z) - \mathbb{E}(x^2z)\mathbb{E}(xz)/\mathbb{E}(x^2) & \mathbb{E}(x^2) - \mathbb{E}(xz)^2 - \mathbb{E}(x^2z)^2/\mathbb{E}(x^2) \end{pmatrix} \\ &\equiv \begin{pmatrix} M_{11} & M_{12} \\ M_{12} & M_{22} \end{pmatrix}, \end{aligned}$$

and so

$$\text{Var}(\hat{\beta}_3) = \frac{M_{11}}{M_{11}M_{22} - M_{12}^2}.$$

The asymptotic value of  $n\hat{\beta}_3$  depends on certain integrals. For the case of primary interest to us with  $x \sim U[-1, 1]$ , and  $p(x) = 1/2$  in the experimental region, these are

$$\begin{aligned} \mathbb{E}(x^2) &= \frac{1}{2} \int_{-1}^1 x^2 dx = \frac{1}{3}, \\ \mathbb{E}(xz) &= \frac{1}{2} \int_{-1}^a (-x) dx + \frac{1}{2} \int_b^1 x dx = \frac{1}{2} - \frac{a^2 + b^2}{4}, \\ \mathbb{E}(z) &= -\frac{1}{2}(a + 1) + \frac{1}{2}(1 - b) = -\frac{a + b}{2}, \quad \text{and} \\ \mathbb{E}(x^2z) &= \frac{1}{2} \int_{-1}^a (-x^2) dx + \frac{1}{2} \int_b^1 x^2 dx = -\frac{a^3 + b^3}{6}. \end{aligned}$$

Table 2 shows  $\text{Var}(\hat{\beta}_3)$  for various designs. The first two are the full experiment and the RDD discussed previously. Next is an experiment on just the bottom half of  $x$ . This strategy



is inadmissible by our criteria. It has more variance than the RDD and also lower allocation efficiency.

Next, the table shows  $\text{Var}(\hat{\beta}_3)$  for an experiment on just the second 10% of the running variable, from the 80th to the 90th percentiles of the  $U[-1, 1]$  distribution. Just below it is an equal sized experiment in the middle. We see that experimenting in the middle is much more informative. Shifting the experimental region to one side reduces the sample size for either the treatment or control level of  $z$ . It also affects the correlations among predictors in the two line model.

The variance for experimenting on the second decile looks large compared to the central experiments. It has within it a central experiment on just the middle third of the data from the 70th to the 100th percentiles of  $x$ . Experimenting on the middle third of  $[-1, 1]$  involves taking  $a = -1/3$  and  $b = 1/3$  which yields  $\text{Var}(\hat{\beta}_3) \doteq 7.36$ . However if we had only experimented over the range 0.4 to 1.0 (with cut points at 0.6 and 0.8) then  $N$  would be only 0.3 times as large as it is in the second decile experiment. Furthermore, reducing the range of  $x$  by a factor of 0.3 multiplies  $\beta_3$  by  $1/0.3$  and  $\text{Var}(\hat{\beta}_3)$  by  $1/0.3^2$ . To adjust for these factors we divide  $7.36/N$  by  $0.3^3$  and get  $272.72/N$ . As a result doing the experiment on the second 10% really is better than just doing a central 1/3 experiment on the top 30%.

One tiny experiment involves just randomizing for one percent of the data centered on the median of  $x$ . We get a variance of  $11.99/N$  for this compared to  $12/N$  for the RDD, so the tiny experiment is almost identical to the RDD. We can move the location of the tiny experiment. Table 2 shows the results for a tiny experiment near the 80<sup>th</sup> and 90<sup>th</sup> percentiles of  $x$ . These are quite similar to skewed RDDs where the cutpoint is off center.

## 9 Discussion

In an incentive plan, a regression discontinuity design rewards the a priori best customers but it has severe disadvantages if one wants to follow up with regression models to measure impact. There is a tradeoff between estimation efficiency and allocation efficiency. Proposition 2 provides a principled way to translate estimates or educated guesses about the present value of the incentives and future value of information into a choice of  $\Delta$  in a hybrid experiment.

In industrial settings, the incentive under study will change over time. Experience with similar though perhaps not identical prior incentive plans then gives some guidance for making the tradeoff.

We have examined a simple linear model because it is easiest to work with and is a reasonable starting point in many contexts. Analysts have many more models at their disposal when the data come in. Section 5 on the quadratic model provides a warning: the RDD becomes very unreliable already with this model which is only slightly more complicated than the two-line model.

