Defining and Identifying Incumbency Effects*

An Application to Brazilian Mayors

Fredrik Sävje†

January 29, 2015

PRELIMINARY DRAFT

Recent studies of the effects of political incumbency on election outcomes have almost exclusively used regression discontinuity designs. This shift from the past methods has provided credible identification, but only for a specific type of incumbency effect: the effect for parties. The other effects in the literature, most notably the personal incumbency effect, have largely been abandoned together with the methods previously used to estimate them. This study aims at connecting the new methodical strides with the effects discussed in the past literature. A causal model is first introduced which allows for formal definitions of several effects that previously only been discussed informally. The model also allows previous methods to be revisited and derive how their estimated effects are related. Several strategies are then introduced which, under suitable assumptions, can identify some of the newly defined effects. Last, using these strategies, the incumbency effects in Brazilian mayoral elections are investigated.

*I would like to extend my gratitude to Sebastian Axbard, Natalia Bueno, Per Johansson, Eva Mörk and Jasjeet Sekhon for comments and suggestions on this study. Needless to say, all errors are my own.

†Department of Economics, Uppsala University and UCLS. E-mail: fredrik.savje@nek.uu.se.

1 Introduction

The question whether political incumbency affects election outcomes has occupied political scientists and economists for more than half a century. Starting with theoretical work in the 1960's, the line of thought has been that winning candidates or parties benefit from holding office, and that this results in a higher vote share in the subsequent elections. While most of the discussed mechanisms—such as greater access to media outlets, improved name recognition and various financial benefits—would suggest a positive effect, it is conceivable that there are mechanisms that affect the outcomes negatively. For example, the electorate could be unwilling to allow long-running incumbents if incumbency increase political connectedness (which in turn could facilitate corruption), or the electorate could simply grow tired seeing the same face in office. Both the sign and magnitude of the effect are ultimately empirical questions.[1]

The initial theoretical discussions were followed by a vast array of empirical studies. A notable early contribution is Erikson (1971), who investigates the incumbency effect in the U.S. House of Representatives by comparing the election outcomes of successful first-time runners with their outcome in the subsequent re-election attempt. Empirical investigations have, however, proven particularly hard to conduct credibly. Most of the earliest methods were plagued by severe biases (see, e.g., the discussion in Gelman and King, 1990).

Starting with Lee (2001, 2008), the currently dominating strand of the literature uses regression discontinuity designs (RDD) to provide credible identification of causal effects. This design exploits the fact that the winning party changes discontinuously with the parties' vote margins. In a two party system, if one of the parties receives just shy of half of the votes, it loses the election. Just at the cut-off, changing only a single vote would change the election outcome. Under the assumption that all other relevant factors are continuous at the zero percent vote margin, any difference in the investigated outcome could be argued to arise due to the change incumbency. Under favorable conditions (Caughey and Sekhon, 2011), this design provides very credible identification.

The study by Lee (2001, 2008) also introduced a formal causal model to define the effect he investigated. This model made apparent that it was a specific incumbency effect that was identified with the RDD, namely the effect of being an incumbent *party*. This was a considerable departure from the previous literature. As the effects were often defined only informally (Gelman and King, 1990, is a notable exception), it is not clear exactly which effects were of interest in the studies prior to Lee, but they seem to have primarily focused on the effect of incumbent *candidates*. The difference between the two effects

[1] In fact, the substantial incumbency advantage found in for U.S. election seems to be a quite recent phenomenon, starting in the early 1960's (Cox and Katz, 1996). Studies of less stable elections setting even find considerable *negative* effects (Titiunik, 2011; Uppal, 2009).

is that a winning candidate has incumbency only if he or she is the party's candidate in the subsequent election, a winning party has incumbency in the subsequent election independently of whether the candidate re-runs. In fact, the past literature has often used open-seat elections as their definition of non-incumbency. Since a party can be incumbent in an open-seat election, some cases which are considered to constitute incumbency under Lee's definition would be defined as non-incumbents in much of the past literature.

In this study, I intend to partly bridge the gap between the two strands of the literature. In more detail, I aim to contribute to the literature in four ways.

First, I introduce a causal model with which different types of incumbency effects can be defined. This model allows me to formally define the two effects discussed above, and two new types of incumbency effects: the *personal incumbency effect* and the *direct party incumbency effect*. These effects are among those that have been discussed informally previously in the literature, but have, to my knowledge, never formally been defined or identified. The formal definitions provide a structured way to think about and discuss incumbency effects. In particular, the investigation reveals that what has been referred to as a single personal incumbency effect are several different effects.

Second, I show that most of the predominate estimands and estimators discussed in the literature can be re-interpreted with this causal model. In this investigation, I grant each estimator its identifying assumptions—the exercise is purely to examine which effects they would estimate if they were to succeed (i.e., deriving the associated estimand). This will aid in the interpretation of these measures and clarify how they are related to other measures. In fact, the exercise reveals that some of the previous studies primarily focus on a conditional version of one of the effects. The exercise also allows me to decompose the estimand from Lee's (2001; 2008) causal model and express it using the incumbency effects defined here. This provides a direct link between the two models.

Third, I show that local versions of the personal incumbency effect and the direct party incumbency effect can be identified using a version of the regression discontinuity design. Specially, I introduce and discuss three different identification strategies with various identifying assumptions. The strategies mainly differ in the degree of "localness" of the estimand and the severity of the identifying assumptions. These range from no additional assumption other than those from the RDD, to a weak version of an independence assumption. As one would expect, making stronger assumptions will allow for identification of a less local effect.

Last, to illustrate these strategies, I use data from recent Brazilian mayoral elections to estimate the personal and direct party incumbency effects. The Brazilian setting is one where the party incumbency effect has been shown to be negative (Titiunik, 2011). At least two possible explanations for the negative effect can be imagined. First, the electorate might punish undesirable past behavior of incumbents. Second, the electorate might want to avoid lame-duck mayors, who, for example, could be more prone to corruption (Ferraz and Finan, 2011). As voters cannot exercise (electoral) disciplinary power over lame-duck mayors, they act preemptively and tend to not grant candidates second terms. As the second explanation pertains to candidates rather than parties, we would expect that the personal incumbency effect is more negative than the direct party effect if this was the main channel of influence. However, the estimated direct party effect (-20.8%) is considerably more negative than the personal effect (-13.4%), indicating that first explanation is more likely to be at play in the Brazilian setting.

2 Defining incumbency effects

The intuitive definition of incumbency effects as the change in election performance due to a party or a candidate being incumbent is rather vague. Exactly what is meant with "incumbent"? Incumbent compared to what other state? And whom are we investigating? A disciplined discussion about the effects requires a clear answer to these and other questions.

Many of the early contributions defined their investigated effects in terms of observed variables, and often in close connection to their designs. The problem with this approach is that the definition in itself necessitates identification in order to have a causal interpretation—identifying and defining the effect therefore, in some sense, become simultaneous. As a result, one cannot ask whether one has identified the causal effect that has been defined, as the defined effect is not causal if it is not identified. This illustrates the benefit of a causal model. With it, we can define the causal effect separately from the observed data, and thereby discuss the effects independently of the details of the design and estimation. In this section, I will extend the prior causal models used for incumbency effects so that many of the previously discussed effects can be defined with it.

2.1 A causal model of incumbency effects

I will construct the causal model in the "potential outcomes framework" or the Neyman-Rubin Causal Model (NRCM), first introduced in experimental settings by Splawa-Neyman et al. (1923/1990) and in observational settings by Rubin (1974).

The units of observation in this model are party-elections denoted by index i. For example, $i = 1$ could denote the Democratic Party in the 2004 House of Representatives elections in California's 13th congressional district. All party-elections are collected in a set denoted by $\mathscr{I}$. For every i, there are two variables that we, in the definitions, consider to be manipulated—i.e., our treatment variables. W_i is a binary indicator of whether the party won the election preceding the election denoted by i, and R_i is a binary indicator of whether the candidate of the party in the previous election runs for office in the election denoted by i. For example, if $i = 2$ refers to the Republican Party in the 2004 presidential election, the observed values would be $W_2 = 1$ and $R_2 = 1$. If $i = 3$ was the Republican Party in the 2008 presidential election, we would instead have $W_3 = 1$ and $R_3 = 0$.

In the thought-experiment where we can control W_i and R_i, we can realize four different worlds representing the four possible combinations of the two variables. For example, we could change the chain of events so that the Republicans lost the 2004 presidential election ($W_3 = 0$), or that George W. Bush did not enter the 2004 presidential election ($R_2 = 0$). Let Y_i denote the observed outcome of interest—often vote shares or a binary indicator of whether the party won the election. The potential outcome will be denoted with $Y_i(w,r)$, where w is whether the party won the preceding election and r whether the candidate re-ran. For example, in the world where i won the election ($W_i = 1$) and the candidate re-runs ($R_i = 1$) for office, $Y_i(1,1)$ would be realized outcome.[2] Figure 1 provides an illustration of the definition of the potential outcomes.

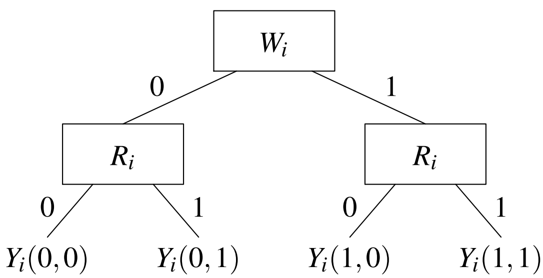

Figure 1. Potential outcomes defined over W_i and R_i. Rectangles indicate variables which are manipulated. Starting in the top node and following the path according to the chosen manipulation we can realize any of the potential outcomes.

The exact manipulations of W_i and R_i are intentionally left vague at this point. Using the model would require that these be made precise so that we can ensure that there is no interference between units or hidden variation of the treatment. In other words, we need to make sure that the assumptions discussed in the introductory chapter are satisfied (e.g., the SUTVA).[3] I will discuss these issues in the following sections, but I will here assume that the

[2]We must here ensure that the outcome is defined in each of the hypothetical worlds (i.e., the observability assumption from the introductory chapter). For example, if an election is uncontested it is not obvious how the victory margin would be coded. To ease exposition, I will disregard these issues in the current section, and assume that all potential outcomes are defined.

[3]There are several issues that need attention. For example, the election winner is a deterministic function of the vote shares, thus manipulating W_i implies manipulating the vote shares. However, large changes in vote share could have fundamentally different interpretations than small changes, thereby potentially violating the consistency assumption. Similarly, exactly how

assumptions hold without motivation; for now, the model can best be seen as a template of a causal model. Already at this stage, the model, however, restricts the potential outcome to a great extent and thereby clarifies the interpretation of them. For example, the hypothetical worlds in this model differ from those in Ansolabehere et al. (2000). In that study, the authors exploits re-districting of election districts to investigate how voters that encounter the candidate for the first time (due to being moved to another election district) vote compared to the "old" voters in the district. While this certainly is informative of an eventual incumbency advantage, it is a fundamentally different effect than the current.

Furthermore, the exact meaning of an "incumbent candidate" is made clear. In this model, it is when the winning candidate from the previous election run for the same office in the current election ($W_i = R_i = 1$). As a consequence, under this definition, Gerald Ford was *not* an incumbent candidate in the 1976 presidential election since he did not run in the preceding election. An alternative definition would be to define incumbency as being the current office holder coming into the election (in which case Gerald Ford would be an incumbent in 1976). While this is a reasonable alternative, and in some ways even preferable due to its closeness to the intuitive concept, it is not clear exactly how we would manipulate office holding in that case. As each type of manipulation affects the interpretation, the effect remain vague under this definition—an unexpected death (Cox and Katz, 2002) or resigning due to threat of impeachment would both change the office holder, but they would probably lead to very different types of incumbency effects. As the vast majority of office holders came into power by winning an election, the current definition is arguably a good balance between clarity and closeness to the intuitive concept.

Yet another, and seemingly intuitive, definition of the potential outcomes is to use the incumbency indicator (I_i) used previously in the literature (Gelman and King, 1990). In a two-party system, this indicator takes value $I_i = 1$ if party i has an incumbent candidate in the election; $I_i = -1$ if the opposing party has an incumbent candidate; and $I_i = 0$ if there is an open-seat election. We would then have three potential outcomes: $Y_i(1)$, $Y_i(0)$ and $Y_i(-1)$. This would, however, unlikely satisfy the consistency assumptions. To see why, note that there is a link between the two models. In a two-party system, let $j : \mathscr{I} \to \mathscr{I}$ be a mapping from a party to its opponent in the same election. For example, if $i = 5$ denotes the Democratic Party in the 2004 presidential election, $j(5)$ gives the index of the Republican Party in the 2004 presidential election. We then have $I_i = W_i R_i - (1 - W_i) R_{j(i)}$. With the incumbency indicator, $Y_i(1)$ maps unambiguously to $Y_i(1,1)$, but $Y_i(0)$ could be any of $Y_i(0,0)$, $Y_i(0,1)$ and $Y_i(1,0)$. In some of these hypothetical worlds, the party

one ensures that the previous candidate re-run for office is not obvious, and there are likely situations where it is impossible to manipulate that variable.

won the previous election, while in others it did not. Since winning the previous election potentially has large effects on the subsequent election outcomes, the consistency assumption, as discussed in the introductory chapter, is unlikely to hold.[4]

2.2 The incumbent legislator effect

The effect on election outcomes for parties, when running with an incumbent legislator holding party incumbency constant.

