

Can Close Election Regression Discontinuity Designs Identify Effects of Winning Politician Characteristics?



John Marshall  Columbia University

Abstract: *Politician characteristic regression discontinuity (PCRD) designs leveraging close elections are widely used to isolate effects of an elected politician characteristic on downstream outcomes. Unlike standard regression discontinuity designs, treatment is defined by a predetermined characteristic that could affect a politician's victory margin. I prove that, by conditioning on politicians who win close elections, PCRD estimators identify the effect of the specific characteristic of interest and all compensating differentials—candidate-level characteristics that ensure elections remain close between candidates who differ in the characteristic of interest. Avoiding this asymptotic bias generally requires assuming either that the characteristic of interest does not affect candidate vote shares or that no compensating differential affects the outcome. Because theories of voting behavior suggest that neither strong assumption usually holds, I further analyze the implications for interpreting continuity tests and consider if and how covariate adjustment, bounding, and recharacterizing treatment can mitigate the posttreatment bias afflicting PCRD designs.*

Verification Materials: The data and materials required to verify the computational reproducibility of the results, procedures, and analyses in this article are available on the *American Journal of Political Science* Dataverse within the Harvard Dataverse Network, at: <https://doi.org/10.7910/DVN/4MZQYH>.

Regression discontinuity (RD) designs have become a staple of the quantitative social scientist's methodological toolkit. RD designs leverage treatment assignments that change discontinuously at a known threshold in a forcing variable to identify treatment effects for observations around that threshold (see Cattaneo and Titiunik 2022). Although external validity can be limited, such designs are often regarded as the observational method best with the greatest internal validity. As researchers have increasingly focused on estimating causal effects, the use of RD designs in political science has exploded over the last decade (de la Cuesta and Imai 2016).

A particularly popular version of the RD design uses close elections to estimate effects of a specific characteristic of elected politicians on downstream electoral, policy,

and constituent outcomes. I will call this application a politician characteristic regression discontinuity (PCRD) design. Studies from across the globe have used PCRD designs to compare narrowly elected politicians who differ in terms of a given predetermined characteristic, usually with the objective of holding observable and unobservable confounders constant.¹ Table A1 in the Supporting Information (p. 2) lists 126 published articles—often in prestigious journals—that estimate downstream effects of ascriptive characteristics (gender, race or ethnicity, clan, religious identity), prior actions of politicians (criminal history, prior incumbency, seniority), labels politicians sort into (party membership, ideology), and institutional status (partisan alignment with other levels of government, term limit status). Indeed, PCRD designs appear to facilitate opportunities to study how

John Marshall, Department of Political Science, Columbia University, International Affairs Building, 420 West 118th Street, New York, NY 10027 (jm4401@columbia.edu).

I thank Jake Bowers, David Broockman, Naoki Egami, Anthony Fowler, Adam Glynn, Don Green, Andy Hall, Shigeo Hirano, Horacio Larreguy, Chris McConnell, Carlo Prato, Pablo Querubin, Alberto Simpser, Tara Slough, Jim Snyder, Jörg Spenkuch, Dan Thompson, Yamil Velez, Yiqing Xu, three anonymous reviewers, and attendees at EGAP and PolMeth presentations for helpful feedback and illuminating discussions.

¹The estimand is rarely stated explicitly. Most studies imply that PCRD designs identify “all else equal” effects of a particular characteristic. The same approach has been used for primary elections.

American Journal of Political Science, Vol. 68, No. 2, April 2024, Pp. 494–510

©2022, Midwest Political Science Association.

DOI: 10.1111/ajps.12741

electoral selection affects representation, accountability, and participation that are limited only by a researcher's capacity to measure politician characteristics of interest.²

Although the appeal of credibly estimating effects of winning candidate characteristics is obvious, whether PCRD designs can isolate the effect of a specific X —the characteristic, or bundle of characteristics, of interest—among politicians in close elections is not. Indeed, this article demonstrates that a nonstandard feature of this application of the RD design causes PCRD designs to generally identify compound treatment effects, rather than the local average treatment effect (LATE) of X . I will show that this confounding can only be avoided by imposing strong—and often implausible or unverifiable—additional assumptions.

The source of bias emanates from the difference between standard RD and PCRD designs. Standard RD designs define treatment as falling above or below a threshold. For example, close elections have been used to vary whether a candidate or party was *elected* to estimate financial returns to holding office (e.g., Eggers and Hainmueller 2009) and incumbent party electoral advantages (e.g., Lee, Moretti, and Butler 2004). In contrast, the treatment in PCRD designs—which instead seek to estimate the LATE of an elected politician characteristic—is defined by possessing (or not) predetermined characteristic X , *conditional on narrowly winning an election*. Beyond targeting different estimands, the mechanics of PCRD designs differ from standard RD designs in two important ways. First, as Sekhon and Titiunik (2012) have noted, close elections do not as if randomly assign characteristic X . The potential correlation between X and other politician characteristics creates the risk of confounding or necessitates reinterpreting the estimand as a compound treatment. Second, and more subtly, restricting attention to close elections entails conditioning on candidate vote shares that may be affected by X . As I show below, the former difference is fairly frequently acknowledged by researchers, but the latter is largely unrecognized. This article focuses on the second issue, showing that PCRD designs generally introduce bias—even when X is independent of other predetermined variables and the weak continuity assumption underpinning standard RD designs holds.

By expressing PCRD designs within the continuity framework of standard RD designs, I show that conditioning on close elections between candidates who differ in terms of characteristic X causes PCRD estimators to identify the LATE of electing a candidate of type X *combined with a (differential-weighted) LATE of*

any compensating differentials. Although PCRD designs ensure continuity across the districts different types of candidates are elected from, the vector of compensating differentials Z that generates this (asymptotic) bias is defined by characteristics of individual candidates that (a) the researcher regards as theoretically distinct from X ; and (b) ensure that candidates of type X remain in close elections with candidates not of type X . For example, in seeking to isolate the effect of gender, competence would be a compensating differential if women in close elections were more competent than men in close races because voters were biased against women.

My main identification result establishes that, even when X is (unconditionally) independent of Z , PCRD designs require strong additional assumptions to isolate the effect solely attributable to characteristic X . Specifically, identification requires that—at the discontinuity—either (i) X does not affect the winning candidate's victory margin or (ii) no compensating differential in Z affects the outcome of interest Y . These assumptions are analogous to the conditions under which the bias associated with conditioning the sample on a posttreatment variable disappears (see Hernán and Robins 2011, Montgomery, Nyhan, and Torres 2018). Where neither additional condition holds, compensating differentials cause PCRD designs to underestimate (overestimate) the LATE of X when the net effect of Z affects the candidate's vote share and the outcome Y in the same (opposite) direction.

I highlight three implications for applied research. First, to claim that PCRD estimates can isolate the effect of X by design, researchers should explicitly state and support one of the additional assumptions just described. However, these assumptions are difficult to empirically substantiate and are theoretically implausible when voters observe X and believe it will affect outcomes they care about. Second, in the likely event that neither assumption can be sustained, strategies for mitigating threats to internal validity vary in their effectiveness. Whereas PCRD estimates that reject a null hypothesis could be combined with candidate-level (dis)continuity tests to bound the LATE of X , covariate adjustment strategies cannot generally prevent biases induced by posttreatment conditioning. Indeed, PCRD designs do not imply that candidate-level covariates should be continuous at the point of discontinuity. Third, researchers might consider reinterpreting PCRD estimates as capturing (weighted) effects of X and Z . Where plausible compensating differentials can be measured, candidate-level discontinuity tests can now help to interpret this compound treatment. However, by focusing on a less well-defined and possibly heterogeneous conception of treatment, researchers cannot isolate the effect of

²Although potentially relevant, multiple testing is not this article's primary focus.

X —which is often desirable from a theoretical or policy perspective.

By clarifying the interpretation of PCRD designs, this article makes several contributions. First, I provide the first systematic account of the challenges that arise when RD designs define treatment by a variable that can also affect the forcing variable. Although some articles have noted study-specific issues relating to how conditioning on close elections introduces compensating differentials (e.g., Gagliarducci and Paserman 2012), 118 of 126 published studies using PCRD designs do not even loosely acknowledge the possibility of compensating differentials. Second, I move beyond proving the inconsistency of PCRD estimators by identifying strong additional assumptions under which the LATE of X can be isolated, establishing when PCRD designs underestimate and overestimate the LATE of X , and evaluating strategies to mitigate bias and reinterpret PCRD estimates. Third, this article illustrates the need to understand why candidates end up in close elections, and thus complements recent work emphasizing the theoretical implications of empirical models (e.g. Ashworth, Berry, and Bueno de Mesquita 2021; Bueno de Mesquita and Tyson 2020; Eggers 2017, Slough Forthcoming) and the large econometric literature documenting how sample selection generates bias (e.g., Heckman 1979).

The posttreatment bias this article highlights differs from prior critiques of RD designs leveraging close elections. Extant studies have examined other ways through which compound treatments can confound causal attribution, including where multiple treatments are assigned at the same threshold (Eggers et al. 2018), where correlated characteristics—such as Black politicians in the United States overwhelmingly being Democrats—are bundled together (e.g., Bucchianeri 2018; Ferreira and Gyourko 2014; Hall 2015), and where treatment affects downstream behaviors such as future candidacy decisions (Eggers 2017; Sekhon and Titiunik 2012). Researchers have also debated whether election outcomes are determined by chance at the discontinuity (Caughey and Sekhon 2011; Eggers et al. 2015) and highlighted the sensitivity of RD estimates and inference to bandwidth sizes and functional form assumptions (Cattaneo and Titiunik 2022; Gelman and Imbens 2019). However, the conceptual problems raised by this article still arise when the standard RD continuity assumption holds and enough data exist for consistent estimation of conditional expectations at the threshold.

