The SAGE Encyclopedia of Educational Research, Measurement, and Evaluation

Regression Discontinuity Analysis

Contributors: Mitchell Campbell & Markus Brauer
Edited by: Bruce B. Frey
Book Title: The SAGE Encyclopedia of Educational Research, Measurement, and Evaluation
Chapter Title: "Regression Discontinuity Analysis"
Pub. Date: 2018
Access Date: February 26, 2018
Publishing Company: SAGE Publications, Inc.
City: Thousand Oaks,
Print ISBN: 9781506326153
Online ISBN: 9781506326139
DOI: http://dx.doi.org/10.4135/9781506326139.n581
Print pages: 1388-1392

Regression discontinuity analysis is a statistical tool that allows researchers to examine the effectiveness of the treatment in such studies. Consider the following: One researcher wants to determine whether tutoring underachieving middle school students improves their math grades; another wonders whether providing financial aid to low-income students has the desired effects on student success and dropout rates; and a third hopes to assess the effectiveness of special support programs for promising high school athletes. These studies fit the definition of a regression discontinuity design, whereby participants who satisfy a chosen criterion are assigned to a certain treatment and some outcome variable is measured later.

Regression discontinuity analysis is used for studies in which participants are assigned to treatment conditions based on a known assignment rule rather than randomly being assigned to conditions. Researchers or practitioners define an a priori *cutoff point* (Z_0) for participants' scores on an assignment variable (Z). Participants below the cutoff point receive the treatment, whereas those above the cutoff point do not (or vice versa). Participants are thus divided into groups defined by a dichotomous treatment variable (X). At a later point, the researchers measure the relevant outcome variable (Y). The goal of the regression discontinuity analysis is to determine whether the treatment has the desired effect on the outcome variable.

The Standard Model

To make these ideas more concrete, the following example will run through this text. In this hypothetical study, the assignment variable is a student's score on a standardized test taken in 10th grade, the treatment is whether the student is enrolled in a standardized test prep class, and the focal outcome measure is "self-efficacy," the student's belief in the student's ability to improve the student's standardized test performance. The school provides the test prep class to students scoring in the lowest 30% on the 10th-grade test. The data from this hypothetical example are displayed in Figure 1.

Figure 1 Students' self-efficacy scores increased significantly as a result of the test prep class