Much of the literature prior to Lee (2001, 2008) focused on the effect of an incumbent candidate on parties' election outcomes. In other words, the question is whether the parties benefit from running with candidates that won the preceding elections. The estimand defined in this section is an effort to formalize this concept. As we will see in following sections, this definition is not new but corresponds exactly to the definition by Gelman and King (1990). The use of "legislator" in the name of this effect is borrowed from Caughey and Sekhon (2011); the effect is not intended to be restricted only to candidates of the legislative branches of government, but to any elected official.

In the current setting, for a party to have an incumbent office holder, two conditions must be true: the party must have won the previous election and the previous candidate must re-run for office. As this implies $W_i = 1$ and $R_i = 1$, the associated potential outcome is clearly $Y_i(1,1)$. The other potential outcome is, however, less clear: the intuitive concept often states "versus not having an incumbent legislator." In principle this could refer to any of $Y_i(1,0)$, $Y_i(0,1)$ and $Y_i(0,0)$.

In an effort to isolate the effect of an incumbent *office holder*, note that two of these potential outcomes entails more than just a change in whether the party has an incumbent candidate. In the hypothetical worlds denoted by $Y_i(0,1)$ and $Y_i(0,0)$, the party is no longer the incumbent party. With any of those potential outcomes, the effect would be compounded by both a change in legislator incumbency and party incumbency. Arguably, the potential outcome that is closest to "not having an incumbent legislator" is thus $Y_i(1,0)$. As an added bonus, $Y_i(1,0)$ unambiguously refer to open-seat election which is the typical contrast in much of the previous literature.

[4] An advantage with the incumbency indicator is that it differentiates between open-seat elections and incumbent elections—a contrast given great importance in the previous literature. The current model does not fully impose that difference as $Y_i(0,0)$ and $Y_i(0,1)$ can refer to both open-seat and incumbent elections (for the opposing party). If that difference is deemed to be of importance, one could define the potential outcomes over W_i, R_i and $R_{j(i)}$, as that would both maintain consistency and make it possible to specify open-seat elections. However, as will we see, making this difference is not fundamental to formalizing the previous concepts, and in an effort to construct a simple model I opt for the current approach.

The incumbent legislator (causal) effect will thus be defined as the difference in election outcomes in the hypothetical worlds that would be realized when we hold W_i constant at 1 and alter R_i. Let $\tau_i^L \equiv Y_i(1,1) - Y_i(1,0)$ be the unit level incumbent legislator effect, and $\tau^L \equiv \mathrm{E}[\tau_i^L] = \mathrm{E}[Y_i(1,1) - Y_i(1,0)]$ the average incumbent legislator effect, where the expectation is taken over $\mathscr{I}$. The definition is illustrated in Figure 2.

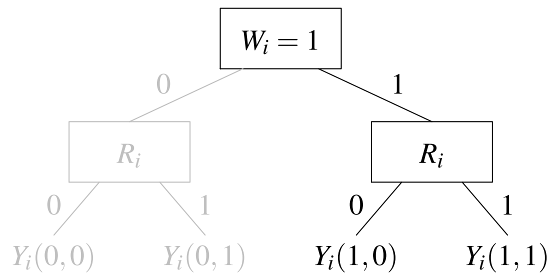

Figure 2. The incumbent legislator effect. The top node is here restricted so we only take the path of $W_i = 1$ and thus end up in either of the two rightmost end-nodes.

The definition invites to some interpretations of the discussions in the previous literature. Finding a positive τ^L would suggest several possible mechanisms. In addition to those already mentioned (e.g., media coverage, financial benefits), an incumbent candidate tend to have greater election experience. When $Y_i(1,0)$ is realized, the party's candidate is taken from the general pool of candidates, while the candidate in $Y_i(1,1)$ has, by definition, participated in a political campaign and been in office during the previous term. Subsequently, part of the legislator effect is the experience gain that the incumbent candidate enjoys. Furthermore, the candidates referred to in $Y_i(1,1)$ are from the pool of candidates that actually had ran for office (since they all ran in the previous election), while the pool of candidates in $Y_i(1,0)$ refer only to potential candidates. As we expect actual candidates to be of higher quality than potential candidates, the legislator effect will include a candidate quality component.

In subsequent sections, I will show that the estimand of several previous studies is a conditional version of τ^L, namely conditional on that the party won the previous election: $\mathrm{E}[\tau_i^L | W_i = 1]$. This conditioning could potentially lead to large changes in the effect, as the candidates in the hypothetical worlds are selected from fundamentally different pools of candidates. When conditioning on $W_i = 1$, candidates referred to by $Y_i(1,1)$ are from the pool with *winning* candidates, while candidates referred by $Y_i(1,0)$ remain largely the pool of potential candidates. While one might concede that actual and potential candidates are on average of equal quality, it would be a stretch to say the same of winning and potential candidates. Subsequently, we would expect $\mathrm{E}[\tau_i^L|W_i = 1]$ to be greater than $\mathrm{E}[\tau_i^L]$.[5]

This fact has implications for the interpretation of previous studies. For example, Cox and Katz (1996) provide an insightful decomposition of the legislator effect into two parts: what they refer to as a direct effect—that is, the benefits directly pertaining to the incumbent candidate—and an indirect, "scare-off" effect due that the opposing candidates of incumbents tend to be of worse quality. If we let Q_i denote the opposing candidate's quality, a scare-off effect would imply $\mathrm{E}[Q_i(1,1) - Q_i(1,0)] < 0$. They explain the existence of the scare-off effect by that high-ability challengers on average have better outside options. Thus, if the high-ability challengers expect to perform badly in the election (due, e.g., to the direct incumbent advantage), they refrain from participating leaving only low quality candidates that lacks attractive outside options. As noted by Cox and Katz (1996), the scare-off effect requires a positive direct effect (unless the challengers are irrationally scared off). However, like most of the past studies, Cox and Katz (1996) estimate the conditional version of the effect. Thus, their scare-off effect is $\mathrm{E}[Q_i(1,1) - Q_i(1,0)|W_i = 1]$. As above, the incumbent party's candidate in $\mathrm{E}[Q_i(1,1)|W_i = 1]$ will be a winning candidate, whereas the candidates in $\mathrm{E}[Q_i(1,0)|W_i = 1]$ in general are not. Since winning candidate are likely to be of higher quality, there could be a conditional scare-off effect even if $E[Q_i(1,1) - Q_i(1,0)] = 0$. While this does not change the fundamental conclusion (i.e., the existence of a scare-off effect), it could change the interpretation. Where Cox and Katz (1996) argues that election experience is the main determinant of the scare-off effect—implying $\mathrm{E}[Q_i(1,1) - Q_i(1,0)] < 0$—their results are consistent with a scare-off effect purely due to candidate quality.

2.3 The re-running loser effect

> *The effect on election outcomes for parties, when running with a candidate that lost the previous election, holding constant that the party lost the previous election.*

The re-running loser effect is in some sense the opposite of the incumbent legislator effect: instead of the effect of running with a previously winning candidate, it is the effect of running with a previously losing candidate. Since neither the party nor the candidate is incumbent, it cannot be interpreted as an incumbency effect. The effect has, to my knowledge, not been discussed

[5] This is different from the selection bias discussed previously in the literature. Where the point here is that $\mathrm{E}[\tau_i^L|W_i = 1]$ might differ from $\mathrm{E}[\tau_i^L]$ (while both being causal effects), the past literature has been concerned with selection as an identification issue; specifically, whether a comparison similar to $\mathrm{E}[Y_i|W_i = 1, R_i = 1] - \mathrm{E}[Y_i|W_i = 1, R_i = 0]$ can be interpreted causally.

previously in the literature. It is nonetheless a causal effect and arguably still of interest.

Like the legislator effect, we alter R_i and fix whether the party won the previous election, but now so it lost ($W_i = 0$). The unit level re-running loser effect is subsequently defined as $\tau_i^R \equiv Y_i(0,1) - Y_i(0,0)$, and the average effect as $\tau^R \equiv \mathrm{E}[\tau_i^R]$.

Some of the factors influencing the legislator effect are active also here. Foremost, the candidate referred to in $Y_i(0,1)$ will in general have greater election experience than candidates in $Y_i(0,0)$. However, the benefits that an incumbent enjoy from holding office (e.g., franking benefits) are absent. Some parts of the previous literature discuss the direct benefits of office holding in excess of any electoral experience gain. It would be difficult to formulate a causal model that encapsulate this notion since it is hard to imagine situation where a candidate holds office without having some election experience, but the closest one might get to capture the idea could be $\tau^L - \tau^R$. This would, however, require that the re-running loser effect completely and solely captures the experience gain. It, for example, precludes that there is a stigma in losing elections so that electorate punishes candidates with previously poor election outcomes.

The selection artifact from the legislator effect is present also here. Where $Y_i(0,1)$ refer to actual candidates, the candidates in $Y_i(0,0)$ are only potentially so. If we believe actual candidates are of a higher quality than potential candidates, we would have $\tau^R > 0$ even if experience is irrelevant. Furthermore, if we were to estimate with effect with methods similar to those used to investigate the legislator effect, the estimand would be the effect conditional on a losing party: $\mathrm{E}[\tau_i^R | W_i = 0]$. Subsequently, the candidates in $Y_i(0,1)$ would exclusively be losing candidates, while candidates in $Y_i(0,0)$ are not. As losing candidates arguably are of lower quality, it is likely that the conditional version is negative, even if the unconditional effect is positive: $\mathrm{E}[\tau_i^R | W_i = 0] < 0 < \mathrm{E}[\tau_i^R]$.

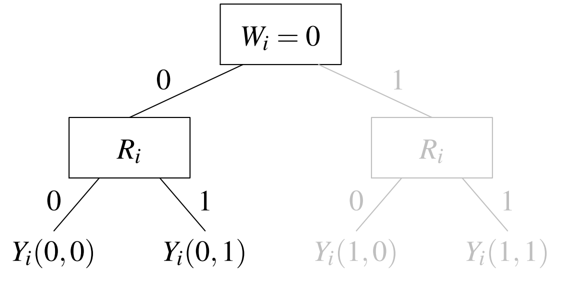

Figure 3. The re-running loser effect. The top node is here restricted so we only take the path of $W_i = 0$ and thus end up in either of the two leftmost end-nodes.

2.4 The personal incumbency effect

The effect on election outcomes for a candidate when running as incumbent office holder.

An effect discussed in the literature dating back at least to Erikson (1971) is how candidates (rather than parties) are affected by incumbency. In particular, we can, for a specific candidate, ask the counterfactual question: what would the election outcome for that candidate have been if he or she ran as the incumbent versus being a non-incumbent runner? This is the question intended to be encapsulated in the personal incumbency effect estimand, as defined in this section.

Note that this effect concerns another unit of observation than the previous effects; instead of parties, we are now interested in candidates. Subsequently, we must again specify what the manipulation is and exactly which candidates we study. The route I choose is to restrict the inquiry to candidates that ran for the same office in two subsequent elections. I take incumbency to mean that the politician won the first of the two elections. While this is not the only possibility, I will argue that it is in many ways the most reasonable.

Similarly to the definition for parties, a broader definition of incumbency would risk violating the consistency assumption. As most incumbent candidates are in power due to winning an election, this definition ensures reasonable clarity while still being relevant. A consequence of this definition is that candidates that are first-time runners do not have a well-defined personal incumbency effect. The first time a candidate runs for office he or she obviously had no chance winning the election preceding it; if one does not participate, one cannot win. These candidates could, therefore, not have been the incumbents in the sense intended here. For first-time runners to be considered incumbents, we are required to alter our imagined manipulation. For example, we could imagine the counterfactual chain of events in which Barack Obama ran for the office of President in 2004, instead for Senator, winning the primary elections instead of John Kerry and winning the presidential elections against George W. Bush, making him the incumbent candidate in the 2008 presidential elections. This would, however, be a very invasive manipulation which would produce a radically different causal effect; an effect that, arguably, is further from the intuitive concept.

Let $\mathscr{C}$ collect the indices of candidate-elections where the candidate ran in the election preceding the current. Let W_c denote whether a candidate $c \in \mathscr{C}$ won the previous election. Let $V_c(1)$ denote the potential outcome when the candidate won that election ($W_c = 1$), and let $V_c(0)$ denote the outcome when he or she lost ($W_c = 0$). Note that, if the candidate does not participate in the election denoted by c, the outcome does not exist—vote shares are only given to participating candidate. To prevent this, we must ensure (or manipulate) the world so that the candidate runs for office independently of the election outcome in the preceding election: the personal incumbency effect does not only imply manipulation of whether the candidate won the previous election, but also whether he or she runs in the current election.[6] Let R_c denote whether the candidate runs in the current election.

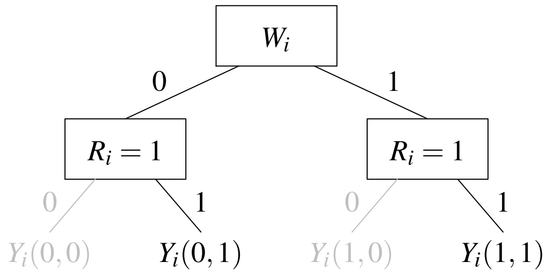

Figure 4. The personal incumbency effect with parties' potential outcomes. The second level nodes (R_i) are here restricted so we only take the paths of $R_i = 1$.