PCRD Designs in Practice

This section builds intuition for the issues that arise when PCRD designs are used to isolate effects of elected politi-

cian characteristics. I first describe the design and its potential problems through the lens of two commonly-studied characteristics—gender and party affiliation. I then review published articles to characterize how PCRD designs have been used and summarize the identification concerns they address.

Electing Women


In my first example, researchers compare outcomes across polities where women and men were elected. Extant studies have used PCRD designs to estimate effects of electing women on policy priorities (e.g., Clots-Figueras 2011; Ferreira and Gyourko 2014), turnout among women and women running for office at future elections (e.g., Broockman 2014; Ferreira and Gyourko 2014), and government instability (Gagliarducci and Paserman 2012). Others have similarly examined the effect of women winning primary elections on general election results (Bucchianeri 2018). Many of these studies define their estimand as the effect of electing a woman instead of a man, often implicitly holding other characteristics of the candidate constant. For example, Ferreira and Gyourko (2014, 24) describe their design as estimating the “effect of gender” and Clots-Figueras (2011, 665) describes her design as estimating “the effect of a legislator’s gender.”

Before illustrating the identification problem at hand, it should be emphasized that defining gender as a treatment that is conceptually distinct from other candidate characteristics is challenging. This is because gender is often viewed as an inherently bundled treatment comprising various correlated features³; women who win close elections may espouse different policies, possess different qualifications, or have different personalities from men who win close elections. To isolate the effect of electing women, a researcher must distinguish the bundle of characteristics that differentiate women and men candidates—the definition of treatment—from the characteristics they regard as distinct from gender—the potential confounders. Gender is a particularly challenging example, but the need to explicitly define treatment applies equally to other characteristics—such as prior experience or partisan alignment with other relevant politicians—that may be easier to distinguish conceptually from potential confounders.

PCRD designs typically then estimate the effect of politician gender in single-member plurality races by comparing outcomes in “treated” districts where a woman was just elected in a race against a man with

³I abstract from whether ascriptive characteristics are manipulable.

FIGURE 1 Hypothetical Examples of PCRD Design Biases



Notes: For a given district d , $Y_d(X_d)$ denotes the potential outcome, $X_d \in \{0, 1\}$ denotes the characteristic of interest of the winning candidate, Δ_d denotes the winning margin of the candidate that possesses characteristic X , Z_d denotes the compensating differential of a winning candidate, and the function b captures the effect of Z_d on the district outcome. τ_{PCRD} denotes the true LATE of X_d and $\hat{\tau}_{PCRD}$ denotes a consistent estimate of the difference between the limits at $\Delta_d = 0$ captured by a typical PCRD estimator.

outcomes in “control” districts where a man was just elected in a race against a woman. Invoking the standard RD assumption of continuity in potential outcomes or local randomization (see Cattaneo and Titiunik 2022), it is usually argued that the two types of districts will, in expectation, be identical in terms of district-level covariates and all other individual-level characteristics of winning candidates at the point of discontinuity. If this were the case, only gender would change at the discontinuity and PCRD designs would identify the effect of electing a woman over an otherwise similar man. To ease exposition, I label the set of potential confounders—the bundle of characteristics that are conceptually distinct from gender—as “competence” and assume that voters prefer more competent candidates because they achieve positive district outcomes—employment in this example.

Using black to indicate the data researchers can observe, Figure 1a plots hypothetical conditional expectations of district employment as a function of a woman’s victory margin Δ_d . Cases to the right of the vertical line, where a woman is elected ahead of a man ($X_d = 1$), are treated. The relevant counterfactual for isolating the effect of gender, shown in gray, is *equally competent* men who win elections against women. The difference in employment at the discontinuity between districts where men and women were elected, τ_{PCRD} , is the LATE of electing a woman over a man in a close election.

Is it reasonable to assume *at the candidate level* that men who narrowly win are equally competent as women who narrowly win? Suppose that—holding competence fixed—voters are more likely to vote for men (e.g., Lawless 2015); this could arise from stereotyping, media attention, or differential support from political elites. To be in close races with women, men must then possess lower levels of competence in expectation than the women against whom they competed. The observed conditional expectation function for the men who won is shown in black to the left of the vertical line, where $b(Z_d|\Delta_d)$ denotes the reduction in employment due to elected men possessing less competence Z_d than elected women at a given vote margin. Even when each function in black is consistently estimated at $\Delta_d = 0$, the PCRD estimate $\hat{\tau}_{PCRD}$ is upwardly (asymptotically) biased in this example because it is confounded by competence—the compensating differential required for women to be in close races with men when voters are biased against women.

Electing Politicians from Different Parties

Analogous challenges apply to estimating differences due to a candidate’s party affiliation. Studies using PCRD designs compare outcomes between political

units that elected candidates from different parties (e.g., Gerber and Hopkins 2011; Lee, Moretti, and Butler 2004; Pettersson-Lidbom 2008). These studies describe the design as capturing the “effect of a Democratic victory” (Gerber and Hopkins 2011, 335) or “causal estimates of the effect of party control” (Pettersson-Lidbom 2008, 1037). This example focuses on the hypothetical effect of electing Democrats over Republicans on district employment.

The researcher must again specify what does and does not constitute part of the partisan affiliation treatment. To illustrate, I assume that party affiliation captures a common set of policy positions or ideology and that Democrats are more popular with voters on average. I again consider candidate “competence” as the compensating differential, and assume that more competent politicians increase employment. Other potential compensating differentials could include alignment with higher level incumbents or prior performance in office.

PCRD designs then compare districts where Democrats and Republicans barely won close elections against candidates from the other party. Because Democrat candidates are more popular than Republican candidates in this example, Republican candidates need to be more competent to counteract this disadvantage. By conditioning on close elections, PCRD designs then compare relatively incompetent Democrats with relatively competent Republicans. As Figure 1b illustrates, the PCRD estimate $\hat{\tau}_{\text{PCRD}}$ understates the effect of being a Democrat in this example because narrowly elected Democrats possess less competence Z_d than narrowly elected Republicans.

Limited Awareness of Compensating Differentials

To examine awareness of these potential issues, I used Google Scholar to identify 126 published articles employing PCRD designs.⁴ The earliest article was published in 2004, but 78% of articles have been published since 2015. Although 38% of studies have focused on executive and legislative elections in the United States, PCRD designs are also commonly applied to the election of individual politicians or changes in legislative representation or control in majoritarian and proportional representation systems in Asia, Europe, and South America. These articles have consistently appeared in prominent journals in political science and economics: 21% were published in the *American Journal of Political Science*, *American Political Science Review*, or *Journal of Politics*, whereas 5% were published in the *American Economic Review*,

Econometrica, the *Quarterly Journal of Economics*, or the *Review of Economic Studies*. According to Google Scholar, these studies had collectively amassed 11,774 citations by March 8, 2022.

After reading each article, I hand-coded whether the article demonstrated awareness of four potential threats to internal validity. Specifically, I coded whether an article (i) assessed continuity in potential outcomes by testing for discontinuities in district- or candidate-level covariates; (ii) assessed the same continuity assumption using density tests to test for sorting around the discontinuity; (iii) recognized that candidate characteristics may come as bundles due to unconditional correlations between characteristics; and (iv) discussed the risk of inducing or altering correlations between candidate characteristics by conditioning on close elections. The first two threats are benchmarks that apply to all RD designs (see Cattaneo and Titiunik 2022), the third could apply to other RD designs but is especially relevant for PCRD designs, and the fourth is specific to PCRD designs. Even brief and suggestive references to an issue were coded positively.

The results in Table 1 indicate that applied researchers are already aware of the importance of validating the continuity assumption and, to a lesser degree, that candidate characteristics come as bundles that are hard to separate. Panel A shows that 91% of articles conducted balance or continuity tests and 75% conducted density tests like the one proposed by McCrary (2008). Panel B shows that both strategies for validating the continuity assumption have become more prevalent over time. Furthermore, 33% of articles demonstrated awareness of the possibility that the candidate characteristic of interest might be unconditionally correlated with other characteristics. As panel D demonstrates, few studies acknowledged that candidates from different parties may also differ in terms of other candidate characteristics. However, the majority of studies seeking to estimate effects of candidate education, gender, ideology, incumbency status, and vocation discussed this issue.

In contrast, very few studies demonstrated any awareness of the issue motivating this article—the risk of inducing or altering correlations between the candidate characteristic of interest and compensating differentials by conditioning on close elections. Indeed, only 6% of articles even loosely mention this issue. These articles usually note that elections could be close because a difference in the characteristic of interest is counterbalanced by differences in terms of other characteristics; in five of the eight cases, this arises from the specific concern that women in close races may differ from men due to voter biases or campaign disadvantages. However, the origins and implications of such compensating

⁴These are listed in Table A1 (Supporting Information).