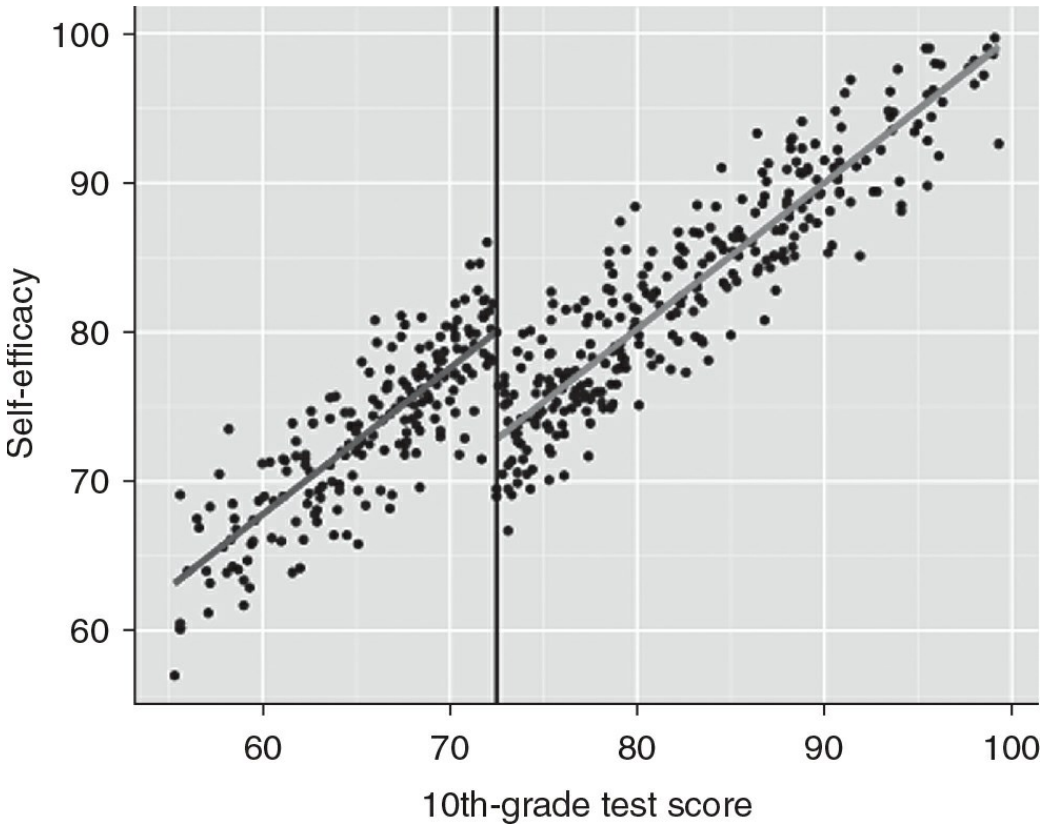

One key to understanding this type of analysis is noting that over and above the effect of the treatment variable, there is a relationship between the assignment variable and the outcome variable. In the given example, even if the test prep class is effective, it is quite possible that the students who attended the class will have lower self-efficacy scores on average than those who did not, simply because they started out at a lower level at the outset of the study. The question is whether the students who receive the class will have higher self-efficacy than would be predicted based on their 10th-grade test scores.

A regression discontinuity design is analyzed as follows: The outcome variable (*Y*, self-efficacy) is regressed on the treatment variable (*X*, attending the test prep class or not) and the assignment variable (*Z*, 10th-grade test score). One thus obtains the following regression equation:

$$Y = b_0 + b_1 X + b_2 Z + e.$$

If the coefficient b_1 is statistically significant, the data suggest that the treatment has an effect on the outcome variable. On the graph, this treatment effect will manifest as a vertical discrepancy between the two parallel regression lines. In the given example, the self-efficacy scores of the students in the test prep class were higher than would be expected based on their 10th-grade test scores. Because the treatment is dichotomous, the treatment effect is exactly equal to the coefficient b_1. Students who attended the test prep class had self-efficacy scores 7.5 points higher as a result of taking the class.

Curvilinear Relationships

Aside from other core model assumptions (discussed elsewhere in this volume), linearity is exceptionally important in estimating the treatment effect without bias. If the relationship

between the assignment and outcome variables is not linear, the b_1 coefficient will not represent the treatment effect accurately. For example, imagine a data set in which there is no treatment effect and a curvilinear relationship between the assignment variable and the outcome variable (see Figure 2a). The data could be accurately described with the following model:

$$Y = b_0 + b_1 Z + b_2 Z^2 + e.$$

Figure 2 (a) There is a curvilinear relationship between students' test scores and their self-efficacy scores (b) Parallel lines fit to curvilinear relationship misestimate treatment effect

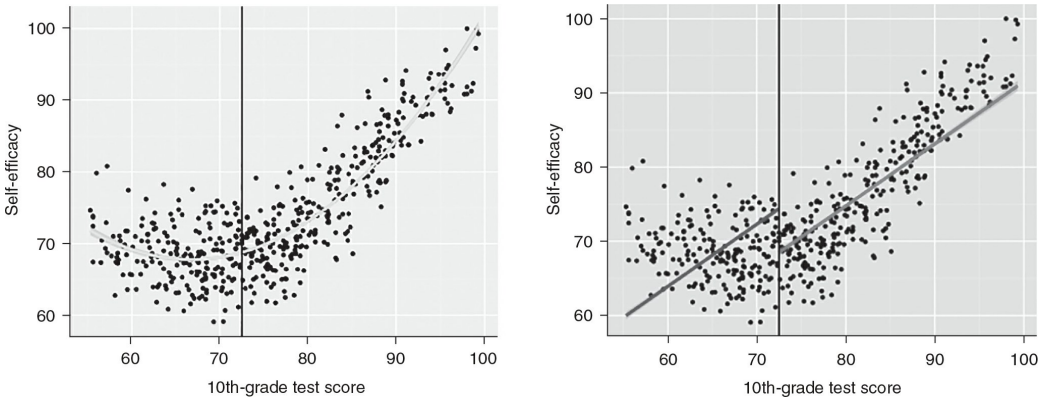

The treatment has no effect here, so the researchers should find a coefficient of 0 if they add treatment as a third predictor to this model. When the researchers ignore this curvilinear relationship and analyze the data with the model described in Equation 1, it is possible to obtain a significant treatment effect (see Figure 2b). However, this effect is due entirely to the fact that straight regression lines are being fitted to a curved data pattern.

When theory and prior studies suggest that there is a curvilinear relationship between assignment variable and outcome variable, it is advised to add a quadratic term to the model described in Equation 1. The full model would then be

$$Y = b_0 + b_1 X + b_2 Z + b_3 Z^2 + e.$$

Like before, the coefficient b_1 represents the treatment effect over and above the (linear and quadratic) effect of the assignment variable.