The effect is the difference between the potential outcomes when alternating W_c while holding R_c constant at one. The unit (candidate-election) level effect is $\tau_c^P \equiv V_c(1) - V_c(0)$, and the average effect $\tau^P \equiv \mathrm{E}_{\mathscr{C}}[V_c(1) - V_c(0)]$ where $E_{\mathscr{C}}[\cdot]$ indicates that the expectation is taken over $\mathscr{C}$. Since we restrict the population to candidates that ran in the previous election, we know that these are the only potential outcomes (unlike first-time runners that neither won nor lost the previous election), and since we ensure what the candidate always run in the current election, both potential outcomes are defined.

At first glance, it seems that we have shifted focus rather substantively: the unit of observation is no longer parties but specific candidates. There is, however, a link between the two models. Let $q : \mathscr{C} \to \mathscr{I}$ be a mapping from candidates to parties so that $i = q(c)$ is the party-election of candidate-election c. For example, if $c = 1$ is Bill Clinton in the 1996 presidential election, $i = q(1)$ is the Democratic Party in the 1996 presidential election. Naturally, the election outcome of the party and its candidate is the same; if the party wins the election, so will the candidate. This implies that $W_c = W_{q(c)}$. Furthermore, as we hold $R_c = 1$, this implies that $R_c = R_{q(c)} = 1$. Taken together, we have $V_c(1) = Y_{q(c)}(1,1)$ and $V_c(0) = Y_{q(c)}(0,1)$, and thus:

$$\tau^P = \mathrm{E}_{\mathscr{C}}[V_c(1) - V_c(0)] = \mathrm{E}_{\mathscr{C}}[Y_{q(c)}(1,1) - Y_{q(c)}(0,1)].$$

Noting that there is a one-to-one correspondence between $\mathscr{I}$ and $\mathscr{C}$, so that for every party there is a candidate and for every candidate there is a party, we

[6]A pressing question, which we return to in later sections, is whether this can be done in a way to maintain the consistency assumption.

can further simplify expression to:

$$\mathrm{E}_{\mathscr{C}}[Y_{q(c)}(1,1) - Y_{q(c)}(0,1)] = \mathrm{E}[Y_i(1,1) - Y_i(0,1)],$$

where the last expectation is taken over $\mathscr{I}$. In other words, the personal incumbency effect, as defined in this section, can also be defined using parties potential outcomes.[7] The definition of the effect is illustrated in Figure 4.

2.5 The direct party incumbency effect

> *The effect on election outcomes for parties when running as incumbent party when the previous candidate does not run for office.*

The main reason why the RDD estimates of incumbency advantage is unlikely to be directly informative of the legislator effect is that parties by themselves could enjoy advantages (or disadvantages) of being the incumbent party. For example, even if the previous candidate does not run for office, he or she could be of help in the first-time candidate's campaign. Similarly, some of the added media coverage might be directed to the incumbent party rather than to its candidate. The direct party incumbency effect is intended to capture these factors.

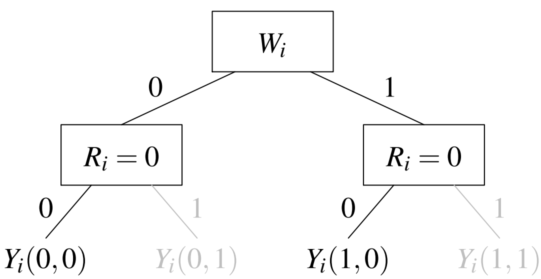

Figure 5. The direct party incumbency effect. The second level nodes (R_i) are here restricted so we only take the paths of $R_i = 0$.

We want here to compare incumbency of a party, when the party does not have an incumbent candidate, to non-incumbency. The potential outcome that refers to party incumbency without candidate incumbency is unambiguously $Y_i(1,0)$. The other potential outcome is less clear, as both $Y_i(0,0)$ and $Y_i(0,1)$ refer to situations where neither the party nor the candidate are incumbents. However, in $Y_i(0,1)$ the candidate has some election experience which the

[7] This result hinges on the one-to-one correspondence between $\mathscr{I}$ and $\mathscr{C}$. This will not hold if parties run with multiple candidates to the same office or if candidates run without a party.

candidate in $Y_i(1,0)$ would, on average, not have. As this type of experience effect arguably is not part of the direct party advantage, $Y_i(0,0)$ would seem to be the natural comparison.[8]

The unit level direct party incumbency effect is $\tau_i^D \equiv Y_i(1,0) - Y_i(0,0)$ and the average effect is $\tau^D \equiv \mathrm{E}[Y_i(1,0) - Y_i(0,0)]$.

3 Previous estimands and estimators

The body of research that investigates incumbency effects is vast, ranging from early efforts in purely describing the difference in election outcomes between incumbents and non-incumbents (see, e.g., Cummings, 1966), to the most recent RDDs. It is beyond the scope of this chapter to provide a complete review of the literature. In this section, I will discuss studies that are close to the current setting. Studies that investigates other types of manipulation than whether the party won or the candidate re-runs, for example re-districting as in Ansolabehere et al. (2000), are not considered, as the causal interpretation differ.

Unlike the previous section, we will here consider observed variables. For each party-election, denoted with i, we observe the tuple $(Y_i, W_i, V_i, R_i, \mathbf{X}_i)$ where Y_i is the observed outcome of interest; W_i is whether the party won the election preceding i; V_i is the two party vote share in the preceding election; R_i is whether the party's candidate from the previous election is its candidate also in i; and $\mathbf{X}_i$ is a vector of election, party or candidate covariates, causally unaffected by W_i and R_i. Some of the previous studies do not fit this general setting, in which case additional variables will be introduced on the go.

I want to emphasis that this is not an exercise of whether the previous studies succeed in identifying their effects, I will grant each study their respective identifying assumptions. Rather, it is an investigation of which effects they aim to estimate *given* that their identification holds; in other words, what their estimands are. Nonetheless, as discussed to great length in the literature, and briefly mentioned in this section, the identifying assumptions in some of the studies are quite restrictive.

3.1 Lee (2008)

Lee (2008, 2001) introduced a regression discontinuity design (RDD) to investigate incumbency effects, and most of the recent studies have followed his example. The method exploits the fact that the winner of an election changes discontinuously at the zero percent vote margin (i.e., at the 50% vote share in

[8] There could be other experience effects. One example would be when the party in $Y_i(1,0)$ is able to recruit a more experienced candidate than in $Y_i(0,0)$. This ought, however, to be included in the direct party effect as it, in some sense, runs through the party.

a two party system). As the variation with this method purely is in whether the party won the previous election (i.e., in W_i), the causal model in Lee (2008) is defined assuming manipulation only in this variable. Let $Y_i(1)$ denote the election outcome of the party if it won the previous election, and $Y_i(0)$ if it lost. The effect that is investigated with the RDD is thus $\tau^{RD} \equiv \mathrm{E}[Y_i(1) - Y_i(0)]$.[9] In this section, I will show how τ^{RD} is related to the causal effects defined above.

Unlike the model above, we are not manipulating R_i. It is thus best seen as a post-treatment variable, and as such, it is potentially affected by W_i. That is, a candidate's choice to run in the subsequent election is, in most cases, made after the previous election, and it could be affected by the election outcome. The candidate might, for example, re-run office only when he or she won the previous election. In this case, we would say that the previous victory caused the candidate to run. Let $R_i(1)$ and $R_i(0)$ be indicators of whether the party's candidate run for office when winning the previous election and when losing, i.e., the potential outcomes of the candidate's running status.

We can connect the two models by making an assumption that the potential outcome when we actively manipulate R_i is the same as when we leave it be, given that we manipulate it to its natural level. This is sometimes referred to as the composition assumption (VanderWeele and Vansteelandt, 2009). When it holds, we have $Y_i(1) = Y_i(1, R_i(1))$ and $Y_i(0) = Y_i(0, R_i(0))$, from which it follows:

$$
\begin{aligned}
Y_i(1) - Y_i(0) &= \{R_i(1)Y_i(1,1) + [1 - R_i(1)]Y_i(1,0)\} \\
&\quad - \{R_i(0)Y_i(0,1) + [1 - R_i(0)]Y_i(0,0)\} \\
&= [Y_i(1,0) - Y_i(0,0)] \\
&\quad + \{R_i(1)[Y_i(1,1) - Y_i(1,0)] \\
&\qquad - R_i(0)[Y_i(0,1) - Y_i(0,0)]\} \qquad (1)
\end{aligned}
$$

Taking expectations then yields:

$$
\begin{aligned}
\tau^{RD} &= \mathrm{E}[Y_i(1) - Y_i(0)] \\
&= \mathrm{E}[Y_i(1,0) - Y_i(0,0)] \\
&\quad + \mathrm{E}\{R_i(1)[Y_i(1,1) - Y_i(1,0)] \\
&\qquad - R_i(0)[Y_i(0,1) - Y_i(0,0)]\}, \\
&= \tau^{D} + \rho^{10}\tau^{L,10} - \rho^{01}\tau^{R,01} + \rho^{11}[\tau^{L,11} - \tau^{R,11}], \qquad (2)
\end{aligned}
$$

[9]In the standard version of the RDD, we are only able to identify the effect exactly at the point where the investigated variable changes discontinuously. Subsequently, the estimand is a local version of $E[Y_i(1) - Y_i(0)]$. I will return to this discussion in later sections, but for the moment I will leave this conditioning implicit to ease exposition.

where,

$$\begin{aligned}\rho^{xy} &= \Pr[R_i(1)=x, R_i(0)=y],\\ \tau^{L,xy} &= \mathrm{E}[Y_i(1,1)-Y_i(1,0)|R_i(1)=x, R_i(0)=y],\\ \tau^{R,xy} &= \mathrm{E}[Y_i(0,1)-Y_i(0,0)|R_i(1)=x, R_i(0)=y].\end{aligned}$$

Despite being a bit cumbersome, the interpretation of expression (2) is rather straightforward. As we manipulate whether the party won the previous election, all units benefit from the direct party incumbency effect (τ^D). This is captured by the first term. The direct party incumbency effect does, however, not account for the fact that some of the parties will also gain incumbent candidates when they win the elections. This effect will depend on exactly how winning the election affects incumbency of the candidate.

Borrowing terminology from the instrumental variable literature (Imbens and Angrist, 1994), we can classify parties into four categories depending on the causal effect of W_i on R_i. A party whose candidate would re-run independently of treatment ($R_i(1)=R_i(0)=1$) will be referred to as an *always-runner*. A party whose candidate runs only if he or she won the previous election ($R_i(1)=1, R_i(0)=0$) is called a *complier*. A *never-runner* is a party with a candidate that would never run in the current election ($R_i(1)=R_i(0)=0$), and a *defier* is a party whose candidate only runs if the previous election was lost ($R_i(1)=0, R_i(0)=1$). Note that all parties fall into exactly one of these four categories.

For never-runner parties (which are of proportion ρ^{00} in $\mathscr{I}$), the party incumbency effect is simply the direct party incumbency effect (τ^D). The previous candidate will never participate in the election, thus any effect must go directly through the party. For parties that are compliers (of proportion ρ^{10}), winning the election *caused* the previous candidate to re-run. In addition to the direct party incumbency effect, these parties will therefore benefit from any legislator incumbency effect ($\tau^{L,10}$). For defiers (of proportion ρ^{01}), winning the election instead causes the party not to run with the previous candidate, or equivalently, losing the election causes the candidate to re-run. When winning, they are thereby affected by the negative of the re-running loser effect ($-\tau^{R,01}$).

Always-runners (of proportion ρ^{11}) are not affected by winning through its effect on whether the candidate re-runs—the candidate runs in both cases. It will, however, have an indirect effect as the setting in which the candidate runs changes; the previous election result will determine if he or she runs as a winner or a loser. In other words, when these parties win elections, they gain incumbent candidates, and avoid losing candidates, in the subsequent election. Like compliers, when winning elections, these parties are benefited by the legislator incumbency effect. However, like defiers, these parties are also affected by the re-running loser effect when losing elections. The total effect is the difference between the two effects ($\tau^{L,11} - \tau^{R,11}$).

Weighting these four effects with the respective proportion in $\mathscr{I}$ produces (2). All types of parties need, however, not to exist in all settings. This would be reflected with a corresponding weight of zero. For example, we suspect that defiers are rare—nearly always, winning the election would make it easier to gain the party's nomination in the next election. This implies that ρ^{01} is zero in many settings. Defiers are, however, not impossible. Candidates that run for the state legislature might, when winning, try to get elected to Congress in the subsequent election, while when losing, this option might not be available and they instead run for the state legislature again.

Net of the direct party incumbency effect, the effects are conditional. If we assume that retirement decisions are ignorable (as sometimes done in the literature), the conditioning does not matter. This is, however, not very likely. For example, we would expect the candidates that manage to re-run for office despite losing the previous election to be of higher than average quality, thus $\tau^{L,11} > \tau^{L}$ and $\tau^{R,11}, \tau^{R,01} > \tau^{R}$. In fact, most of the effect that is captured in the conditional versions of the legislator incumbency and re-running loser effects might not be effects of incumbency or election experience as usually understood, but rather a selection effect as discussed in the preceding section. While they are still causal effects, their interpretations are rather different.