TABLE 1 PCRD Designs Demonstrating Awareness of Different Threats to Internal Validity

	Number of articles	Include covariate continuity tests	Include sorting/density tests	Demonstrate awareness of correlated characteristics	Demonstrate awareness of compensating differentials
Panel A: All articles					
All articles	126	115	94	42	8
Panel B: Articles by five-year period					
2002–2006	1	1	0	0	0
2007–2011	10	6	1	3	0
2012–2016	39	33	29	15	5
2017 to forthcoming	76	75	64	24	3
Panel C: Articles by region					
Africa	0	0	0	0	0
Asia	21	20	19	7	1
Europe	26	24	24	8	2
Middle East	4	4	4	1	0
North America	50	44	27	13	3
Oceania	0	0	0	0	0
South America	24	22	19	13	2
Cross-continental	1	1	1	0	0
Panel D: Articles by politician characteristic					
Partisan alignment across tiers of government	17	15	14	4	1
Criminal history	4	4	4	1	0
Education	4	3	2	2	1
Gender	24	23	18	17	5
Ideology	2	2	1	2	0
Incumbency, term limit status, or seniority	11	9	7	8	2
Partisan affiliation	58	54	42	5	0
Race, ethnicity, religion, or clan	7	6	6	4	0
Preoffice vocation	5	5	4	3	0
Panel E: Articles by type of electoral discontinuity					
Individual politician (executive or legislator)	101	92	73	36	8
Legislature majority	13	11	11	1	0
Legislature seat share	13	13	10	5	0
Party representation threshold	2	2	2	0	0

Notes: Data are based on the author's hand-coded classifications. Some articles fall into multiple categories.

differentials received limited discussion and were often quickly dismissed, despite the fact that compensating differentials between candidates will generally exist in PCRD designs—as this article demonstrates.

One article delves deeper into this issue. Gagliarducci and Paserman (2012, 1031) note that identifying the effect of elected women using a PCRD design requires that “the vote share of each candidate must not depend directly on gender.” They argue that this assumption is plausible in their study examining the effect of women narrowly elected as mayors on early municipal government termination in Italy, if voters do not select candidates to maintain government stability, are unaware that mayor gender affects government stability, or only select candidates on the basis of factors unrelated to gender that could affect government stability. This article generalizes the conditions under which PCRD designs identify the desired estimand, characterizes the nature of asymptotic bias when these strong conditions do not hold, and discusses bias mitigation strategies and alternative conceptualizations of treatment.

Theoretical Analysis

This section first recaps how standard RD designs work in the context of plurality elections in single member districts, before explaining how PCRD designs differ.⁵ I then show how a posttreatment bias introduced by these differences prevents PCRD designs from isolating the effect of the characteristic of interest in a stylized example. I finally provide general results demonstrating that additional assumptions—which are far stronger than the standard RD continuity assumption—are required to identify the effects often attributed to PCRD designs by applied researchers.

Standard RD Designs

In the close election application of RD designs, each candidate i in district d receives share $V_{id} \in [0, 1]$ of the votes cast between the top two candidates. The continuously distributed forcing variable is the difference between V_{id} and the vote share V_{jd} of the other most popular candidate $j \neq i$ in the

district: $\Delta_{id} := V_{id} - V_{jd} \in [-1, 1]$. The following treatment variable then indicates whether candidate i won the election in district d :⁶

$$T_{id} := \begin{cases} 1 & \text{if } \Delta_{id} > 0 \\ 0 & \text{if } \Delta_{id} \leq 0. \end{cases} \quad (1)$$

In addition to observing T_{id} based on which candidate wins the race, researchers also observe an outcome variable Y_{id} for each candidate. The potential outcome $Y_{id}(T_{id}) \in \mathbb{R}$ depends on whether a candidate wins office. This representation encodes the SUTVA assumption that i 's potential outcomes are not affected by the treatment status of other candidates and that there is a single version of treatment. Since only one potential outcome can be observed, the observed outcome is related to potential outcomes by $Y_{id} = T_{id}Y_{id}(1) + (1 - T_{id})Y_{id}(0)$.

The standard RD design requires the following weak continuity assumption:⁷

Assumption 1. Potential outcomes $Y_{id}(T_{id})$ satisfy:

- (a) *Continuity from above:* $\lim_{v \downarrow 0} \mathbb{E}[Y_{id}(1) | \Delta_{id} = v] = \mathbb{E}[Y_{id}(1) | \Delta_{id} = 0]$;
- (b) *Continuity from below:* $\lim_{v \uparrow 0} \mathbb{E}[Y_{id}(0) | \Delta_{id} = v] = \mathbb{E}[Y_{id}(0) | \Delta_{id} = 0]$.

This assumption states that, at the point of discontinuity, potential outcomes do not vary discontinuously in any way other than whether a given candidate won the election. In the case of close elections, this is plausible because factors exogenous to candidate characteristics, such as election day weather, generate random variation in which candidate wins; Eggers et al. (2015) provide evidence supporting this claim from 10 countries across the world.

When continuity holds, the LATE of T_{id} at the point of discontinuity—denoted by $\tau_{RD} := \mathbb{E}[Y_{id}(1) - Y_{id}(0) | \Delta_{id} = 0]$ —can be identified by comparing observed outcomes between candidates that narrowly won and narrowly lost. Applied researchers typically employ an RD estimator of the following form:

$$\hat{\tau}_{RD} = \hat{\mu}_+(0) - \hat{\mu}_-(0), \quad (2)$$

where $\hat{\mu}_+(s)$ and $\hat{\mu}_-(s)$ are estimators of $\lim_{v \downarrow s} \mathbb{E}[Y_{id} | \Delta_{id} = v]$ and $\lim_{v \uparrow s} \mathbb{E}[Y_{id} | \Delta_{id} = v]$, respectively. The state of the art involves estimating $\hat{\mu}_+(0)$ and $\hat{\mu}_-(0)$ using local polynomial regressions either side of the discontinuity and correcting for a consistent estimate of the misspecification bias that arises from approximating the unknown functional form of $\mathbb{E}[Y_{id} | \Delta_{id} = v]$ (Calonico,

⁵ Designs adapted to legislative chambers (Clots-Figueras 2011), proportional representation elections (Folke 2014), and to leverage discontinuities in control of legislative bodies (Pettersson-Lidbom 2008) differ in some respects. However, similar challenges apply to isolating effects of elected politician characteristics.

⁶ For simplicity, I assume that i does not win if $V_{id} = V_{jd}$.

⁷ The local randomization approach imposes stronger assumptions (Cattaneo and Titiunik 2022).

Cattaneo, and Titiunik 2014; Calonico et al. 2019). To trade off the finite sample biases and precision gains of including observations further from the discontinuity, researchers often use a data-driven procedure to select the bandwidth that minimizes the mean squared error of $\hat{\tau}_{RD}$. See Cattaneo and Titiunik (2022) for an excellent review of RD estimation and inference methods.

To focus on the asymptotic bias—the bias as the sample size tends to infinity—that arises with PCRD designs, this article abstracts from estimation challenges. Specifically, for a random sample of n elections drawn from a large population, I assume that

Assumption 2. For any conditioning set W , $\hat{\mu}_+(0|W)$ and $\hat{\mu}_-(0|W)$ are consistent estimators with bounded variance.

Following Hahn, Todd, and Van der Klaauw (2001) and Imbens and Lemieux (2008), it is well established that

Proposition 1. Under Assumptions 1 and 2, $\hat{\tau}_{RD}$ is a consistent and asymptotically unbiased estimator of τ_{RD} .

Proof. See Supporting Information Appendix B (pp. 3–5) for all proofs. \square

Within political science, this type of RD design has proved popular for estimating the consequences of *being elected*. One strand of this literature has explored the effect of being elected to office on a candidate's downstream wealth (e.g., Eggers and Hainmueller 2009). Another strand has studied the effect of winning elections on subsequent election outcomes. Since the decision to run in future elections often depends on winning a close election, researchers interested in isolating the effect of winning office, conditional on running, on future electoral success have focused on the electoral outcomes of parties linked to the winning and losing candidates where parties always run for office. This has spawned a large literature measuring party incumbency advantages. In the United States, narrowly elected incumbent parties are substantially more likely to win future elections in the same district (e.g., Lee, Moretti, and Butler 2004). These studies are united by comparing downstream outcomes *between election winners and losers*.

How PCRD Designs Differ from Standard RD Designs

Whereas close election applications of standard RD designs estimate effects of a politician getting elected, PCRD designs instead seek to isolate effects of a specific characteristic of elected politicians on downstream outcomes. Rather than compare winning and losing

candidates, PCRD designs compare politicians who all narrowly won elections in separate districts but differ according to a predetermined binary characteristic denoted by $X_{id} \in \{0, 1\}$.⁸ Characteristics of empirical interest have included gender, race, vocational experience, criminality, prior incumbency, partisan affiliation, and partisan alignment.

Isolating the effect of a characteristic of interest is difficult because politicians are defined by many characteristics that tend to be correlated. For example, elected women in the United States are more likely to be Democrats, politicians who have engaged in corruption are more likely to be aligned with higher level politicians, or politicians from traditional parties are more likely to be experienced. If the characteristic—or bundle of characteristics—of interest X_{id} is correlated with a vector of K distinct candidate-level characteristics $\mathbf{Z}_{id} \in \mathbb{R}^K$, any effect of X_{id} may be confounded by the effects of \mathbf{Z}_{id} . To define their target estimand, researchers must decide which characteristics should and should not be included in their treatment of interest; put differently, they must decide which characteristics are conceptually distinct from the characteristic of interest. This is an inexact science. Some researchers (explicitly or implicitly) assert a characteristic of interest and seek to substantiate the claim that only this characteristic drives any effect. Others define their treatment as a bundle of correlated characteristics that are distinct from some other characteristics.⁹

I abstract from the challenge of interpreting correlated characteristics, which is already acknowledged in a third of articles reviewed in Table 1. Rather, I show that PCRD designs introduce a form of posttreatment bias even when characteristic X_{id} is (unconditionally) independent of other relevant characteristics, e.g. if X_{id} was randomly assigned. Accordingly, I will at times assume that

Assumption 3. Characteristic X_{id} is independent of \mathbf{Z}_{id} and \mathbf{Z}_{jd} : $X_{id} \perp\!\!\!\perp \mathbf{Z}_{id}, \mathbf{Z}_{jd}$.