In some applications, the allocation variable may be the output of a scoring model based on many customer variables. We expect that incorporating randomness into the design will give better data for refitting such an underlying scoring model, but following up that point is outside the scope of this article. The effects are likely to vary considerably from problem to problem.

## References

- Atila Abdulkadiroglu, Joshua D Angrist, Yusuke Narita, and Parag A Pathak. Impact evaluation in matching markets with general tie-breaking. Technical report, National Bureau of Economic Research, 2017. URL <http://www.nber.org/papers/w24172>.
- Leona S Aiken, Stephen G West, David E Schwalm, James L Carroll, and Shenghwa Hsiung. Comparison of a randomized and two quasi-experimental designs in a single outcome evaluation: Efficacy of a university-level remedial writing program. *Evaluation Review*, 22(2):207–244, 1998.
- Joshua Angrist, Sally Hudson, and Amanda Pallais. Leveling up: Early results from a randomized evaluation of post-secondary aid. Technical report, National Bureau of Economic Research, 2014. URL <http://www.nber.org/papers/w20800.pdf>.
- Joshua D. Angrist and Jorn-Steffen Pischke. *Mostly Harmless Econometrics*. Princeton University Press, Princeton, 2009.
- Joshua D. Angrist and Jorn-Steffen Pischke. *Mastering Metrics*. Princeton University Press, Princeton, 2014.
- George E. P. Box, William Gordon Hunter, and J. Stuart Hunter. *Statistics for experimenters*. John Wiley and Sons, New York, 1978.
- Joseph C. Cappelleri and William M. K. Trochim. Cutoff designs. In Marcel Dekker, editor, *Encyclopedia of Biopharmaceutical Statistics*. CRC Press, 2003. doi: 10.1081/E-EBS12000734. URL <https://www.socialresearchmethods.net/research/Cutoff%20Designs%202003.pdf>.
- Andrew Gelman and Guido Imbens. Why high-order polynomials should not be used in regression discontinuity designs. *Journal of Business & Economic Statistics*, 0(0), 2017. URL <http://www.nber.org/papers/w20405>.
- A. S. Goldberger. Selection bias in evaluating treatment effects: Some formal illustrations. Technical Report Discussion paper 128–72, Institute for Research on Poverty, University of Wisconsin–Madison, 1972.
- Guido Imbens and Thomas Lemieux. Regression discontinuity designs: a guide to practice. *Journal of Econometrics*, 142(2):615–635, 2008. URL [www.nber.org/papers/w13039.pdf](http://www.nber.org/papers/w13039.pdf).
- Guido W Imbens and Donald B Rubin. *Causal inference in statistics, social, and biomedical sciences*. Cambridge University Press, 2015.
- Robin Jacob, Pei Zhu, Marie-Andrée Somers, and Howard Bloom. A practical guide to regression discontinuity. *MDRC*, 2012a.
- Robin Tepper Jacob, Pei Zhu, Marie-Andre, Somers, and Howard Bloom. A practical guide to regression discontinuity. *MDRC Publicatoins*, July 2012b. URL <https://www.mdrc.org/publication/practical-guide-regression-discontinuity>.

- Wilbert Van Der Klaauw. Regression–discontinuity analysis: A survey of recent developments in economics. *LABOUR*, 22(2):219–245, 2008. URL <http://citeseerx.ist.psu.edu/viewdoc/download?doi=10.1.1.466.956&rep=rep1&type=pdf>.
- David S. Lee and Thomas Lemieux. Regression discontinuity designs in economics. *Journal of Economic Literature*, 48:281–355, June 2010. URL <https://www.princeton.edu/~davidlee/wp/RDDEconomics.pdf>.
- Steven L. Scott. Multi-armed bandit experiments in the online service economy. *Applied Stochastic Models in Business and Industry*, 31(1):37–45, 2015.
- Student. The Lanarkshire milk experiment. *Biometrika*, 23(2/3):398–406, 1931.
- D. L. Thistlethwaite and D. T. Campbell. Regression-discontinuity analysis: An alternative to the ex post facto experiment. *Journal of Educational psychology*, 51(6):309, 1960.
- C. F. Jeff Wu and Michael S. Hamada. *Experiments: planning, analysis, and optimization*. John Wiley & Sons, New York, 2011.