3.2 The Sophomore Surge and the Retirement Slurp

The Sophomore Surge is estimated by comparing the election outcome of a newly elected official, in his or her first winning election, with the election outcome in the subsequent election. In the first election, the candidate did not run as an incumbent, while in the second election he or she did. The candidate, thereby, only enjoys any benefits that incumbency provides in the second election. Since the identity of the candidate is unchanged, the argument is that the change in election outcomes must be due to the incumbency effect. As previously noted (Erikson, 1971; Gelman and King, 1990), this comparison is unlikely to capture a causal effect. There are mainly three issues that could bias the results. First, we condition the analysis on that the candidate won the first election, thereby introducing a regression toward the mean artifact that would lead to a negative bias. Second, we also condition the analysis on that the candidate re-runs for office. It is conceivable that the candidate to some degree can forecast the election results and, when suspecting a negative outcome, withdraw his or her candidacy instead of suffering the expected humiliating defeat. This would introduce a positive bias. Last, implicit in the analysis is a stability assumption that if the candidate had not won the first election, the outcome in the second election would (on average) have been the same as in the first. In this section, I will disregard all of these pressing issues and focus on which effect the Sophomore Surge estimator corresponds to given that it could identify it.

The Sophomore Surge has been interpreted in several ways in the past literature. Erikson (1971) and Caughey and Sekhon (2011) seem to see it as a measure of the personal incumbency effect, as defined here—i.e., the effect of incumbency for specific candidates. Gelman and King (1990) on the other hand interpret it both as a (biased) estimator of the legislator incumbency effect (p. 1145) and what they call the personal incumbency effect (p. 1153). Their definition of the personal effect differs, however, in many ways from the current. I argue that the Sophomore Surge estimand is best seen as a mix of the legislator and direct party incumbency effects. In order to separate the identification problems from the definition of the estimand, I will, for illustration, presume that whether a candidate wins is a deterministic function of his or her characteristics (thus solving regression to the mean), that whether a candidate re-runs is random (solving strategic resigning) and that the potential outcomes are constant over time for each candidate.

Let $E[Y_{i,t}-Y_{i,t-1}|W_{i,t}R_{i,t}=1,R_{i,t-1}=0,W_{j(i),t-1}R_{j(i),t-1}=0]$ be the population quantity that the Sophomore Surge estimator tries to estimate. As above, $j(i)$ gives the opposing party of i in a two party election. The variables are defined as above, but with a time index for clarity. Specifically, $Y_{i,t-1}$ is the election outcome in the first election and $Y_{i,t}$ the outcome in the second. Conditioning on $R_{i,t-1}=0$ ensures that the candidate is a first-time runner in $t-1$. Which, together with $W_{j(i),t-1}R_{j(i),t-1}=0$, ensures that the election was an open-seat election. Finally, $W_{i,t}R_{i,t}=1$ ensures that the candidate won the first election and re-ran for office. In other words, the conditioning set gives us the Sophomore Surge estimator. For brevity, let $\mathscr{A}$ collect all units fulfilling these conditions:

$$\mathscr{A}=\{i\in\mathscr{I}:W_{i,t}R_{i,t}=1,R_{i,t-1}=0,W_{j(i),t-1}R_{j(i),t-1}=0\}.$$

Now consider which potential outcomes the observed outcomes correspond to. Due to the condition $W_{i,t}R_{i,t}=1$ (i.e., that the party won the election in $t-1$ and that its candidate re-runs in t), we have that $Y_{i,t}$ is realization of the potential outcome $Y_i(1,1)$. We can drop the time index as we have assumed that the potential outcomes are stable over time. It is, however, not unambiguous which potential outcome $Y_{i,t-1}$ corresponds to. For the election in $t-1$, we only condition on that the candidate is a first-time runner ($R_{i,t-1}=0$), but we do not restrict whether the party won the election preceding the one in $t-1$; that is, $W_{i,t-1}$ could be both one and zero. Thus, the outcome can be a realization of both $Y_i(1,0)$ and $Y_i(0,0)$. Let $\gamma=Pr[W_{i,t-1}=1|\mathscr{A}]$ be the proportion of elections in $t-1$ where parties in $\mathscr{A}$ won the preceding election.[10] With the law of iterated expectations, we have:

$$E[Y_{i,t-1}|\mathscr{A}]=\gamma\mathrm{E}[Y_i(1,0)|\mathscr{A},W_{i,t-1}=1]+(1-\gamma)\,\mathrm{E}[Y_i(0,0)|\mathscr{A},W_{i,t-1}=0],$$

[10] Due to the conditioning, this proportion need not be 50% as we otherwise would expect.

and thus:

$$
\begin{aligned}
\mathrm{E}[Y_{i,t} - Y_{i,t-1}|\mathscr{A}] &= \mathrm{E}[Y_i(1,1)|\mathscr{A}] - \gamma \mathrm{E}[Y_i(1,0)|\mathscr{A}, W_{i,t-1} = 1] \\
&\quad - (1-\gamma)\,\mathrm{E}[Y_i(0,0)|\mathscr{A}, W_{i,t-1} = 0], \\
&= \mathrm{E}[Y_i(1,1)|\mathscr{A}] - \gamma \mathrm{E}[Y_i(1,0)|\mathscr{A}, W_{i,t-1} = 1] \\
&\quad - (1-\gamma)\,\mathrm{E}[Y_i(0,0)|\mathscr{A}, W_{i,t-1} = 0] \\
&\quad + \mathrm{E}[Y_i(1,0)|\mathscr{A}] - \mathrm{E}[Y_i(1,0)|\mathscr{A}], \\
&= \mathrm{E}[Y_i(1,1) - Y_i(1,0)|\mathscr{A}] \\
&\quad + (1-\gamma)\,\mathrm{E}[Y_i(1,0) - Y_i(0,0)|\mathscr{A}, W_{i,t-1} = 0] \\
&= \mathrm{E}[\tau_i^L|\mathscr{A}] + (1-\gamma)\,\mathrm{E}[\tau_i^D|\mathscr{A}, W_{i,t-1} = 0].
\end{aligned}
$$

The implicit estimand is a mixture of conditional versions of the legislator and direct party incumbency effects, where the proportion depends on the specific election setting (i.e., on γ). Alas, even if identification was unproblematic, the interpretation of the estimator's effect is not obvious. This vagueness could possibly explain why different scholars have interpreted it in different ways.

The Retirement Slump estimator is the difference between the election result of an incumbent candidate in his or her last election before retirement, and the result of the party in the subsequent election when running with a first-time candidate. An investigation similar to that for the Sophomore Surge would reveal that, granted identification, the Retirement Slump estimates a conditional version of the legislator incumbency effect. Intuitively, in the first election where the incumbent candidate runs for office, the outcome is a realization of $Y_i(1,1)$. In the subsequent election, the candidate resigns ($R_i = 0$), but the party still won the past election ($W_i = 1$); as a result, the outcome corresponds to the potential outcome $Y_i(1,0)$. Their difference would be the legislator incumbency effect for the units included in the comparison.

3.3 Gelman and King (1990)

The estimand in Gelman and King (1990), and those in studies adapted from their model (Cox and Katz, 1996; Levitt and Wolfram, 1997), is, as previous mentioned, a version of the legislator incumbency effect. As the authors provide a causal model similar to the current, the connection is quite direct. Specifically, they define (p. 1143) their potential outcome when incumbent ($w^{(I)}$ in their notation) as the "proportion of the vote received by the incumbent legislator in his or her district." This corresponds directly to $Y_i(1,1)$ above. Their potential outcome when not incumbent ($w^{(O)}$ in their notation) is defined as the "proportion of the vote received by the incumbent party in [the same] district, if the incumbent legislator does not run [...]" Clearly, they imagined a treatment where we held victory in the previous election constant; the potential outcome corresponds unambiguously to $Y_i(1,0)$.

The unit level effect in the model of Gelman and King (1990) is, thus, the same as τ_i^L. However, they aggregate the units' effects not by averaging over all parties in all elections, but only over the Democratic Party. As a result, it is not obvious how their estimand is connected to τ^L. As we will see, the effect is a conditional version of τ^L, namely legislator effect for winning parties.[11]

To see this, we will turn to the estimator of Gelman and King (1990) and grant it its identifying assumptions. Their estimator tries to model the conditional expectation function of the Democratic Party's election outcome based on the previous election winner and incumbency status of the candidate. Let D_i be an indicator taking on value 1 if i is the Democratic Party, and 0 otherwise. Let P_i be an indicator taking on value 1 if the Democratic Party won the preceding election, and -1 otherwise. Finally, let I_i be an indicator of incumbency status, where value 1 indicates that the Democratic Party has an incumbent candidate, value -1 that the Republican Party has an incumbent candidate and value 0 if neither party has an incumbent candidate. The population function that Gelman and King (1990) estimates is, in our notation:

$$E[Y_i|D_i = 1, P_i = p, I_i = \ell] = \beta_0 + \beta_2 p + \psi\ell,$$

where ψ is the coefficient intended to capture the legislator effect.[12]

Gelman and King (1990) make two assumptions that will be used in the current investigation. They first assume (p. 1143) that the average incumbency effects for Democrats and Republicans are the same. As they note, this assumption is not necessary but will simplify the investigation. The second assumption (p. 1152) is that the decision to re-run (R_i in our notation) is exogenous. While slightly stronger than necessary, I will operationalize these assumptions so that D_i and R_i is mean independent of the potential outcomes. With these two assumptions, we can decompose the following conditional expectation function:

$$\begin{aligned}
E[Y_i|D_i = d, W_i = 1, R_i = r] = &\\
&= rE[Y_i(1,1)|D_i = d, W_i = 1, R_i = 1]\\
&\quad + (1-r)E[Y_i(1,0)|D_i = d, W_i = 1, R_i = 0],\\
&= rE[Y_i(1,1)|W_i = 1] + (1-r)E[Y_i(1,0)|W_i = 1],\\
&= E[Y_i(1,0)|W_i = 1] + rE[Y_i(1,1) - Y_i(1,0)|W_i = 1],\\
&= \alpha + \tau^{L,1} r, \qquad (3)
\end{aligned}$$

where $\alpha = E[Y_i(1,0)|W_i = 1]$ and $\tau^{L,1} = E[Y_i(1,1) - Y_i(1,0)|W_i = 1]$. The second equality follows from mean independence of D_i and R_i.

[11] This fact is alluded to in their footnote 5. The purpose of the current section is thus only to make this fact explicit.

[12] They also include a covariate of the vote share of the Democratic Party which I omit to ease exposition. Its inclusion might be important for identification but can, when investigating the definitions, safely be disregarded.

Turning again to the expectation function of interest in Gelman and King (1990), note that it can be decomposed as follows:

$$\begin{aligned} &E[Y_i|D_i = 1, P_i = p, I_i = \ell] = \\ &\quad = (1+p)/2E[Y_i|D_i = 1, P_i = 1, I_i = \ell] \\ &\quad\quad + (1-p)/2E[Y_i|D_i = 1, P_i = -1, I_i = \ell]. \end{aligned} \quad (4)$$

We will investigate these two terms separately. Starting with the first term, note that since we condition on the Democratic Party ($D_i = 1$) and that the Democratic Party won ($P_i = 1$), we have $W_i = 1$. Furthermore, when the Democratic Party won, I_i depends on whether the Democratic candidate reruns. Subsequently, if $P_i = 1$, I_i will be equal to R_i.[13] This implies, together with (3), that:

$$\begin{aligned} E[Y_i|D_i = 1, P_i = 1, I_i = \ell] &= E[Y_i|D_i = 1, W_i = 1, R_i = \ell], \\ &= \alpha + \tau^{L,1}\ell, \end{aligned}$$

The second term in (4) is slightly trickier. We will again make use of the function $j(i)$ that maps to the opposing party of i. In a two-party election (or if the outcome is defined as the share of the two-party vote), we have $Y_i = (1 - Y_{j(i)})$. Furthermore, in the sample of Gelman and King (1990), the opposing party of Democrats is always Republican, and *vice versa*, thus $D_i = (1 - D_{j(i)})$. Since P_i and I_i are election specific variables, rather than party specific, we have: $P_i = P_{j(i)}$ and $I_i = I_{j(i)}$. We can therefore express the second term as:

$$\begin{aligned} &E[Y_i|D_i = 1, P_i = -1, I_i = \ell] = \\ &\quad = E[(1 - Y_{j(i)})|(1 - D_{j(i)}) = 1, P_{j(i)} = -1, I_{j(i)} = \ell], \\ &\quad = 1 - E[Y_{j(i)}|D_{j(i)} = 0, P_{j(i)} = -1, I_{j(i)} = \ell], \\ &\quad = 1 - E[Y_i|D_i = 0, P_i = -1, I_i = \ell], \end{aligned}$$

where the last equality follows from that $j(i)$ is a permutation of the party indices. Similarly to the first term, when we condition on the Republican Party ($D_i = 0$) and on that it won the previous election ($P_i = -1$), we have that $W_i = 1$. Furthermore, when $P_i = -1$ and $D_i = 0$, we have that $I_i = -R_i$. Again with (3), this yields:

$$\begin{aligned} 1 - E[Y_i|D_i = 0, P_i = -1, I_i = \ell] &= 1 - E[Y_i|D_i = 0, W_i = 1, R_i = -\ell], \\ &= 1 - (\alpha - \tau^{L,1}\ell), \\ &= 1 - \alpha + \tau^{L,1}\ell. \end{aligned}$$

[13] More formally, we have $(D_i = 1, P_i = 1) \Leftrightarrow (D_i = 1, W_i = 1)$ and $(D_i = 1, P_i = 1) \Rightarrow (I_i = R_i)$.