This assumption will clarify that biases emerge in PCRD designs even in a “best case scenario” where—at least among politicians who could end up in close races— X_{id} is independent of i 's conceptually distinct characteristics \mathbf{Z}_{id} and the conceptually distinct characteristics \mathbf{Z}_{jd} of their chief competitor j . However, Assumption 3 is relaxed for more general theoretical results.

⁸Nonbinary characteristics could compare any two characteristic discrete values or bins.

⁹I analyze the case where researchers view *all* characteristics as a single bundle in the “Expanding the Conception of Treatment” section.

In shifting attention to the *type* of politician who wins, the unit of analysis in PCR designs is the district. The district-level forcing variable is then $\Delta_d := V_{1d} - V_{0d} \in [-1, 1]$, where V_{1d} and V_{0d} respectively denote the vote shares of the most popular politician of type $X_{id} = 1$ and $X_{id} = 0$ in district d . Districts where the top two candidates are of the same type are excluded. The corresponding district-level treatment indicates whether a candidate of type $X_{id} = 1$ won the election:

$$X_d := \begin{cases} 1 & \text{if } \Delta_d > 0 \\ 0 & \text{if } \Delta_d \leq 0. \end{cases} \quad (3)$$

District-level potential outcomes depend on X_d , which reflects the individual-level potential outcomes of the type of politician who won: $Y_d(X_d) = X_d Y_{1d}(1) + (1 - X_d) Y_{0d}(1)$. For example, $Y_d(1)$ and $Y_d(0)$ could correspond to the district-level outcome when the elected candidate was a woman and a man, respectively. In PCR designs comparing observations of Y_d across districts, the politician-level potential outcome $Y_{id}(0)$ is neither relevant nor well defined because losing politicians do not enter office.

The LATE of interest in PCR designs is the difference in potential outcomes across districts with close elections where politicians of different types were elected: $\tau_{\text{PCR}} := \mathbb{E}[Y_d(1) - Y_d(0) | \Delta_d = 0]$. This is typically estimated using the following PCR estimator:

$$\begin{aligned} \hat{\tau}_{\text{PCR}} &= \lim_{v \downarrow 0} \mathbb{E}[\widehat{Y}_d | \Delta_d = v] - \lim_{v \uparrow 0} \mathbb{E}[\widehat{Y}_d | \Delta_d = v] \\ &= \hat{\mu}_+(0 | X_{id} = 1, X_{jd} = 0) - \hat{\mu}_+(0 | X_{id} = 0, \\ &\quad X_{jd} = 1). \end{aligned} \quad (4)$$

The second line rewrites district-level outcomes Y_d in terms of candidate-level outcomes Y_{id} to distinguish standard RD designs from PCR designs: whereas standard RD designs compare candidates either side of the threshold in the forcing variable, PCR designs compare narrow winners on one side of the threshold who differ in terms of X_{id} . Consequently, this nonstandard RD design conditions on a predetermined difference between X_{id} and X_{jd} that could also affect the forcing variable Δ_d . I next show how this distinction prevents $\hat{\tau}_{\text{PCR}}$ from consistently estimating τ_{PCR} , even when Assumptions 1–3 hold.

Bias in PCR Designs with a Single Compensating Differential

To build intuition, I start with a simple case where a single compensating differential $Z_{id} - Z_{jd}$ ensures that the race between candidates i and j in district d is close despite only one candidate possessing characteristic X_{id} . Let

characteristic X_{id} help candidate i win votes, for example, by being the incumbent, representing a popular party, or not suffering gender-based biases. The compensating differential will offset this advantage, for example, because the candidate for whom $X_{id} = 1$ is less competent, more malfeasant, or less politically connected than the candidate for whom $X_{id} = 0$. In addition to affecting candidate vote shares, the compensating differential $Z_{id} - Z_{jd}$ will also affect district-level outcomes that depend on the winning candidate's level of Z_{id} .

My stylized example captures these roles of X_{id} and $Z_{id} - Z_{jd}$ in the following functional forms:

$$V_{id} = \alpha \frac{X_{id} - X_{jd}}{2} + \beta \frac{Z_{id} - Z_{jd}}{2} + \frac{\varepsilon_{id} - \varepsilon_{jd}}{2}, \quad (5)$$

$$Y_{id}(1) = \tau X_{id} + \gamma Z_{id} + \nu_d, \quad (6)$$

where $\alpha \geq 0$ and $\beta > 0$ imply that possessing more of characteristic X_{id} or Z_{id} increases a candidate's vote share,¹⁰ whereas τ and γ , respectively, denote (constant) effects of X_{id} and Z_{id} on district-level outcomes. I further assume that $Z_{id} - Z_{jd} \sim N(0, \sigma_Z^2)$ and $\varepsilon_{id} - \varepsilon_{jd} \sim N(0, \sigma_\varepsilon^2)$ are normally distributed and drawn independently of both X_{id} and each other,¹¹ while ν_d is district-level noise that is drawn independently of all other variables. By embedding Assumption 3 in these distributional assumptions, any bias in $\hat{\tau}_{\text{PCR}}$ must emerge from the PCR design.

Derivation of Asymptotic Bias. Although X_{id} is independent of $Z_{id} - Z_{jd}$, PCR designs can generate a correlation by conditioning on close elections where two narrow winners with different characteristics obtain similar vote shares. To see why, note that the limiting case of close elections—where candidates within a given district receive equal numbers of votes—implies:

$$\Delta_d = \alpha + \beta(Z_{1d} - Z_{0d}) + \varepsilon_{1d} - \varepsilon_{0d} = 0, \quad (7)$$

where the electorally advantaged candidate of type $X_{id} = 1$ is denoted by $i = 1$ and the candidate of type $X_{id} = 0$ is denoted by $i = 0$. A tie between these candidate types occurs because there is a countervailing compensating differential ($Z_{1d} < Z_{0d}$) and because candidate 1 encountered unfortunate random shocks ($\varepsilon_{1d} < \varepsilon_{0d}$).

¹⁰ α and β are positive for simplicity, but need not be restricted. These coefficients are common across candidates because X_{id} and Z_{id} characterize differences between candidates.

¹¹ $V_{id} \in [0, 1]$ can be violated with normal distributions. However, the general results do not impose unbounded distributions, while $V_{id} \in [0, 1]$ almost always holds when σ_Z^2 and σ_ε^2 are small.

The PCRD estimator then converges to the following quantity:

$$\begin{aligned}
 \hat{\tau}_{\text{PCRD}} &\xrightarrow{p} \lim_{v \downarrow 0} \mathbb{E}[Y_{id} | \Delta_{id} = v, X_{id} = 1, X_{jd} = 0] \\
 &\quad - \lim_{v \downarrow 0} \mathbb{E}[Y_{id} | \Delta_{id} = v, X_{id} = 0, X_{jd} = 1] \\
 &= \mathbb{E}[Y_{id}(1) | \Delta_{id} = 0, X_{id} = 1, X_{jd} = 0] \\
 &\quad - \mathbb{E}[Y_{id}(1) | \Delta_{id} = 0, X_{id} = 0, X_{jd} = 1] \\
 &= \tau + \gamma \mathbb{E}[Z_{1d} - Z_{0d} | \alpha + \beta(Z_{1d} - Z_{0d}) \\
 &\quad + \varepsilon_{1d} - \varepsilon_{0d} = 0] + \mathbb{E}[v_d | \Delta_{id} = 0, \\
 &\quad X_{id} = 1, X_{jd} = 0] \\
 &\quad - \mathbb{E}[v_d | \Delta_{id} = 0, X_{id} = 0, X_{jd} = 1] \\
 &= \tau - \underbrace{\gamma \frac{\alpha \beta \frac{\sigma_Z^2}{\sigma_\varepsilon^2}}{1 + \beta^2 \frac{\sigma_Z^2}{\sigma_\varepsilon^2}}}_{\text{asymptotic bias relative to the LATE of } X_{id}}, \tag{8}
 \end{aligned}$$

where the first line follows from consistent estimation of conditional expectations (Assumption 2), the second line follows from continuity (Assumption 1(a)) and the district-level outcome being the potential outcome of the winning candidate in that district, the third line substitutes the functional forms for potential outcomes and vote shares, and the final line applies the distributional assumptions on $Z_{id} - Z_{jd}$, $\varepsilon_{id} - \varepsilon_{jd}$, and v_d .¹² Since PCRD designs compare candidates that win elections, only part (a) of Assumption 1 is needed. Although this weakens the standard RD continuity assumption, it is difficult to imagine contexts where part (a) held but part (b) did not.