Interactions Between Treatment and Outcome

In certain cases, the treatment's effectiveness depends on individuals' scores on the assignment variable. Two cases are common: (1) individuals with scores on the assignment variable close to the cutoff point benefit *less* from the treatment (see Figure 3a) and (2) individuals with scores on the assignment variable close to the cutoff point benefit *more* from the treatment (see Figure 3b). Both cases are problematic for the classic regression discontinuity model, which forces the two regression lines representing the model predictions to be parallel. The model is thus misspecified. In addition, certain observations may have large residuals that decrease the statistical power to detect a treatment effect.

Figure 3 (a) Students scoring further from the cutoff point benefit more from the

treatment than those who score close to the cutoff point. (b) Students scoring closer to the cutoff point benefit more from the treatment than those who score further from the cutoff point

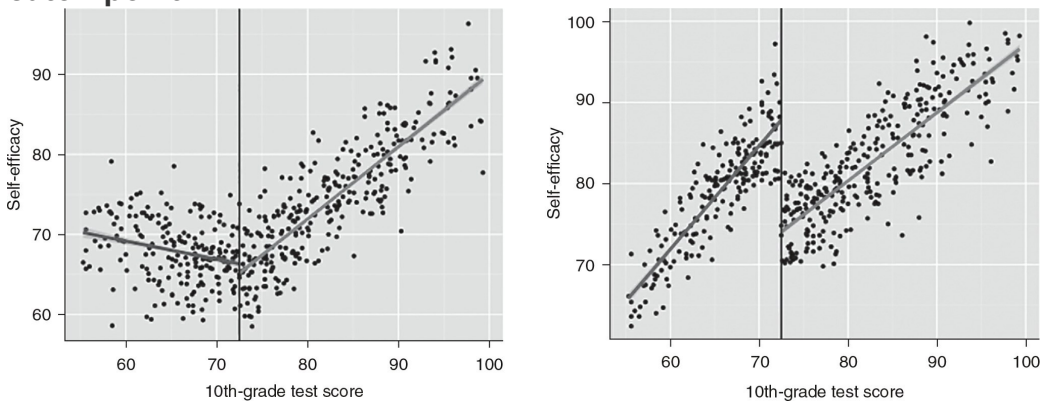

The solution is to estimate an interactive model in which the *Y*–*Z* relationship is allowed to vary between the treated and the untreated groups. This can be achieved with the following model:

$$Y = b_0 + b_1 X + b_2 Z + b_3 XZ + e.$$

If the coefficient b_3 is statistically significant, the data suggest that the relationship between the assignment variable and the outcome variable is not the same in the two groups. Like in every interactive model, b_1 represents the treatment effect for a participant with a score of 0 on the assignment variable. If the assignment variable in this example were included in its raw form (i.e., uncentered), the coefficient b_1 would estimate the treatment effect for a student with a score of 0 on the 10th-grade test. Clearly, this coefficient, and its associated *F*- and *p* values, would be rather meaningless.

To address this issue, many texts on regression discontinuity analysis suggest centering the assignment variable around the cutoff point by subtracting the value of the cutoff point (here 71) from every student's score on the assignment variable. Now, b_1 represents the treatment effect for a student with a 10th-grade test score of 71. Note that this effect is *not* the average treatment effect, making this a suboptimal approach. This approach will lead researchers to underestimate the average treatment effect when individuals whose scores on the assignment variable are close to the cutoff point benefit comparatively less from the treatment (Figure 3a) and overestimate the average treatment effect when individuals whose scores on the assignment variable are close to the cutoff point benefit comparatively more from the treatment (Figure 3b).

A better data-analytic strategy is to center the assignment variable around the average score in the treatment group (here 65). With this form of centering, the coefficient b_1 will represent the treatment effect for the typical person within the treatment group, accurately reflecting the average treatment effect. Regardless of the type of centering that is done with the assignment variable, the coefficient b_3 indicates whether the relationship between the assignment and outcome variables is different in the treatment and no treatment conditions.

In practice, it is virtually impossible to distinguish the case in which there is a curvilinear relationship between assignment variable and outcome variable and no treatment effect (Figure 2a) and the case in which there is a linear relationship between assignment variable

and outcome variable, an average treatment effect, and an assignment variable by treatment interaction caused by the fact that the treatment is less effective for participants with scores close to the cutoff point (Figure 3a). Both the Polynomial Model 2 and the Interactive Model 4 will fit the data quite well. Depending on the spread of the scores on the assignment variable, the researchers may be able to demonstrate that there is no evidence for curvilinearity among individuals in the untreated group. They may attempt to demonstrate that the interactive model fits the data better: showing that it has a smaller sum of squared errors, fewer outliers, and violates fewer model assumptions or by conducting a log-likelihood test showing that the observed results are more likely under the interactive hypothesis than under the curvilinear hypothesis. But researchers should be aware that the two interpretations are hard to distinguish empirically in a given data set, and ultimately they have to use theoretical arguments and refer to prior studies if they end up favoring one interpretation over the other. For example, in the hypothetical example, it makes little sense that individuals with very low scores on the 10th-grade test would score more highly on the self-efficacy measure if the test prep class had no effect. The researchers could argue that the interactive model is more logical than the quadratic model.