Substituting the terms in (4) with the derived expressions, we get:

$$\begin{aligned}
E[Y_i|D_i = 1, P_i = p, I_i = \ell] = \\
&= (1+p)/2\, E[Y_i|D_i = 1, P_i = 1, I_i = \ell] \\
&\quad + (1-p)/2\, E[Y_i|D_i = 1, P_i = -1, I_i = \ell], \\
&= (1+p)/2\,(\alpha + \tau^{L,1}\ell) \\
&\quad + (1-p)/2\,(1-\alpha+\tau^{L,1}\ell), \\
&= 0.5 + (\alpha - 0.5)\,p + \tau^{L,1}\ell.
\end{aligned}$$

Comparing the coefficients in this version of the conditional expectation function with the coefficients specified by Gelman and King (1990), we see that $\beta_0 = 0.5$, $\beta_2 = (\alpha - 0.5)$ and $\psi = \tau^{L,1}$. In other words, their estimand is our legislator effect conditioned on being the winning party:

$$E[Y_i(1,1) - Y_i(1,0)|W_i = 1]. \tag{5}$$

As noted in previous sections, this estimand might differ quite substantially from the unconditional version and will partly capture different mechanisms.

3.4 Erikson and Titiunik (2013)

In a recent working paper by Erikson and Titiunik (2013), the personal incumbency effect is investigated using a regression discontinuity design. To my knowledge this is the only estimator, apart from the Sophomore Surge, that is intended to investigate the personal incumbency effect. In this section, I will investigate their strategy using the current causal model. The exercise will reveal that their estimand is best interpreted as a legislator incumbency effect. While the term "personal incumbency advantage" sometimes been used to refer to the incumbent legislator effect in the previous literature, the authors contrast their estimand with Gelman and King (1990), so to my reading, it is intended to capture an effect similar to what I refer to as the personal incumbency effect.

The authors model the conditional expectation function of the Democratic vote share infinitesimally close to the RDD cut-off as:

$$\lim_{v \downarrow 0.5} E[Y_i|V_i = v, D_i = 1, I_i = \ell] \;=\; Par^w + (\theta+\sigma)\ell, \tag{6}$$

$$\lim_{v \uparrow 0.5} E[Y_i|V_i = v, D_i = 1, I_i = \ell] \;=\; Par^l + (\theta+\sigma)\ell, \tag{7}$$

where V_i is the vote share of the party denoted by i in the election preceding i, and the other variables are defined as above. I have dropped the time index since it does not affect the analysis. Par^w and Par^l are interpreted as the average baseline vote for the Democratic Party (i.e., in absence of an incumbent candidate) and $(\theta+\sigma)$ is intended to capture the personal incumbency effect, which consists of the direct personal incumbency effect (θ) and the scare-off effect (σ).[14]

Note that whether the party won the preceding election (W_i) is a deterministic function of the vote share, we therefore have:

$$\begin{aligned}\lim_{v\downarrow 0.5} E[Y_i|V_i=v,D_i=1,I_i=\ell] &= \lim_{v\downarrow 0.5} E[Y_i|V_i=v,W_i=1,D_i=1,I_i=\ell],\\ \lim_{v\uparrow 0.5} E[Y_i|V_i=v,D_i=1,I_i=\ell] &= \lim_{v\uparrow 0.5} E[Y_i|V_i=v,W_i=0,D_i=1,I_i=\ell].\end{aligned}$$

Erikson and Titiunik (2013) make three assumptions that we will use. First, they make the simplifying assumption (p. 10 in the online Appendix) that the personal incumbency effect is the same for both Democrats and Republicans. Second, they assume (p. 13) that the candidate's decision to re-run is non-strategic. Last, that the RDD assumptions holds (p. 12), which implies that W_i is ignorable at the cut-off. I will again operationalize these so that D_i, R_i and W_i are mean independent of the potential outcomes at the RD cut-off. This gives us:

$$\begin{aligned}&\lim_{v\downarrow 0.5} E[Y_i|V_i=v,D_i=d,R_i=r] =\\ &\quad = r\lim_{v\downarrow 0.5} E[Y_i|V_i=v,D_i=d,R_i=1]\\ &\qquad +(1-r)\lim_{v\downarrow 0.5} E[Y_i|V_i=v,D_i=d,R_i=0]\\ &\quad = rE[Y_i(1,1)|V_i=0.5]\\ &\qquad +(1-r)E[Y_i(1,0)|V_i=0.5]\\ &\quad = E[Y_i(1,0)|V_i=0.5]\\ &\qquad +rE[Y_i(1,1)-Y_i(1,0)|V_i=0.5]\\ &\quad = \alpha^{rd}+\tau^{L,rd}r,\end{aligned}$$

where $\alpha^{rd}=E[Y_i(1,0)|V_i=0.5]$ and $\tau^{L,rd}=E[Y_i(1,1)-Y_i(1,0)|V_i=0.5]$. This expression will aid us translate the conditional expectation functions in Erikson and Titiunik (2013) to the current causal model. Note that, like in the previous section, $(W_i=1,D_i=1)$ implies $I_i=R_i$. It follows that:

$$\begin{aligned}\lim_{v\downarrow 0.5} E[Y_i|V_i=v,D_i=1,I_i=\ell] &= \lim_{v\downarrow 0.5} E[Y_i|V_i=v,D_i=1,R_i=\ell],\\ &= \alpha^{rd}+\tau^{L,rd}\ell. \qquad (8)\end{aligned}$$

[14]The expressions on page 13 in Erikson and Titiunik (2013) have the quality differentials, D^w-R^w and D^l-R^l in their notation, rather than σ. However, on the following pages, they state that in their setting $D^w=R^l=0$. They also define $R^w=D^l=-\sigma$ when there is an incumbent ($I_i\neq 0$), and $R^w=D^l=0$ when there is not ($I_i=0$). As we will see, in the first function we always have $I_i\in\{0,1\}$, while in the second $I_i\in\{0,-1\}$. As a result, the quality differentials in both equations are equivalent with σI_i.

Comparing the definition of Erikson and Titiunik (2013) in (6) with the derived expression in (8), we see that $(\theta+\sigma)=\tau^{L,rd}$ for the upper limit of the RDD estimator.

Continuing with the lower limit, we use the function $j(i)$ which maps to the opposing party of i:

$$\begin{aligned}
&\lim_{v\uparrow 0.5} E[Y_i|V_i=v, D_i=1, I_i=\ell] = \\
&\qquad = \lim_{v\uparrow 0.5} E[(1-Y_{j(i)})|(1-V_{j(i)})=v, D_{j(i)}=0, I_{j(i)}=\ell], \\
&\qquad = 1-\lim_{v\uparrow 0.5} E[Y_{j(i)}|V_{j(i)}=1-v, D_{j(i)}=0, I_{j(i)}=\ell], \\
&\qquad = 1-\lim_{v\downarrow 0.5} E[Y_{j(i)}|V_{j(i)}=v, D_{j(i)}=0, I_{j(i)}=\ell], \\
&\qquad = 1-\lim_{v\downarrow 0.5} E[Y_i|V_i=v, D_i=0, I_i=\ell],
\end{aligned}$$

where the last equality follows from that $j(i)$ is a permutation of party indices. Recognizing that $(W_i=1, D_i=0)$ implies $I_i=-R_i$, we have:

$$\begin{aligned}
&1-\lim_{v\downarrow 0.5} E[Y_i|V_i=v, D_i=0, I_i=\ell] = \\
&\qquad = 1-\lim_{v\downarrow 0.5} E[Y_i|V_i=v, D_i=0, R_i=-\ell], \\
&\qquad = 1-\alpha^{rd}+\tau^{L,rd}\ell. \qquad (9)
\end{aligned}$$

By comparing (7) with (9), we see that $(\theta+\sigma)=\tau^{L,rd}$ also for the lower limit. Subsequently, under their assumptions, the parameter of interest is not the personal incumbency effect, as defined here, but rather the legislator incumbency effect.

4 Local identification with experimental variation in W_i

In this section, I will investigate which of the defined effects can be identified when the assignment of W_i is ignorable. One such situation would be in an RDD setting, where W_i is ignorable at the cut-off; another would be if we somehow can randomly assign W_i; and yet another, if W_i is ignorable conditional on a set of covariates. For the moment, I will not further specify exactly why W_i is ignorable, in order to keep the analysis simple. At the end of this section, I will discuss the particularities when ignorability of W_i is gained through an RDD, which also is the setting of the illustration in the last section.

First, note that when only W_i is ignorable, R_i is best seen as a post-treatment variable. I will specify two potential outcomes of R_i in the same way as in Section 3.1. Specifically, let $R_i(0)$ be an indicator of whether the candidate re-ran when the party lost the previous election, and let $R_i(1)$ be indicator whether he or she ran when the party won. If we assume that $(R_i(0), R_i(1))$ are independent of the potential outcomes, our task is simple: we have $E[Y_i(x,y)] =$

$E[Y_i|W_i = x, R_i = y]$ and identification is trivial. The personal incumbency effect could, for example, be identified with:

$$\begin{aligned}\tau^P = E[Y_i(1,1) - Y_i(0,1)] &= E[Y_i(1,1)|W_i = 1, R_i(1) = 1] \\ &\quad - E[Y_i(0,1)|W_i = 0, R_i(0) = 1], \\ &= E[Y_i|W_i = 1, R_i = 1] - E[Y_i|W_i = 0, R_i = 1].\end{aligned}$$

This assumption is, however, unlikely to hold. Consider, for example, a situation where there are high and low quality candidates. Assume that high quality candidates tend to re-run for office both when they win and lose the election, while weak candidates only do so when they win. Now, consider the identification strategy of τ^P in the previous paragraph. $E[Y_i(1,1)|W_i = 1, R_i = 1]$ would consist of both high and low quality candidates, whereas $E[Y_i(0,1)|W_i = 0, R_i = 1]$ consists only (or mostly) of high quality candidates. If the quality of the candidate matter for election performance, this contrast will not have a causal interpretation.

We could, sometimes greatly, reduce the severity in this assumption by condition on a set of covariates; thereby, we only require conditional independence of $(R_i(0), R_i(1))$. While one of the identification strategies I present uses a weak version of conditional independence, I will start by asking what one could do when R_i is in no way ignorable.

The situation is not unlike that of an instrumental variable (IV) studied by Imbens and Angrist (1994). In both cases, we have post-treatment variable that is not in our direct control (in our case, it is R_i, with an IV, it is the treatment variable), but the variable is affected by another that are under our control (in our case, W_i, with an IV, the instrument). Like the IV setting, we can only observe the values of R_i that is given by W_i, and, as a result, we are restricted to investigate the effect only for units which are affected by W_i in a particular way. In other words, we can only study the effect conditionally on the causal effect of W_i on R_i: a local average treatment effect (LATE).

Unlike the IV setting, however, we cannot safely assume that the "instrument" (i.e., W_i) has no direct effect on the outcome. In this setting, we suspect that winning the previous election potentially will have large effects on subsequent results. This rules out investigating the effect of R_i using W_i as an instrument. The only way we would have variation in R_i is through variation in W_i. In the complier and defier groups, the two variables will be perfectly correlated, and we can, therefore, impossibly separate the two effects. As a result, with exogenous variation only in W_i, we cannot identify the legislator and losing re-runner effects.

The personal and direct party incumbency effects *do not* require variation in R_i. On the contrary, they require that R_i is fixed. Assume for the moment that we can observe $(R_i(0), R_i(1))$ for all units.[15] Consider a conditional version

[15] Not even in this setting could we identify the legislator and losing re-runner effect without additional assumptions.

of the personal incumbency effect:

$$\begin{aligned}
\tau^{P,11} &\equiv E[Y_i(1,1) - Y_i(0,1)|R_i(1) = R_i(0) = 1], \\
&= E[Y_i(1,1)|W_i = 1, R_i(1) = R_i(0) = 1] \\
&\qquad - E[Y_i(0,1)|W_i = 0, R_i(1) = R_i(0) = 1], \\
&= E[Y_i|W_i = 1, R_i(1) = R_i(0) = 1] \\
&\qquad - E[Y_i|W_i = 0, R_i(1) = R_i(0) = 1].
\end{aligned}$$

Since the effect does not consider variation in R_i, we can identify it solely with ignorability in W_i—given that we observe the potential outcomes of R_i. With the terminology from the previous sections, we can potentially identify the personal incumbency advantage for parties that are always-runners. Obviously, we do not observe both potential outcomes of R_i—only the realized values. However, we can identify the effect that W_i has on R_i, and thereby we can possibly gain enough traction to identify the effect of interest. In the following subsections, I will investigate under which assumptions we can identify this effect. The exercise results in three identification strategies. When describing the strategies, I will focus on the personal incumbency effect. With minor changes, the strategies could, however, be used to investigate the direct party incumbency effect—which is also done in the illustration of the methods.