This derivation shows that PCRD designs can induce a form of posttreatment bias that yields inconsistent estimates of the effect of characteristic X_{id} . Where X_{id} affects V_{id} , this is because the event $\Delta_{id} = 0$ upon which the LATE is conditioned is itself affected by X_{id} . Satisfying $\Delta_{id} = 0$, and thus $\Delta_d = 0$, when $X_{id} \neq X_{jd}$ generally requires the existence of a compensating differential $Z_{id} \neq Z_{jd}$, which can in turn affect district outcomes. Supporting Information Appendix C (pp. 5–6) shows that similar insights emerge where there are multiple (possibly correlated) compensating differentials.

The asymptotic bias can be avoided if one of the following three conditions holds. First, the PCRD estima-

tor consistently recovers τ_{PCRD} when no compensating differential is needed ($\alpha = 0$). This requires voters not to select politicians on the basis of X_{id} . Second, there is no asymptotic bias when the compensating differential does not affect the outcome ($\gamma = 0$). This occurs when $Z_{id} - Z_{jd}$ affects which candidate voters prefer, but Z_{id} does not ultimately affect the district outcome of interest. Third, $\hat{\tau}_{\text{PCRD}}$ consistently estimates τ_{PCRD} when the difference in vote share due to X_{id} is exactly offset by noise ($\alpha + \varepsilon_{1d} - \varepsilon_{0d} = 0$). This knife-edge condition is closely related to the previous condition, since ε_{id} are factors that affect candidate i 's vote share without affecting their behavior in office. It is only plausible where the signal to noise ratio $\frac{\sigma_Z^2}{\sigma_\varepsilon^2}$ is sufficiently low that candidate vote shares are predominantly determined by chance, rather than systematic differences in attributes. As I discuss below, theories of voting behavior suggest that none of these conditions usually holds.

How Do the Direction and Magnitude of Bias Vary?

The asymptotic bias of PCRD estimates can be upward or downward, depending on how X_{id} and Z_{id} affect district outcomes. Where the direction of the effect of each variable agrees—either $\tau, \gamma > 0$ or $\tau, \gamma < 0$ —Equation (8) shows that $\hat{\tau}_{\text{PCRD}}$ is downwardly biased in magnitude. PCRD designs will thus *underestimate* the effect of X_{id} when both candidate characteristics appeal to voters and both characteristics also lead to better (or worse) district-level outcomes once a politician enters office. This would occur where voters select candidates with characteristics that they correctly anticipate will produce better outcomes in office from the perspective of most voters, such as greater economic performance, security, or redistribution toward a majority group. Intuitively, $\hat{\tau}_{\text{PCRD}}$ captures a lower bound on the effect of X_{id} because differences in Y_d due to electing a candidate possessing desirable characteristic X_{id} are counteracted by electing a candidate possessing relatively less of the also desirable compensating differential Z_{id} . The partisan affiliation example above illustrates such a case.

PCRD designs instead *overestimate* the effect of X_{id} where the signs of τ and γ disagree. In the gender example above, opposing effects can arise when voters incorrectly believe that women will perform worse. Opposing effects could similarly occur if vote buying efforts win votes for candidates whose nonprogrammatic policies later reduce voter welfare.

Each direction of bias creates different challenges for hypothesis testing. Where underestimation occurs, the sign of a PCRD estimate that rejects the null hypothesis remains correct because $\hat{\tau}_{\text{PCRD}}$ is a lower bound. Conversely, a failure to reject the null hypothesis is relatively

¹²Since $Z_{id} - Z_{jd}$ and $\varepsilon_{id} - \varepsilon_{jd}$ are normally distributed, $\mathbb{E}[Z_{idk} - Z_{jdk} | \Delta_d] = \mathbb{E}[Z_{idk} - Z_{jdk}] + \frac{\text{Cov}[Z_{idk} - Z_{jdk}, \Delta_d]}{\mathbb{V}[\Delta_d]} (\Delta_d - \mathbb{E}[\Delta_d])$. In this application, $\mathbb{E}[Z_{id} - Z_{jd}] = 0$, $\text{Cov}[Z_{idk} - Z_{jdk}, \Delta_d] = \beta \sigma_Z^2$, $\mathbb{V}[\Delta_d] = \beta^2 \sigma_Z^2 + \sigma_\varepsilon^2$, and $\Delta_d - \mathbb{E}[\Delta_d] = -\alpha$. By independence, $\mathbb{E}[v_d | \Delta_{id} = 0, X_{id} = 1, X_{jd} = 0] = \mathbb{E}[v_d | \Delta_{id} = 0, X_{id} = 0, X_{jd} = 1]$.

uninformative because it is consistent with both X_{id} not affecting district outcomes and a positive effect of possessing X_{id} cancels out with a negative effect of possessing relatively less of Z_{id} . These challenges operate in reverse where overestimation occurs.

The second term in Equation (8) also illustrates when the asymptotic bias is greatest. The magnitude of the bias increases in γ because the compensating differential has a larger effect on district outcomes, and similarly increases in α because more of the compensating differential is required to ensure close elections. The bias also increases in $\frac{\sigma_z^2}{\sigma_\varepsilon^2}$, as noise becomes relatively less important than candidate characteristics in determining whether an election is close. The influence of β on the size of bias is ambiguous because a greater impact of compensating differentials on candidate vote shares reduces the importance of noise in generating close elections but also reduces the size of the compensating differential needed to overcome the difference in X_{id} .

Bias in PCRD Designs in General

To demonstrate these results more generally, I relax the functional form and distributional assumptions imposed on V_{id} and Y_{id} . First, candidate i 's vote share is now an unrestricted function $v(X_{id}, X_{jd}, \mathbf{Z}_{id}, \mathbf{Z}_{jd}, \varepsilon_{id}, \varepsilon_{jd})$, where the realization of two independent and identically distributed shocks $(\varepsilon_{id}, \varepsilon_{jd})$ is independent of all other variables. Second, I assume that potential outcomes are additively separable between X_{id} and \mathbf{Z}_{id} , but otherwise allow effects of these variables to vary across districts:

Assumption 4. *Candidate i 's potential outcome if elected to office is $Y_{id}(1) = \tau_d X_{id} + g(\mathbf{Z}_{id}) + \nu_d$, where ν_d is distributed independently of all other variables.*

Additive separability excludes the possibility that the effect of X_{id} varies with compensating differentials \mathbf{Z}_{id} . This assumption is not necessary for asymptotic bias to emerge in PCRD designs, but facilitates a simple decomposition of the bias.

The following proposition establishes the quantity that the PCRD estimator converges to:

Proposition 2. *Under Assumptions 1(a), 2, and 4:*

$$\hat{\tau}_{\text{PCRD}} \xrightarrow{p} \tau_{\text{PCRD}} + \mathbb{E}[g(\mathbf{Z}_{id})|\Delta_{id} = 0, X_{id} = 1, X_{jd} = 0] - \mathbb{E}[g(\mathbf{Z}_{id})|\Delta_{id} = 0, X_{id} = 0, X_{jd} = 1].$$

This result demonstrates that PCRD designs generally yield asymptotically biased estimates of the LATE of X_{id} . This bias emerges when \mathbf{Z}_{id} is unconditionally correlated with X_{id} and correlations with \mathbf{Z}_{id} are induced or altered

by conditioning on close elections when compensating differentials are required for elections to be tied.

To focus on the asymptotic bias introduced by posttreatment conditioning, I impose Assumption 3. Extending the intuitions from the example with a single compensating differential, the following proposition establishes three sufficient conditions for PCRD designs to identify and consistently estimate the LATE of X_{id} :

Proposition 3. *Under Assumptions 1(a)–4, $\hat{\tau}_{\text{PCRD}} \xrightarrow{p} \tau_{\text{PCRD}}$ if one of the following conditions holds:*

- (i) $V_{id} \perp\!\!\!\perp X_{id}, X_{jd}$ among candidates that could enter close races;
- (ii) $\mathbb{E}[g(\mathbf{Z}_{id})|\Delta_{id} = 0, X_{id} = 1, X_{jd} = 0] = \mathbb{E}[g(\mathbf{Z}_{id})|\Delta_{id} = 0, X_{id} = 0, X_{jd} = 1]$;
- (iii) whenever $v(1, 0, \mathbf{z}, \mathbf{z}', \varepsilon_{id}, \varepsilon_{jd}) = v(0, 1, \mathbf{z}', \mathbf{z}, \varepsilon_{jd}, \varepsilon_{id})$, $\mathbf{z} = \mathbf{z}'$.

Similarly, the following identification result holds under one of the three preceding conditions:

$$\begin{aligned} &\mathbb{E}[Y_d(1) - Y_d(0)|\Delta_d = 0] = \\ &\lim_{\nu \downarrow 0} \mathbb{E}[Y_{id}|\Delta_{id} = \nu, X_{id} = 1, X_{jd} = 0] \\ &- \lim_{\nu \downarrow 0} \mathbb{E}[Y_{id}|\Delta_{id} = \nu, X_{id} = 0, X_{jd} = 1]. \end{aligned}$$

Condition (i) amounts to ensuring that candidate vote shares are not correlated with X_{id} and X_{jd} ; this means that compensating differentials are not required for elections to be close. Condition (ii) allows for compensating differentials to counteract the electoral advantage of X_{id} if they do not affect district outcomes in expectation. This condition encompasses two cases: no compensating differential affects Y_d ; or the net effect of compensating differentials exactly canceling out. Condition (iii) similarly ensures that noise—factors that benefit a candidate which are not features of the candidate themselves—exclusively compensates for the electoral advantage of X_{id} in close races.