Statistical and Practical Considerations

The researchers might imagine that students' scores will also be affected by factors like parents' educational level or motivation to attend college. Covariates like these can simply be added to the regression equation. If a covariate coefficient is significant, it indicates that there is a relationship between the covariate and the outcome over and above the effect of the treatment and the assignment variable.

If the researchers are interested in exploring the effects of the treatment on individuals who are not part of the focal treatment group, they may choose to use a probabilistic assignment rule to decide who receives the treatment. This method contains elements of both a known assignment rule and random assignment: A certain proportion of participants on either side of the cutoff is given the treatment. For example, a researcher may decide that half of the students who scored in the bottom 40% on the test and one sixth of those in the upper 60% are randomly chosen for the test prep class.

Note that it is impossible to use a dichotomous variable as the assignment variable because such a variable will be perfectly confounded with the treatment variable: For example, in a study in which gender is the assignment variable, it would be impossible to say whether the observed differences are driven by treatment or gender.

The conclusion validity of studies with a regression discontinuity design is lower than that of randomized experiments, and a much larger sample size is required to achieve the same level of statistical power for two reasons. First, finding an effect of the treatment on the outcome variable over and above the effect of the assignment variable is difficult given that the treatment and assignment variables are, by definition, highly correlated. When the effects of multiple correlated predictors are estimated, the standard errors of regression coefficients are large, resulting in lower statistical power. Second, in practice, the treatment and comparison groups tend to be very different in size (e.g., people with an IQ over 150, families in the bottom 10% of household income). This imbalance between groups also decreases the power of the analysis.

The conclusions that can be drawn from results of studies with a regression discontinuity design are limited in scope because any treatment effects can only be assumed to hold for

the treatment group and often not all alternative explanations can be ruled out. Including relevant covariates can help reveal the true effect of a treatment on the outcome of interest: If the relationship between the test prep class and self-efficacy persists when controlling for parents' education level, the researchers can have more confidence that the treatment is having the observed effect.

A number of ethical and practical concerns make research utilizing regression discontinuity designs rare or challenging. First, universal application of an assignment rule is difficult and, in some cases, unethical. By setting a cutoff point, a researcher ultimately decides who is deserving of a given treatment. Perhaps a student who scored just above the cutoff point for receiving the test prep class will lose motivation to go to college. Furthermore, when a set cutoff point is made public, it begins to lose its meaning (e.g., people lying on their taxes to qualify for government programs). As a result, the scores on the assignment variable may contain a lot of error, hampering one's ability to reach accurate conclusions. Finally, in many situations where a regression discontinuity design is being used, a randomized experiment would be more effective, depending on the questions the researchers are interested in answering. If the school is hoping to assess whether the test prep class should be made mandatory for all students, a randomized experiment would make more sense than a regression discontinuity design because conclusions from such a study can be assumed to hold for all individuals in the population of interest.

The regression discontinuity design belongs to the family of the so-called quasi-experimental designs. Other designs in this family are the nonequivalent control group design and the interrupted time series design. Like other quasi-experimental designs, the regression discontinuity design has less internal validity and less conclusion validity than a randomized experiment. However, it allows researchers to draw causal conclusions with greater confidence than a post-only correlational design or a simple pretest–posttest design. Yet, some researchers doing field research employ these latter designs when random assignment is not feasible, perhaps unaware that better alternatives, such as the regression discontinuity design, are available to them.

See alsoCorrelation; Multicollinearity; Power Analysis; Quasi-Experimental Design

Mitchell CampbellMarkus Brauer
http://dx.doi.org/10.4135/9781506326139.n581
10.4135/9781506326139.n581

Further Readings

Cook, T. D., & Campbell, D. T. (1979). Quasi-experimentation: Design & analysis issues for field settings. Boston, MA: Houghton Mifflin.
Glover, S., & Dixon, P. (2004). Likelihood ratios: A simple and flexible statistic for empirical psychologists. Psychonomic Bulletin & Review, 11(5), 791–806.
Judd, C. M., & Kenny, D. A. (1981). Estimating the effects of social intervention. Cambridge, UK: Cambridge University Press.
Shadish, W. R., & Cook, T. D. (2009). The renaissance of field experimentation in evaluating interventions. Annual Review of Psychology, 60, 607–629. doi:http://dx.doi.org/10.1146/annurev.psych.60.110707.163544