With the first strategy, which I will refer to as *always-runner stratification*, we will try to identify covariate strata which contain only always-runners. This strategy does not require any additional assumptions, but the identified effect is for an even smaller subpopulation than for the always-runners. In fact, depending on the exact covariates used in the analysis, the subpopulation might not contain a single unit.

With the second strategy, *non-complier stratification*, I will make a monotonicity assumption (similar to the one made with in IV). The assumption requires that the directionality of the effect of W_i on R_i is the same for all units—e.g., $R_i(0) \leq R_i(1)$ for all i. With this assumption, the effect can be identified in strata which do not contain any compliers, thus a larger part of the population than with the previous strategy.

The last strategy, *running-on-observables*, imposes an independence assumption, in addition to the monotonicity assumption, where one of the potential outcomes of R_i is assumed to be conditionally independent of $Y_i(1,1)$. This is a strong assumption, but still weaker than the typical independence assumptions; independence is only needed with respect to one of the potential outcomes. With this approach, the effect is identified for the complete subpopulation of always-runners.

This analysis bear close resemblance to the problems with *principal stratification*, as discussed in Frangakis and Rubin (2002). When considering the candidate as the unit of observation, the realization of the post-treatment variable (i.e., R_i in this case) determines whether we can observe the outcome of interest: if the candidate does not re-run, we will not observe the vote share in the election he or she did not participate in. Viewed from this perspective, the issue is not an inferential problem, but rather a definitional concern. That is, the personal incumbency effect might not even be defined for the complete population. With this interpretation, the reason we focus on always-runners is because these are the only units for which the personal effect is well-defined. With candidates as the units of observation, we must directly manipulate R_i in order for the personal effect to be defined for all candidates. But, such manipulations are hard to imagine in the general case; how do we force a retired, or dead, candidate to run in an election? While the two perspectives lead to different interpretations of the effects, both would result in empirical strategies similar to those presented here.

4.1 Always-runner stratification

With this strategy, we restrict our attention to a small part of the always-runners, namely those that are in covariate strata with only other always-runners. By limiting our focus to this group, we can identify the effect without additional assumptions.

Let $\mu_1(\mathbf{x}) = E[R_i(1)|\mathbf{X}_i = \mathbf{x}]$ be the fraction of parties in stratum $\mathbf{x}$ whose candidates would re-run for office when winning the preceding election. If $\mu_1(\mathbf{x}) = 1$, all parties' candidates in the corresponding stratum will re-run when their parties' win—the stratum consist of only always-runners and compliers. Similarly, let $\mu_0(\mathbf{x}) = E[R_i(0)|\mathbf{X}_i = \mathbf{x}]$ be the fraction of re-runners in the stratum when the parties lose their elections. If $\mu_0(\mathbf{x}) = 1$, the stratum consists only of always-runners and defiers. Combining the two, we get that strata with $\mu_1(\mathbf{x}) = \mu_0(\mathbf{x}) = 1$ consist of only always-runners. Let $\mathscr{A} = \{\mathbf{x} : \mu_1(\mathbf{x}) = \mu_0(\mathbf{x}) = 1\}$ be the set of all covariate vectors that correspond to strata containing only always-runners.

The estimand we focus on with this strategy is the personal incumbency for units within these strata, namely:

$$\tau^{P,A} \equiv E[Y_i(1,1) - Y_i(0,1)|\mathbf{X}_i \in \mathscr{A}].$$

Showing identification is fairly straightforward. Since W_i is ignorable, we get:

$$\tau^{P,A} = E[Y_i(1,1)|W_i = 1, \mathbf{X}_i \in \mathscr{A}] - E[Y_i(0,1)|W_i = 0, \mathbf{X}_i \in \mathscr{A}].$$

Remember that for all units with covariates in $\mathscr{A}$, we have $R_i(1) = R_i(0) = 1$:

$$\begin{aligned}\tau^{P,A} &= E[Y_i(1,1)|W_i = 1, R_i(1) = R_i(0) = 1, \mathbf{X}_i \in \mathscr{A}] \\ &\quad - E[Y_i(0,1)|W_i = 0, R_i(1) = R_i(0) = 1, \mathbf{X}_i \in \mathscr{A}], \\ &= E[Y_i|W_i = 1, \mathbf{X}_i \in \mathscr{A}] - E[Y_i|W_i = 0, \mathbf{X}_i \in \mathscr{A}].\end{aligned}$$

In other words, if we know $\mathscr{A}$, we can identify $\tau^{P,A}$.

In some cases, we might have *a priori* knowledge about $\mathscr{A}$, but it is seldom complete.[16] The set can, however, be identified from the observable variables. Note that since W_i is ignorable, we have:

$$\begin{aligned}\mu_1(\mathbf{x}) &= E[R_i(1)|\mathbf{X}_i = \mathbf{x}] \\ &= E[R_i(1)|\mathbf{X}_i = \mathbf{x}, W_i = 1], \\ &= E[R_i|\mathbf{X}_i = \mathbf{x}, W_i = 1].\end{aligned}$$

A similar exercise can be done with $\mu_0(\mathbf{x})$. As both $\mu_1(\mathbf{x})$ and $\mu_0(\mathbf{x})$ are identified, $\mathscr{A}$ is also identified, which enables us to identify $\tau^{P,A}$.

The main strength of this strategy is that it does not need any identifying assumption (in addition to those that provide ignorability of W_i). However, the estimand is the effect for a very local group of units. There might be, and probably are, strata that consist of a mix of always-runner and other types of units. All these units are discarded with this strategy. As a result $\tau^{P,A}$ might not be the estimand of interest, even if it captures the qualitative concept of interest. In the worst case, $\mathscr{A}$ is empty, and the estimand is then undefined. If the covariates are few and not informative of whether the candidate re-runs, this could happen even if most units are always-runners.

4.2 Non-complier stratification

With this strategy, I will assume monotonicity: that the causal effect of winning the previous election affects whether the candidate re-runs in the same direction for all parties, in particular that $R_i(1) \geq R_i(0)$. This allows for identification for a greater subpopulation than with the previous strategy.[17]

As a result of the monotonicity assumption, we know that all parties with $R_i(0) = 1$ are always-runners. The only other type of parties with $R_i(0) = 1$ are defiers, but monotonicity ensures that they do not exist. However, parties with $R_i(1) = 1$ still consist of both always-runners and compliers. Under this assumption, $\mu_1(\mathbf{x})$ is the proportion of always-runners and compliers in stratum $\mathbf{x}$, while $\mu_0(\mathbf{x})$ is the proportion of always-runners in the same stratum. As a consequence, whenever $\mu_1(\mathbf{x}) = \mu_0(\mathbf{x})$ the strata contains no compliers. Let $\mathscr{N} = \{\mathbf{x} : \mu_1(\mathbf{x}) = \mu_0(\mathbf{x})\}$ be the set of all covariate vectors that correspond to strata which do not contain any compliers.

The estimand in focus here is the personal incumbency effect for always-runners in these strata:

$$\tau^{P,N} \equiv E[Y_i(1,1) - Y_i(0,1)|R_i(1) = R_i(0) = 1, \mathbf{X}_i \in \mathscr{N}].$$

Identification follows in many ways the same pattern as the previous strategy. Since $\mu_1(\mathbf{x})$ and $\mu_0(\mathbf{x})$ are identified, as previously shown, we have also iden-

[16]For example, as in Section 6, term limits could indicate which units are in $\mathscr{A}$.

[17]The direction of the monotonicity assumption does not matter neither for this or the next strategy, if appropriate changes are made.

tified $\mathscr{N}$. Since all parties with covariates in $\mathscr{N}$ are either always-runners or never-runners (i.e., $R_i(1) = R_i(0)$), observing $R_i = 1$ for a party in $\mathscr{N}$ implies that it is an always-runner. With ignorability of W_i, we have:

$$
\begin{aligned}
\tau^{P,N} &= E[Y_i(1,1)|R_i = 1, \mathbf{X}_i \in \mathscr{N}] \\
&\qquad - E[Y_i(0,1)|R_i = 1, \mathbf{X}_i \in \mathscr{N}], \\
&= E[Y_i(1,1)|W_i = 1, R_i = 1, \mathbf{X}_i \in \mathscr{N}] \\
&\qquad - E[Y_i(0,1)|W_i = 0, R_i = 1, \mathbf{X}_i \in \mathscr{N}], \\
&= E[Y_i|W_i = 1, R_i = 1, \mathbf{X}_i \in \mathscr{N}] \\
&\qquad - E[Y_i|W_i = 0, R_i = 1, \mathbf{X}_i \in \mathscr{N}].
\end{aligned}
$$

Both terms in the last expression can be estimated with observed data, thus the effect is identified.

The monotonicity assumption lets us identify the effect for a subpopulation that is weakly bigger than the previous subpopulation. If a stratum only contains always-runners, as in the first strategy, it naturally contains no compliers; we have $\mathscr{A} \subseteq \mathscr{N}$. The subpopulation is, however, still likely to be small relative to the complete population. It might, therefore, still not be the estimand of ultimate interest. While less likely than before, the worst case is that $\mathscr{N}$ is empty.

4.3 Running-on-observables

With this strategy, I will make a conditional independence assumption which will allow for identification of the effect for the complete subpopulation of always-runners. As always, assuming independence in non-experimental settings is a strong assumption. The assumption needed with this strategy is, however, weaker than the usual ignorability assumption. I will still assume monotonicity, as in the previous section.

Since we investigate the full subpopulation of always-runners, the estimand is as presented above:

$$
\begin{aligned}
\tau^{P,11} &= E[Y_i(1,1) - Y_i(0,1)|R_i(1) = R_i(0) = 1], \\
&= E[Y_i(1,1)|R_i(1) = R_i(0) = 1] - E[Y_i(0,1)|R_i(1) = R_i(0) = 1].
\end{aligned}
$$

Note that the second term of this expression is identified without an independence assumption. With monotonicity, we have that parties with $R_i(0) = 1$ are always-runners, and for all parties that lost the preceding election we observe $R_i(0)$. Consequently, if we observe $R_i = 1$ for a party that lost, it must be an always-runner. Together with ignorability of W_i, this gives us:

$$
\begin{aligned}
E[Y_i(0,1)|R_i(1) = R_i(0) = 1] &= E[Y_i(0,1)|R_i(1) = R_i(0) = 1, W_i = 0], \\
&= E[Y_i|R_i(1) = R_i(0) = 1, W_i = 0], \\
&= E[Y_i|R_i = 1, W_i = 0].
\end{aligned}
$$

Using the law of iterated expectations, it will later prove useful write this as:

$$E[Y_i|R_i=1,W_i=0] = E_{\mathbf{X}}[E(Y_i|R_i=1,W_i=0,\mathbf{X}_i)|R_i=1,W_i=0],$$

where $E_{\mathbf{X}}[\cdot]$ indicates that we take the expectation over the covariate distribution, in this case conditional on $R_i=1$ and $W_i=0$.

We are less fortunate with the first term of the estimand. The monotonicity assumption does not ensure that all units with $R_i=1$ and $W_i=1$ are always-runners—some will be compliers. This is where the independence assumption is needed. We assume that any systematic difference in election outcomes between always-runners and compliers, when winning the previous election, can be described by differences in their covariate distributions. In other words, an always-runner and complier with the same covariate values are expected to have the same election outcome if they won the previous election. The assumption formalized would be:

$$Y_i(1,1) \perp R_i(0)|R_i(1)=1,\mathbf{X}_i,$$

or a mean-independence version thereof. With this assumption, we can identify the first term conditionally:

$$\begin{aligned}
E[Y_i(1,1)|R_i(1)=R_i(0)=1,\mathbf{X}_i] &= E[Y_i(1,1)|R_i(1)=1,\mathbf{X}_i],\\
&= E[Y_i(1,1)|R_i(1)=1,W_i=1,\mathbf{X}_i],\\
&= E[Y_i|R_i=1,W_i=1,\mathbf{X}_i],
\end{aligned}$$

where the first equality follows from the independence assumption, and the second from ignorability of W_i.