Implications for Applied Research

The preceding analysis demonstrated that PCRD designs seeking to isolate effects of an elected politician characteristic require that researchers invoke stronger additional assumptions or accept post-treatment bias (and possibly correlated characteristics bias). This section explores the implications for applied research, starting by considering the viability of imposing one of the conditions in Proposition 3. Since these conditions are unlikely to hold in many applications, I then consider strategies to mitigate the threat to internal validity. I finally discuss the

implications of redefining treatment to encompass X_{id} and all compensating differentials.

Invoking an Additional Assumption

Perhaps the most appealing method for addressing the inconsistency of PCRD estimates is to explicitly invoke and substantiate one of the conditions in Proposition 3. Beyond imposing Assumption 3 around the discontinuity, this entails assuming—at the discontinuity—that X_{id} does not affect candidate vote shares (condition (i)), that compensating differentials induced by variation in X_{id} do not affect the outcome of interest (condition (ii)), or that compensating differentials are not required because idiosyncratic electoral shocks counterbalance compensating differentials (condition (iii)). I focus on conditions (i) and (ii). Condition (iii) is similar to condition (ii) in claiming that other factors affecting election results do not shape postelection district outcomes; moreover, this knife-edge condition almost surely fails to hold when condition (i) does not hold. Unfortunately, as I next explain, neither condition (i) nor (ii) is easily validated and both assumptions conflict with theories of voting behavior.

Empirical Challenges. Where compensating differentials are observed or assumed to exist, showing that no compensating differential affects the outcome of interest is particularly difficult. First, strong support for condition (ii) requires that a researcher further show that *each* compensating differential has no effect on the outcome around the discontinuity. Finding identification strategies for all potential compensating differentials—or even just the most plausible compensating differentials—is unlikely to be feasible. Moreover, if the effects of compensating differentials are heterogeneous with respect to V_{id} , this challenge is exacerbated by the need for these estimates to be local to close elections. Second, because compensating differentials like candidate competence are often difficult to measure, it is hard to confidently claim that unobserved compensating differentials are not affecting the outcome.

Validating that characteristic X_{id} does not affect candidate vote shares around the discontinuity is more attainable. Since X_{id} and V_{id} are both observed, a single test demonstrating that X_{id} does not affect V_{id} among candidates around the discontinuity can support condition (i). A compelling test showing that candidate i 's gender or party affiliation does not affect their vote share requires an additional research design exogenously varying X_{id} among candidates that end up in close races. At the expense of external validity, conjoint experiments

(Hainmueller, Hopkins, and Yamamoto 2014) could help to establish the electoral value of X_{id} . A less compelling test might instead show a limited correlation between X_{id} and V_{id} . However, researchers cannot simply use the ex post sample of close elections, where X_{id} and V_{id} are uncorrelated by construction when $\Delta_{id} = 0$. To capture the local effect of X_{id} among candidates that end up in close races, researchers might consider estimating treatment effects among districts where elections were *predicted* to be close.

Theoretical Challenges. Compounding the empirical challenge of validating conditions (i) and (ii), both assumptions are often theoretically implausible. Researchers using PCRD designs are usually interested in characteristics like gender or partisan affiliation because they expect these characteristics to impact outcomes that voters also care about. For characteristic X_{id} not to influence candidate vote shares—as condition (i) requires—when it does affect district outcomes, voters would need to be oblivious to, or not vote on the basis of, the characteristic's expected impact on the outcome of interest to the researcher *or other outcomes voters are concerned about*. For example, Gagliarducci and Paserman (2012) argue that Italian voters are unlikely to vote on the basis of government termination risks and are poorly informed about whether gender might precipitate termination. Nevertheless, termination risk could still correlate with other outcomes that influence vote choice and voters attribute to gender. For condition (ii), the existence of compensating differentials that do not affect the district outcome of interest would require voters to wrongly believe that these characteristics affect an outcome they care about or not care about the outcome of interest to the researcher.

Most theories of voting behavior suggest that at least some voters observe candidate characteristics and understand how such characteristics affect outcomes they care about. Even where electorates are only partially informed about the link between characteristics and outcomes, candidates with identical vote shares should produce identical welfare outcomes—broadly construed to encompass any outcome that matters to voters which different candidates could affect—in expectation when limited information is aggregated across a population (Fowler 2018). Equal vote shares could reflect equally effective candidates committing to policies that converge on the median voter's preferred policy or comparative advantages of one candidate on some dimensions being counteracted by the comparative advantages of other candidates on other dimensions. The latter explanation does not prevent characteristic X_{id} from affecting an out-

come of interest, just that other characteristics produce offsetting effects on other outcomes that leave the median voter indifferent between two different candidates. Even if voters are not fully rational, the possibility that easily observed candidate characteristics that impact important outcomes are simultaneously uncorrelated with candidate vote shares is implausible in many contexts.

Mitigating Threats to Internal Validity

Where one of conditions (i)–(iii) is not invoked, PCRD designs lack a compelling foundation for identifying the effect of characteristic X_{id} . This is because candidates that narrowly win *must* differ, in expectation, in other consequential ways too.¹³ I next discuss three potential strategies to combat the post-treatment bias inherent to PCRD designs, and their limitations.

Continuity Tests. As Table 1 shows, most studies using PCRD designs conduct continuity tests to validate that potentially confounding district- or candidate-level characteristics do not vary discontinuously at the point of discontinuity. This entails estimating $\lim_{v \downarrow \nu} \mathbb{E}[Z_{1dk} | \Delta_d = \nu] - \lim_{v \uparrow \nu} \mathbb{E}[Z_{0dk} | \Delta_d = \nu]$ to test the null hypothesis that $\mathbb{E}[Z_{idk} | \Delta_d = 0, X_{id} = 1, X_{jd} = 0] = \mathbb{E}[Z_{idk} | \Delta_d = 0, X_{id} = 0, X_{jd} = 1]$ for each observable covariate k .

In standard RD designs, finding only differences consistent with statistical chance corroborates Assumption 1 (Cattaneo and Titiunik 2022; Imbens and Lemieux 2008). In PCRD designs, *district-level* characteristics remain useful for balance tests because—as Sekhon and Titiunik (2012) note—the standard RD continuity assumption implies continuity in district-level characteristics. Intuitively, this is because district-level characteristics can determine the types of races that are close and the degree to which compensating differentials are required by affecting electoral advantages, but cannot vary across candidate types within a given race.¹⁴ District-

¹³The asymptotic bias of PCRD designs is smaller than for designs that do not restriction attention to comparisons between relatively similar candidates when the posttreatment bias is small relative to district- and candidate-level differences between districts that differ in X_{id} ; see Supporting Information Appendix D (p. 7).

¹⁴For example, assume V_{id} does not depend on $Z_{id} = Z_{jd} = Z_d$ in Equation (5). However, a district-level characteristic could *differentially* affect candidate type $X_{id} = 1$, for example, if $V_{id} = \alpha \frac{X_{id} - X_{jd}}{2} + \beta_1 \frac{(Z_{id1} - Z_{jd1})}{2} + \beta_2 \frac{(X_{id} - X_{jd})Z_{d2}}{2} + \frac{\varepsilon_{id} - \varepsilon_{jd}}{2}$. Nevertheless, the effect of Z_d on Y_{id} —given by $\mathbb{E}[Z_d | \alpha + \beta_1(Z_{id1} - Z_{jd1}) + \beta_2 Z_{d2} + \varepsilon_{1d} - \varepsilon_{0d} = 0]$ —is the same when types $X_{id} = 1$ and $X_{id} = 0$ narrowly win, and thus cancels out. The magnitude of $Z_{id1} - Z_{jd1}$ required to compensate for $\alpha + \beta_2 Z_{d2}$, rather than α , increases though.

level continuity tests are rightly common in PCRD applications, but do not imply continuity in candidate-level covariates.

In contrast, continuity tests for *candidate-level* covariates operate differently. If neither condition (i) nor condition (ii) can be invoked,¹⁵ there *must* exist at least one (observable or unobservable) compensating differential. Consequently, detecting discontinuities in theoretically plausible compensating differentials—as Gagliarducci and Paserman (2012) do for age, education, and vocational experience—can now serve as a manipulation check guiding researcher interpretation of PCRD estimates. As I discuss below, characterizing compensating differentials—and thus potential sources of bias—can inform efforts to bound estimates or reinterpret treatments. Conversely, failing to reject continuity in observable candidate-level covariates does not necessarily validate a PCRD design. This is because continuity in observable covariates is consistent with condition (i) or (iii) holding *as well as* the existence of unobserved compensating differentials, a lack of statistical power to detect observable compensating differentials, or false positive results.

Covariate Adjustment. Where the assumptions necessary for identification do not obviously hold, a strategy common to many statistical approaches is covariate adjustment. This involves adjusting for predetermined potential confounders to the greatest extent possible using observable covariates. Gagliarducci and Paserman (2012) address imbalances between men and women who narrowly won mayoral in elections in Italy by adjusting for various covariates, including those on which significant imbalances were observed. Covariate adjustment can be implemented by adjusting for a subset of compensating differentials $Z_{id}^{cond} \subset Z_{id}$ using local polynomial estimators (Calonico et al. 2019).