The identified quantities are conditional on $\mathbf{X}_i$, but we want the unconditional expectation for always-runners; we need to take the expectation over $\mathbf{X}_i$ for always-runners. The parties with $(R_i=1,W_i=1)$ consist, however, of both always-runners and compliers. That subpopulation cannot inform us about the distribution of $\mathbf{X}_i$ for always-runner. However, parties with $(R_i=1,W_i=0)$ provide that information:

$$\begin{aligned}
&E[Y_i(1,1)|R_i(1)=R_i(0)=1] =\\
&\quad = E_{\mathbf{X}}[E[Y_i(1,1)|R_i(1)=R_i(0)=1,\mathbf{X}_i]|R_i(1)=R_i(0)=1],\\
&\quad = E_{\mathbf{X}}[E[Y_i(1,1)|R_i(1)=R_i(0)=1,\mathbf{X}_i]|R_i(1)=R_i(0)=1,W_i=0],\\
&\quad = E_{\mathbf{X}}[E[Y_i(1,1)|R_i(1)=R_i(0)=1,\mathbf{X}_i]|R_i=1,W_i=0],
\end{aligned}$$

where the first equality follows from the law of iterated expectations, the second from ignorability of W_i and the third from monotonicity. Substituting the inner expectation for the expression we derived above, we get:

$$\begin{aligned}
&E[Y_i(1,1)|R_i(1)=R_i(0)=1] =\\
&\qquad = E_{\mathbf{X}}[E[Y_i(1,1)|R_i(1)=R_i(0)=1,\mathbf{X}_i]|R_i=1,W_i=0],\\
&\qquad = E_{\mathbf{X}}[E(Y_i|R_i=1,W_i=1,\mathbf{X}_i)|R_i=1,W_i=0].
\end{aligned}$$

Finally, by joining the two terms, we have identified the estimand:

$$
\begin{aligned}
\tau^{P,11} &= E[Y_i(1,1)|R_i(1)=R_i(0)=1] - E[Y_i(0,1)|R_i(1)=R_i(0)=1], \\
&= E_{\mathbf{X}}[E(Y_i|R_i=1,W_i=1,\mathbf{X}_i)|R_i=1,W_i=0] \\
&\quad - E_{\mathbf{X}}[E(Y_i|R_i=1,W_i=0,\mathbf{X}_i)|R_i=1,W_i=0], \\
&= E_{\mathbf{X}}[E(Y_i|R_i=1,W_i=1,\mathbf{X}_i) \\
&\qquad\qquad -E(Y_i|R_i=1,W_i=0,\mathbf{X}_i)|R_i=1,W_i=0],
\end{aligned}
$$

4.4 Identification using RDD

In the previous section, I assumed that W_i was globally ignorable. This is, however, seldom the case. An RDD would provide local ignorability, but then we need slight modifications to the analysis. In this section, I will briefly outline how an RDD can be employed to identify the personal incumbency effect.

In the most common set-up, the RDD only requires that the potential outcomes are continuous at the RDD cut-off (Hahn et al., 2001). This weak assumption enables us to identify the effect at the cut-off; we compare the limits of the expected value of the observed outcome conditionally on the running variable as it approaches the cut-off from either side. The added level of complexity, however, makes this route impractical with the current identification strategies. To accurately estimate the limit conditionally on covariates would require more data than we usually are blessed with. To gain more leverage in estimation, I will therefore rely on a slightly stronger assumption to provide identification in the RDD setting.

I will interpret the RDD as a local random experiment similar to the discussion in Lee (2008). Whereas Lee (2008) interpreted the experiment taking place exactly at the cut-off, I will extend the assumption so that we can consider the experiment to take place in a neighborhood around the cut-off. An initiated discussion of the interpretation of the RDD as a localized experiment can be found in Cattaneo, Frandsen and Titiunik (2013), from where I have drawn inspiration for the current set-up.

Using the notation from the previous sections, where V_i denotes the two-party vote share of party i in the previous election, I will assume that W_i and V_i are independent of all potential outcomes in some neighborhood $\mathscr{V}$ around the RD cut-off:

$$
Y_i(1,1), Y_i(0,1), Y_i(1,0), Y_i(0,0), R_i(1), R_i(0) \perp W_i, V_i | V_i \in \mathscr{V}. \tag{10}
$$

With this assumption, the strategies can be used in an RDD setting if we restrict the estimands to units with $V_i \in \mathscr{V}$. With the running-on-observables strategy, this restriction will make the estimand more local. The change in localness is, however, less clear for the other two strategies. On the one hand, there will be fewer parties in each stratum, leading to a more local effect. On the other hand, as we now require that the strata only contain always-runners in the studied neighborhood, the number of admissible strata might increase.

The assumption of local randomness is stronger than the ordinary RDD assumptions. It should, however, be noted that the effect can never be estimated only with units at the cut-off in finite samples. Even if identification is proved at the cut-off, units in the neighborhood of the cut-off must be used for estimation. Oftentimes, this neighborhood is quite large. In practice, the two approaches do not differ as much as one would initially expect. As an example, in the illustration in the following section, I restrict the analysis to either a one or four percentage points vote margin window on either side of the cut-off, in Lee (2008) the smallest window is five percentage points.[18]

Finally, note that the RDD provides a setting where the consistency assumption is reasonable to hold. We would be suspicion about a manipulation that turns a party that had a land-slide victory into a losing party; it would entail such an invasive change of the history of events. Such manipulation would arguably capture much more than just the intended incumbency effect. At or around the RD cut-off, we can, however, imagine manipulations with few side-effects that change the election winner by slightly changing the vote shares. An RDD, thereby, clarify the intended manipulation in the causal model, and the consistency assumption is more reasonable in this setting.

5 Inference

The main focus of this study is in the definitions and identification of incumbency effects, substantially less focus will be given to estimation and hypothesis testing. In this section, I will briefly outline how I will estimate the population quantities of interest (and their distribution under a null hypothesis) in the following illustration.

The two first strategies, the always-runner and non-complier stratification, will be considered as a two-step estimation problem: first estimate the sets of strata, $\mathscr{A}$ and $\mathscr{N}$, and then estimate the effect in these estimated sets. The main challenge, with respect to point estimation, is the first step. Once these sets are found, the effect can be estimated simply by comparing mean responses in the two treatment groups.

Estimating $\mathscr{A}$ could be seen as a type of extreme value estimation: a single non-running unit would exclude a stratum from $\mathscr{A}$. As such, it is far from trivial to estimate. I will opt for a simple solution using a matching-like method akin to kernel regression. For each party with $R_i = 1$ (a potential always-runner), I match it to the k nearest neighbors based on the Mahalanobis distance of its covariates, separately in the treatment and control groups. If *all* these $2k$ units also have $R_i = 1$, the party is considered an always-runner and

[18] In some ways, the current approach can be seen as moving the choice of bandwidth in the RDD from a question about estimation, to a question about identification.

added to $\hat{\mathscr{A}}$ (the matched units are not added unless they also fulfill this condition). Intuitively, under a smoothness condition and asymptotically in sample size ($n \to \infty$), if k grows at a rate so that $k \to \infty$ and $k/n \to 0$, then $\hat{\mathscr{A}}$ should approach $\mathscr{A}$.

With non-complier stratification, we cannot exclude strata based on single observations—not all parties in the strata must be re-runners for it to be included in $\mathscr{N}$. Instead of a non-parametric estimator, I will model the response surfaces of the re-running variable separately for winners and losers—i.e., $\mu_1(\mathbf{x})$ and $\mu_0(\mathbf{x})$—using a logistic function that depends on all covariates and their second power. All parties where the differences between their fitted values of the response surfaces are lower than a small ε are added to $\hat{\mathscr{N}}$. Intuitively, if the parameterizations of the functions are correct and ε approaches zero, $\hat{\mathscr{N}}$ should approach $\mathscr{N}$ asymptotically in the sample size.

With the last strategy, running-on-observables, a more classical matching estimation method can be used. With the monotonicity assumption, we know that all parties with $R_i = 1$ and $W_i = 0$ are always-runners. Each party in this subsample is matched to a party with $R_i = 1$ and $W_i = 1$ based on how similar their covariates are. A point estimate can then be derived by comparing the outcomes between the matched pairs. To construct matches, I use the GenMatch algorithm (Diamond and Sekhon, 2012) with the minimum p-value of paired Fisher's exact tests of all covariates as the balance measure.[19]

Hypothesis testing will exploit the view of the RDD as a local experiment. Specifically, I will use Fisher's exact tests with the treatment group contrast as test statistic when the assignment of W_i is permuted 20,000 times. With the running-on-observables strategy, treatment is permuted within matched pairs. With this approach, the relevant null hypothesis is sharp in the sense that it tests whether there exists any effect of incumbency, rather than the existence of an average effect. Another consequence is that the population that inference is drawn about is the sample at hand, rather than some wider group of parties. Furthermore, the preprocessing steps are disregarded with these tests. If one wants to draw inference to larger groups, the current tests are likely to underestimate the true uncertainty due to both variability in sampling and the preprocessing steps.

The vote shares in an election will, by construction, sum to one. As a result, there is dependence between parties' outcomes in an election. This dependence must be accounted for when drawing inferences. The standard approach, in a two-party system, is to condition the analysis on one of the parties (e.g., restricting the analysis to the Democratic Party as in Lee, 2001, 2008). Since party identity is a covariate—thus unaffected by treatment—this conditioning will not break the causal interpretation. Furthermore, as shown in Section 3, since there is a perfect correlation between the outcomes in a

[19]To speed up calculations, I first run GenMatch with a paired t-test and then refine the resulting matches using Fisher's exact test.

two-party system, the estimate will still capture the average treatment effect for both parties.

When investigating the personal incumbency effect, it is, however, not possible to condition the analysis on one of the parties. As we in this case condition on whether the parties' candidates re-run, the mirroring discussed in Section 3 does not hold. For example, if we condition on the Democratic Party and its candidate does not re-run, the Republican Party in that election will be excluded from the analysis even if its candidate re-runs. Unless we are ready to assume that the effect is the same for all parties, we cannot restrict the analysis to a single party. Fortunately, the use of a sharp null enables us to disregard any dependence that exists between the outcomes. As treatment is assumed to have no effect under the null, no other assignment would have produced different outcomes. Thus, any influence between units would remain constant with any assignment, and the test is valid when we include all parties even if there is dependence.

6 Incumbency effects for Brazilian mayors

To illustrate the discussed identification strategies, I will investigate the incumbency effects in Brazilian mayor elections. Since the 1988 constitution, the more than 5500 Brazilian municipalities have substantial autonomy and the main responsibility of local service provision, including public transport, education and health services (Titiunik, 2011). The executive power of the municipality is wielded by a directly elected mayor (*Prefeito*), while the legislative body (*Câmara Municipal*) consists of a council of elected aldermen (*Vereador*). The mayoral office is, thereby, an important part of the Brazilian political system, and we would expect voters to be highly affected by their mayor's behavior.

Brazilian mayors are elected in the general municipal elections held every four years. In most municipalities, the mayor is elected by a first-past-the-post voting system. In large municipalities (population over 200,000) where no candidate acquires a majority of the total votes, a runoff election is conducted between the two leading candidates from the first round. A candidate can serve as mayor for at most two consecutive terms.

I will focus on the personal and direct party incumbency effects, as defined above. The (overall) party incumbency effect has been investigated by Titiunik (2011). In summary, she finds that incumbent parties are affected negatively by their incumbency. She discusses a possible mechanism for this finding: the relatively weak party system in Brazilian municipalities limits parties' ability to control their candidates while they are in office. Taken together with the large resources that Brazilian mayors control, and the two term limit, mayors might act in their self-interest rather than provide the best services and policies for the municipality. Titiunik (2011) argues that these facts lead to voters expressing their dissatisfaction by punishing the incumbent party, resulting in a negative party incumbency effect. In other words, this is a punishing mechanism where voters react on past behavior of the candidate.

An alternative explanation would be that the electorate wants to avoid lame-duck mayors, as discussed in the introduction of this chapter. Voters can discipline a first-term mayor by not granting him or her a second term. Second-term mayors, on the other hand, will never run for a third term, due to the term limit, and voters lack any disciplinary power over such candidates. Mayors are therefore more likely to act in line with their self-interest in a second term compared to their first term. Voters are, thus, reluctant to grant mayors an additional terms. They do not what to vote all incumbents out of office, as the threat of not being granted a second term would then be an empty threat; they would lose all their disciplinary power. They can, however, demand that incumbent candidates are of higher quality in order to grant them a second term. In other words, this alternative explanation is a preventive mechanism where voters react on the future, expected behavior of the candidate.

While the (overall) party incumbency effect is expected to be negative under both of these mechanisms, the personal and direct party effects could differ. If voters act preventive, we would not expect the direct party effect to be negative as this refers to parties running with a candidate that would serve his or her first term (i.e., those least likely to act according to their self-interest). The personal effect with preventive voters is, on the contrary, very likely to be negative as this refers to candidates that run for their second-term.

In contrast, if voters act punitively against the party, we would suspect the personal and direct party effect to be of similar magnitude. The direct party effect could even be more negative than the personal if mayors act more in line with their self-interest in the second term. The two explanations have different implications, and our investigation could shed some light on which is more likely.[20] As we will see, the direct party effect is more negative than the personal effect consistent with the punitive mechanism discussed by Titiunik (2011).

6.1 Data

The data is obtained from the Electoral Data Repository (*Repositório de Dados Eleitorais*) maintained by the Brazilian Superior Electoral Court (*Tribunal Superior Eleitoral*). The repository contains information over candidates, parties, basic electorate demographics and election results for elections in 1994 and onwards. Including all levels of government, the repository contains nearly fifty thousand elections and more than half a million unique individuals running for office. The election data was largely collected using electronic voting machines that were widely used from the 1998 election. Subsequently,

[20]This exercise can however not rule out explanations other than the two considered here.

the data on municipal election prior to 1998 contain only a small number of municipalities and candidates, and will not be used in the analysis. The municipal elections in 2000, 2004, 2008 and 2012 result in 61,254 party-election observations.