However, adjusting for candidate-level characteristics does little to address the posttreatment bias that arises in PCRD designs. Since Equation (7) must always hold, covariate adjustment does not increase the plausibility of condition (i) because conditioning on Z_{id}^{cond} induces or accentuates the need for compensating differentials in terms of other covariates $Z_{id} \setminus Z_{id}^{cond}$ that are not adjusted for. For example, a researcher using a PCRD design to estimate effects of electing university-educated politicians might condition on ideology because they are concerned that better educated politicians are in close races because they espouse unpopular policy positions.

¹⁵Or another condition ensuring $\mathbb{E}[g(Z_{id}) | \Delta_{id} = 0, X_{id} = 1, X_{jd} = 0] = \mathbb{E}[g(Z_{id}) | \Delta_{id} = 0, X_{id} = 0, X_{jd} = 1]$.

Even if covariate adjustment breaks the correlation between education and ideology, university-educated politicians in close races with nonuniversity-educated politicians with similar ideologies must still, in expectation, differ in other ways to remain in close races. Covariate adjustment can increase the plausibility of conditions (ii) and (iii) by increasing the share of variation in candidate vote share explained by noise, that is, reducing $\frac{\sigma_{Z|Z^{cond}}^2}{\sigma_\varepsilon^2}$. However, candidate vote shares only differ due to noise once a researcher has fully adjusted for *all* compensating differentials that affect district-level outcomes, including compensating differentials induced by conditioning.

Bounding and Correcting Effect Magnitudes. Where none of the additional conditions that yield consistent PCRD estimates are plausible, a more promising strategy—at least for more limited researcher objectives—is to use $\hat{\tau}_{PCRD}$ to bound the effect of X_{id} or correct estimates of its effects. Such strategies can help establish the direction of the effect or the direction and degree of bias that would nullify or reverse a directional finding.

The preceding discussion of when PCRD designs over and underestimate effects of X_{id} illuminates the benefits and drawbacks to bounding. Indeed, underestimation—which enables researchers to claim that an effect of X_{id} is not smaller in magnitude than $\hat{\tau}_{PCRD}$ —occurs when the net effect of all compensating differentials affects the outcome in the same way as X_{id} at the discontinuity. This relatively strong conclusion for nonnull findings could be substantiated using continuity tests to identify compensating differentials and then providing theoretical or empirical evidence to argue that X_{id} and Z_{idk} affect Y_d in the same direction at the discontinuity. For example, if primary voters are averse to ideological extremists and such candidates must compensate by being more competent on average, then Hall's (2015) results might understate the general election penalty associated with selecting extreme candidates in primary elections. However, establishing the direction of an effect is harder when $\hat{\tau}_{PCRD}$ fails to reject the null hypothesis because we cannot be sure if underestimation accounts for accepting the null.

In the spirit of Rosenbaum (2002), bias correction may be possible where compensating differentials are observable and plausible estimates of their effects on district outcomes can be imputed. If g is a linear function and $\hat{\gamma}_k$ is a credible estimate of the LATE of each compensating differential Z_{idk} at the point of discontinuity, then Proposition 2 implies that the PCRD esti-

mate could be corrected to obtain the LATE of X_{id} as follows¹⁶:

$$\hat{\tau}_{PCRD}^{corr} = \hat{\tau}_{PCRD} - \sum_k \hat{\gamma}_k \hat{\delta}_k, \quad (9)$$

where each $\hat{\delta}_k$ consistently estimates $\mathbb{E}[Z_{idk}|\Delta_d = 0, X_{id} = 1, X_{jd} = 0] - \mathbb{E}[Z_{idk}|\Delta_d = 0, X_{id} = 0, X_{jd} = 1]$ using a (dis)continuity test for observable compensating differentials. Even without estimates of γ_k , researchers could examine the sensitivity of their results to plausible values of γ_k . In the gender example, this could involve estimating the difference in candidate competence between men and women who narrowly win and then $\hat{\tau}_{PCRD}$ for plausible values of γ_k . Supporting Information Appendix E (p. 8) provides examples of how attention to theorized mechanisms can inform γ_k .

Expanding the Conception of Treatment

Given the challenges of isolating effects of politician characteristic X_{id} , an alternative approach is to explicitly redefine the estimand to include compensating differentials. Specifically, researchers might target a compound treatment effect incorporating effects of X_{id} and all compensating differentials induced or altered by PCRD designs at the point of discontinuity. Hall (2015, 24) adopts this type of approach when noting that his PCRD estimate of the effect of selecting ideological extremists in U.S. primary elections on general election outcomes “includes the component of the overall effect that comes from the change in ideology, but also includes any other factors that differ between the two types of candidates.” He thus distinguishes between a specific individual characteristic—*extremism*—and the bundle of correlated characteristics characterizing a typical *extremist*. This logic extends to distinguishing the effects of electing a candidate representing the positions of the Democratic party from electing candidates that are Democrats. In the context of women winning primary elections, Bucchianeri (2018, 445) similarly defines his estimand as the “causal effect of nominating a female candidate, *not* the causal effect of gender,” and notes that this bundle could include compensating differentials that ensure women remain in close races with men. The key advantage of redefining the treatment of interest to include correlated characteristics and compensating differentials is that

¹⁶Letting Z_{id} include higher order polynomials and interactions between characteristics, the Weierstrass approximation theorem ensures that $\sum_k \gamma_k Z_{idk}$ approximates $g(Z_{id})$.

PCRD designs can now yield consistent estimates, albeit for a different estimand.

Formally, this reconceptualization entails focusing on joint potential outcomes $Y_d(X_d, \mathbf{Z}_d) = y(X_d, \mathbf{Z}_d) + \nu_d$, where \mathbf{Z}_d remains a vector of other characteristics of the winning candidate in district d . The following proposition characterizes the compound treatment effect that a PCRD estimator converges to under the standard RD continuity assumption:

Proposition 4. *Under Assumptions 1(a) and 2:*

$$\hat{\tau}_{\text{PCRD}} \xrightarrow{p} \int [y(1, \mathbf{z}) - y(0, \mathbf{z})] f_c(\mathbf{z}) d\mathbf{z} + \int y(1, \mathbf{z}) f_1(\mathbf{z}) d\mathbf{z} - \int y(0, \mathbf{z}) f_0(\mathbf{z}) d\mathbf{z},$$

where $f_{\mathbf{Z}_{id}|\cdot}(\mathbf{z})$ is the conditional probability density function of \mathbf{Z}_{id} , the pointwise common component of the density is $f_c(\mathbf{z}) := \min\{f_{\mathbf{Z}_{id}|\Delta_{id}=0, X_{id}=1, X_{jd}=0}(\mathbf{z}), f_{\mathbf{Z}_{id}|\Delta_{id}=0, X_{id}=0, X_{jd}=1}(\mathbf{z})\}$, and $f_m(\mathbf{z}) := f_{\mathbf{Z}_{id}|\Delta_{id}=0, X_{id}=m, X_{jd}=1-m}(\mathbf{z}) - f_c(\mathbf{z})$ is the excess density among politicians of type $X_{id} = m$ that win close elections.

When the distribution of \mathbf{Z}_{id} differs across candidates of type $X_{id} = 1$ and $X_{id} = 0$ that win close elections, this result shows that $\hat{\tau}_{\text{PCRD}}$ captures effects of both X_{id} and \mathbf{Z}_{id} . The proposition expresses this in terms of a LATE of X_{id} , weighted by the common distribution of \mathbf{Z}_{id} characteristics across candidate types, and effects of differences in the distribution of \mathbf{Z}_{id} across candidate that differ in X_{id} . For Hall (2015), X_{id} represents ideological extremism and \mathbf{Z}_{id} captures all other characteristics of extremists—both those that naturally correlate with X_{id} and those induced, accentuated, or attenuated by conditioning the estimand on close elections that are affected by X_{id} .

Reconceptualizing potential confounders as part of the PCRD estimand again implies a nonstandard role for candidate-level continuity tests. Rather than validating Assumption 1, candidate-level covariate tests yield estimates of $\hat{\delta}_k$ that now help to characterize the compound treatment. Substantial differences in $Z_{1dk} - Z_{0dk}$ at the point of discontinuity suggest that Z_{idk} may be an important component of the comparison captured by $\hat{\tau}_{\text{PCRD}}$, whereas the reverse holds for covariates where discontinuities are not detected. Hall (2015) adopts such an approach by examining whether extremist primary winners differ from nonextremist winners on other dimensions in his analysis of mechanisms.

There are, however, three notable drawbacks to broadening the notion of treatment; the importance of each drawback will vary by application. First, it is hard to fully characterize treatment in many empirical applications. This is both because compensating differ-

entials should generally exist but continuity tests may lack the statistical power to detect differences in \mathbf{Z}_{id} and because researchers may struggle to adequately measure relevant elements of \mathbf{Z}_{id} . Beyond the label “ X and all its compensating differentials,” PCRD designs lack clarity about the bundle of characteristics that constitute the treatment.

Second, the external validity and interpretability of PCRD estimates may be limited because the design is unlikely to capture typical or homogeneous bundles of characteristics. PCRD designs identify compound treatments defined by the correlations between characteristics that exist *after* conditioning on close elections where X_{id} affects vote shares. Extremists that win close races against nonextremists may thus be atypical of extremists that narrowly win any type of primary election, in addition to being atypical of extremists in general. Moreover, because many permutations of $(\mathbf{Z}_{id}, \mathbf{Z}_{jd})$ can produce close elections, narrowly winning candidates of type X_{id} can experience different values of \mathbf{Z}_{id} —a violation of the treatment uniformity component of SUTVA. For example, some extremists that overcome an electoral penalty associated with their ideological extremism to win may be unusually competent and others may offer more appealing platforms.

Third, bundled treatments limit the degree to which PCRD designs can test specific theories or inform certain policy decisions. Theories often specify “all else equal” comparative statics for different variables that PCRD designs cannot distinguish because all candidate-level characteristics, albeit to differing degrees, are considered part of a compound treatment. PCRD designs therefore cannot reveal whether ideological extremists lose general elections because of their policy positions, differences in competence between extremist and nonextremist candidates that narrowly win, or some combination of both.

The extent to which this inability to distinguish the contribution of different elements of the compound treatment limits the relevance of PCRD estimates to policy makers likely depends on the policy tools available. On one hand, policy makers with constrained choice sets may not care which part of the treatment matters, only that policies that encourage (or discourage) politicians of type X_{id} to run for office or help such candidates win elections should be favored. For example, local party committees might alter candidacy rules to avoid selecting extremist candidates that lose general elections. On the other hand, the limited information conveyed by PCRD estimates is less helpful where policy makers are picking between or devising more fine-grained policies that can encourage candidates of type X_{id} instead of type

Z_{idk} . Understanding the mechanism driving PCRD estimates could be consequential for reformers investigating whether they should adopt gender quotas or require more specific competencies of their candidates.

Conclusions

This article demonstrates that PCRD designs—a popular approach used to estimate effects of a specific characteristic, or bundle of characteristics, of elected politicians on downstream outcomes—generally require imposing substantially stronger assumptions than standard RD designs. This is because the treatment variable in this nonstandard RD application is defined both by winning close elections and a candidate characteristic that can affect selection into the set of narrow election winners of different types. I have shown that such posttreatment conditioning causes PCRD designs to capture the effect of the specific characteristic of interest together with all the compensating differentials required for candidates with the characteristic of interest to remain in close races.

Even when the characteristic of interest is unconditionally independent of other characteristics, PCRD designs generate inconsistent estimates of the LATE exclusively attributable to the characteristic that defines treatment, except under two strong additional assumptions: (i) the characteristic of interest did not affect the candidate's vote share; or (ii) no compensating differential affected the outcome of interest. Unfortunately, neither condition is plausible in many contexts and both are difficult to empirically validate. Accordingly, PCRD designs cannot generally isolate impacts of specific politician characteristics on outcomes relating to political representation, accountability, and participation.

Researchers can attempt to combat this challenge in several ways. One approach is to explicitly accept and then mitigate threats to internal validity using theory and data to bound or sign treatment effects. Another approach broadens the definition of treatment to exclude the possibility of candidate-level confounding by redefining the estimand to include both the characteristic of interest and all compensating differentials. Both approaches entail trade-offs, either in terms of internal validity or the generality of treatment, but could nevertheless illuminate hypotheses in certain settings—even though PCRD designs cannot isolate effects in the same way as standard RD designs.

References

- Ashworth, Scott, Christopher R. Berry, and Ethan Bueno de Mesquita. 2021. "Modeling Theories of Women's Underrepresentation in Elections." Working paper.
- Broockman, David E. 2014. "Do Female Politicians Empower Women to Vote or Run for Office? A Regression Discontinuity Approach." *Electoral Studies* 34:190–204.
- Bucchianeri, Peter. 2018. "Is Running Enough? Reconsidering the Conventional Wisdom about Women Candidates." *Political Behavior* 40(2):435–66.
- Bueno de Mesquita, Ethan, and Scott A. Tyson. 2020. "The Commensurability Problem: Conceptual Difficulties in Estimating the Effect of Behavior on Behavior." *American Political Science Review* 114(2):375–91.
- Calonico, Sebastian, Matias D. Cattaneo, Max H. Farrell, and Rocío Titiunik. 2019. "Regression Discontinuity Designs Using Covariates." *Review of Economics and Statistics* 101(3):442–51.
- Calonico, Sebastian, Matias D. Cattaneo, and Rocío Titiunik. 2014. "Robust Nonparametric Confidence Intervals for Regression-Discontinuity Designs." *Econometrica* 82(6):2295–2326.
- Cattaneo, Matias D., and Rocío Titiunik. 2022. "Regression Discontinuity Designs." *Annual Review of Economics* 14:821–51.
- Caughey, Devin, and Jasjeet S. Sekhon. 2011. "Elections and the Regression Discontinuity Design: Lessons from Close US House Races, 1942–2008." *Political Analysis* 19(4): 385–408.
- Clots-Figueras, Irma. 2011. "Women in Politics: Evidence from the Indian States." *Journal of Public Economics* 95(7–8): 664–90.
- de la Cuesta, Brandon, and Kosuke Imai. 2016. "Misunderstandings about the Regression Discontinuity Design in the Study of Close Elections." *Annual Review of Political Science* 19:375–96.
- Eggers, Andrew C. 2017. "Quality-Based Explanations of Incumbency Effects." *Journal of Politics* 79(4):1315–28.
- Eggers, Andrew C., Anthony Fowler, Jens Hainmueller, Andrew B. Hall, and James M. Snyder, Jr. 2015. "On the Validity of the Regression Discontinuity Design for Estimating Electoral Effects: New Evidence from over 40,000 Close Races." *American Journal of Political Science* 59(1): 259–74.
- Eggers, Andrew C., Ronny Freier, Veronica Grembi, and Tommaso Nannicini. 2018. "Regression Discontinuity Designs Based on Population Thresholds: Pitfalls and Solutions." *American Journal of Political Science* 62(1):210–29.
- Eggers, Andrew C., and Jens Hainmueller. 2009. "MPs for Sale? Returns to Office in Postwar British Politics." *American Political Science Review* 103(4):513–33.
- Ferreira, Fernando, and Joseph Gyourko. 2014. "Does Gender Matter for Political Leadership? The Case of US Mayors." *Journal of Public Economics* 112:24–39.
- Folke, Olle. 2014. "Shades of Brown and Green: Party Effects in Proportional Election Systems." *Journal of the European Economic Association* 12(5):1361–95.

- Fowler, Anthony. 2018. "A Bayesian Explanation for the Effect of Incumbency." *Electoral Studies* 53:66–78.
- Gagliarducci, Stefano, and M. Daniele Paserman. 2012. "Gender Interactions within Hierarchies: Evidence from the Political Arena." *Review of Economic Studies* 79(3):1021–52.
- Gelman, Andrew, and Guido Imbens. 2019. "Why High-Order Polynomials Should not Be used in Regression Discontinuity Designs." *Journal of Business & Economic Statistics* 37(3):447–56.
- Gerber, Elisabeth R., and Daniel J. Hopkins. 2011. "When Mayors Matter: Estimating the Impact of Mayoral Partisanship on City Policy." *American Journal of Political Science* 55(2):326–39.
- Hahn, Jinyong, Petra Todd, and Wilbert Van der Klaauw. 2001. "Identification and Estimation of Treatment Effects with a Regression-Discontinuity Design." *Econometrica* 69(1):201–209.
- Hainmueller, Jens, Daniel J. Hopkins, and Teppei Yamamoto. 2014. "Causal Inference in Conjoint Analysis: Understanding Multidimensional Choices via Stated Preference Experiments." *Political Analysis* 22(1):1–30.
- Hall, Andrew B. 2015. "What Happens When Extremists Win Primaries?" *American Political Science Review* 109(1):18–42.
- Heckman, James J. 1979. "Sample Selection Bias as a Specification Error." *Econometrica* 47(1):153–61.
- Hernán, Miguel A., and James M. Robins. 2011. *Causal Inference*. Boca Raton, FL: CRC Press.
- Imbens, Guido W., and Thomas Lemieux. 2008. "Regression Discontinuity Designs: A Guide to Practice." *Journal of Econometrics* 2(142):615–35.
- Lawless, Jennifer L. 2015. "Female Candidates and Legislators." *Annual Review of Political Science* 18:349–66.
- Lee, David S., Enrico Moretti, and Matthew J. Butler. 2004. "Do Voters Affect or Elect Policies? Evidence from the US House." *Quarterly Journal of Economics* 119(3):807–59.
- McCrary, Justin. 2008. "Manipulation of the Running Variable in the Regression Discontinuity Design: A Density Test." *Journal of Econometrics* 142(2):698–714.
- Montgomery, Jacob M., Brendan Nyhan, and Michelle Torres. 2018. "How conditioning on posttreatment variables can ruin your experiment and what to do about it." *American Journal of Political Science* 62(3):760–775.
- Pettersson-Lidbom, Per. 2008. "Do Parties Matter for Economic Outcomes? A Regression-Discontinuity Approach." *Journal of the European Economic Association* 6(5):1037–56.
- Rosenbaum, Paul R. 2002. *Observational Studies*. Berlin: Springer.
- Slough, Tara. Forthcoming. "Phantom Counterfactuals." *American Journal of Political Science*.
- Sekhon, Jasjeet S., and Rocío Titiunik. 2012. "When Natural Experiments Are neither Natural nor Experiments." *American Political Science Review* 106(1):35–57.

Supporting Information

Additional supporting information may be found online in the Supporting Information section at the end of the article.

Appendix A: Published articles using PCRD designs

Appendix B: Proof of propositions

Appendix C: Example with two compensating differentials

Appendix D: Do PCRD designs reduce bias relative to observational designs?

Appendix E: Differential empirical implications of theorized mechanisms