Relative to the reference election (i.e., the election that the RDD vote margin is measured), I will use the preceding election to construct covariates and the subsequent election for outcomes. For example, for an election in 2004, the RDD vote margin refer to the 2004 election; the 2000 election provides covariates; and the 2008 election provides the outcomes. As a consequence only elections in 2004 and 2008 were included in the sample—in total 29,740 observations. A wide array of covariates describing the candidate, party and municipality was appended to these observations.[21] These covariates will be investigated in detail in subsequent sections, but, in short, they include the candidates' occupation, their election experience, if the candidate is on his or her second term as mayor, campaign contribution, district demographics, and previous party performance in the district and at higher regional levels.

In addition to the covariates, the final sample contains information on current and future election participation and performance. Of particular interest is the RDD running variable, the vote margin, which was calculated as the percentage point difference to the nearest party that would cause a change in victory status for the party. For parties that won the election, this is the difference between its vote share and the vote share of the runner-up party. For all other parties, it is the difference between its vote share and the share of the winner. For elections with two rounds, the second round was used for these calculations. This variable can potentially run between -1 (where the party lost the election, and the winning party received all the votes in the municipality) to 1 (where the party itself received all votes). However, in practice, most parties (64.7%) are positioned in the interval from -0.25 to 0.25.

The variable of whether the party's candidate re-runs in the subsequent election (R_i), was constructed by comparing the reported characteristics of the candidates in the two consecutive elections. The vast majority of candidates were matched by a unique ID number. To account for unreported and misreported IDs, the remaining candidates were matched by name and birth year.[22]

Party turn-over is high in the Brazilian setting: only 43% of parties in the sample participated in the subsequent election, and in a five percentage point vote margin window around the cut-off, this increases only to 55%. If election outcomes affect whether parties participate in the subsequent elections in a systematic way, the identifying assumptions are unlikely to hold for the same reasons that we cannot estimate the personal incumbency effect in the standard RDD. While this could be a threat to identification when investigating

[21] 167 observations, or 0.6% of the sample, had missing value on one or more of these variables and was therefore dropped from the analysis.

[22] Name matching was done using the generalized Levenshtein edit distance implemented in the `agrep` command in R.

the (overall) party incumbency effect, it will not pose any *additional* problem when investigating the already conditional versions of the incumbency effect such as the personal effect.[23] It does, however, require changes to the monotonicity assumption as discussed below.

Depending on how data-demanding the strategies are, three different vote margin windows will be used. The running-on-observables strategy requires least amount of data, and it therefore uses either a one or a two percentage point window around the cut-off; this results in 1,091 and 2,012 observations, respectively. The two other strategies are considerably more data demanding, and the window will be extended to a four percentage point window, containing 4,447 observations. These sample sizes refer to the unconditional sizes. When applying each strategy's conditioning set, the number of observations shirks considerably. As the interpretation of the RDD as local experiment is less likely to hold in a bigger window, identification with the two strategies using a four percentage point window is less credible.

6.2 Specification tests

The RDD provides a setting where the identifying assumptions are reasonably weak. Its main strength is, however, that violations of these assumptions often have observable consequences. This provides useful falsification tests of the design.

An indicative test is to study the density of observations around the cut-off (McCrary, 2008). If parties are positioned along the RDD scale in a non-continuous fashion, and especially if there are asymmetries at the cut-off, it would indicate that parties can exercise detailed control over the running variable (i.e., the vote margin). While the absence of (exact) control of their position is neither sufficient nor necessary for the RDD to be valid, it would raise suspicions if they could. For example, if some elections are subject to election fraud, and parties that cheat are not representative of the other parties, the assumptions underlying the RDD would be violated.

To investigate this, I plot the histogram over parties' vote margin around the cut-off in Figure 6. As seen in the first panel, the density is fairly uniform in vote margin windows used for estimation. While there are some density spikes close to the cut-off, they are not in the bins closest to the cut-off and not of a notable magnitude. In an ordinary RDD setting, this would indicate that the units cannot sort along the running variable. However, due to the dependence in vote shares in an election, the density will be symmetric almost by

[23] In a setting where party turn-over is high, let P_i be a binary indicator denoting whether the party's candidate re-runs. We simply alter the above analysis by exchanging R_i for P_iR_i. While this change will not change the derivations themselves, it will change the implication of the identifying assumptions; especially in the running-on-observable strategy, where we now require that the covariates are informative of both whether the parties and the candidates re-run.

Panel A: All parties.

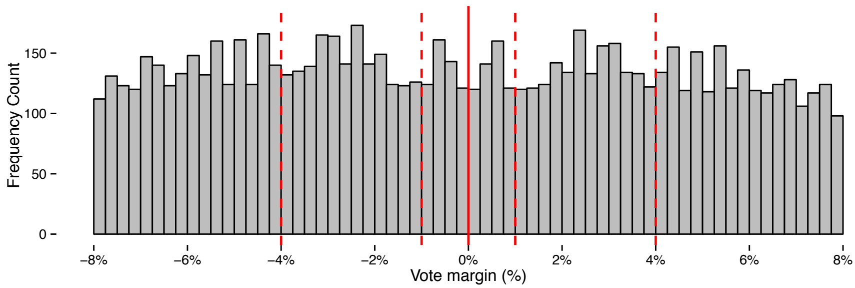

Panel B: Conditional on incumbency.

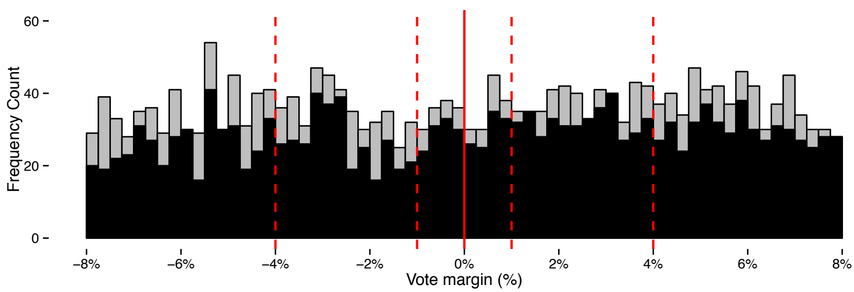

Figure 6. Histograms of the party vote margin. The first panel plots the density of the party vote margin for all parties in the sample. The second panel plots the density for incumbent parties (in gray) and parties with an incumbent mayor as their candidate (in black). In both panels, the bin width is 0.25%, the solid line at 0% indicate the RD cut-off, and the dashed lines at $\pm$ 1% and $\pm$ 4% indicate the main sample restriction used in the analysis.

construction when using vote margin as the running variable.[24] In the standard incumbency RDD (e.g., Lee, 2001, 2008), the automatic symmetry is broken by conditioning on party identity. This tests whether either party has greater ability to control the vote margin than the other. As discussed in Caughey and Sekhon (2011), it is questionable whether this is a relevant test; we expect there are other factors that are more predictive of whether the party has influence of the vote margin. Instead of breaking the symmetry by party identity, I will focus on the factor found most problematic in the past literature: incumbency status.

In the second panel of Figure 6, the density in vote margin for parties that are incumbents coming into the RDD election (i.e., if the vote margin is from 2004, these parties won the election in 2000) and parties with incumbent candidates are plotted. The density is fairly uniform in both cases. There are a worrying low density region around the -2% vote margin mark, especially for incumbent candidates. However, it is reasonably far from the cut-off—if incumbents could influence the margin, we would expect the greatest difference be just at the cut-off. In the 1% vote margin window, there is no density difference or notable discontinuity.

The interpretation of the RDD as a local random experiment implies that covariates are balanced in the neighborhood around the cut-off.[25] To investigate this, I examine the balance of the complete set of covariates, by comparing the average value in the two treatment groups. If the assumption holds, we would expect the difference between the groups to be small, and the p-values from hypothesis tests with a null of no difference to be distributed uniformly on the unit interval.[26]

Balance tests on candidate and party covariates in a 1% vote margin window are reported in Figure 7 and 8. The parties' performance in the council elections, taking place at the same time as the RDD election, is also included in Figure 8. Since voters tend to vote similarly in mayoral and council election, it is not clear whether these can be interpreted as covariates with respect to the RDD election outcome. District covariates (which are balanced by construction), and tests for the 4% window are reported in Appendix A. As expected, covariate balance is markedly worse in the larger estimation window.

Overall, the differences between treatment groups are small, and there is no systematic pattern in the p-values. Five covariates display p-values of 0.1

[24] In a two-party system, symmetry is exact by construction. In a multi-party system, symmetry holds only approximately.

[25] Strictly, the independence assumption in (10) does not require balance in the covariates, but rather balance in the potential outcome. We can, however, call the independence assumption into question if there are imbalances in covariates that are likely to be associated with the outcomes.

[26] Many of the presented covariates are correlated. This means that the informational content of the test is lower than if all covariates where independent. But, these correlations do not change the fact that we expect the p-values to be uniformly distributed.

	Mean Winners	Mean Losers
Miscellaneous		
Female	0.0867	0.0984
Age	48.7	48.9
Married	0.834	0.774
Same birth state	0.843	0.858
Same birth district	0.435	0.485
Campaign contribution		
Own funds	10804	10867
Private persons	12450	11472
Comparnies	11131	10243
Political org.	4083	3261
Other	659	667
Total	39127	36511
Electoral experience		
Incumbent mayor	0.220	0.215
Party's prev. candidate	0.299	0.302
Any election exp.	0.546	0.601
Mayoral el. exp.	0.419	0.455
Council el. exp.	0.111	0.118
Any office holding	0.334	0.339
Mayoral of. holding	0.218	0.215
Council of. holding	0.103	0.107
Ever changed party	0.179	0.233
Education		
Primary or less	0.212	0.204
Secondary	0.317	0.302
University	0.470	0.494
Occupation		
Government	0.137	0.122
Professional	0.284	0.297
White collor	0.214	0.211
Public	0.0978	0.0820
Blue collor	0.177	0.204
Other	0.0904	0.0838

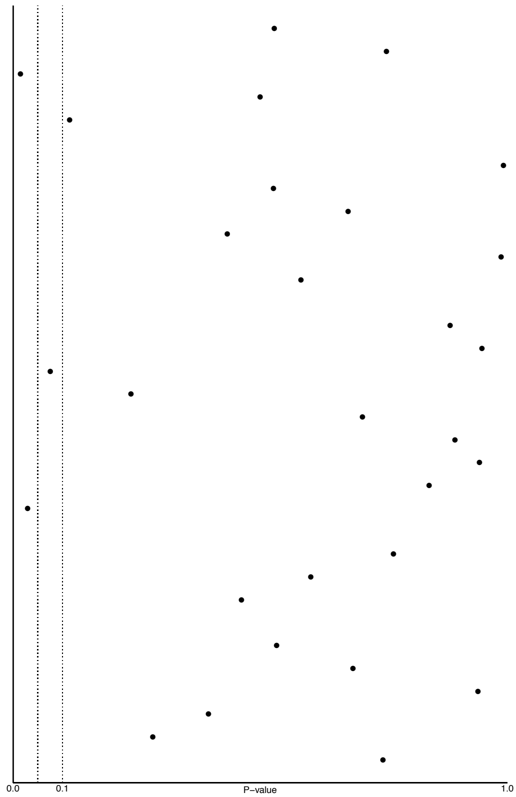

Figure 7. Balance tests for candidate covariates. Each row represents a covariate. The first two columns present the average of the covariate in the treatment and control groups. Circles indicate the p-values from two-sided Fisher's exact tests.

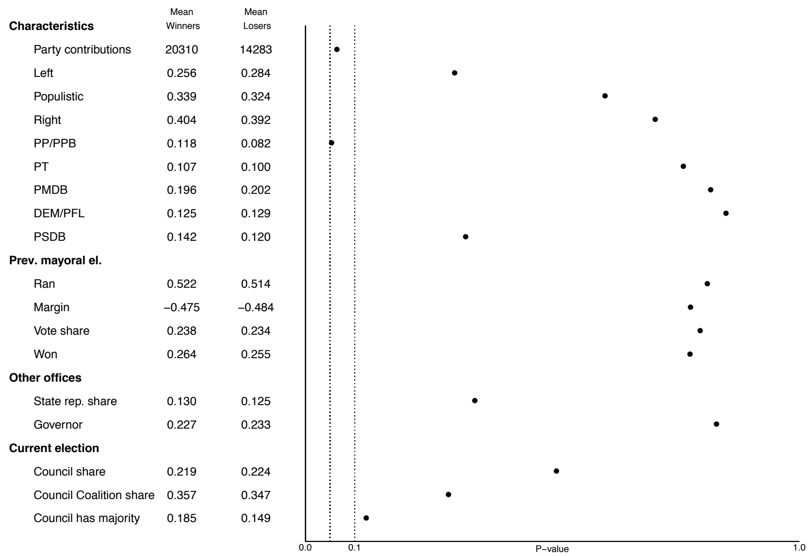

Figure 8. Balance tests for party covariates. Each row represents a covariate. The first two columns present the average of the covariate in the treatment and control groups. Circles indicate the p-values from two-sided Fisher's exact tests.

or less. Considering the large number of tested covariates, this is not unexpected. Some of these covariates (e.g., whether the candidate is married) are unlikely to be correlated with the vote margin in the population and are thus probably due to an unfortunate treatment assignment. There is, however, one covariate which is worrying: party contributions. As seen in Figure 8, close winners tend to have substantially larger contributions than close losers. If resources can be used to influence the vote margin, this is the result we would expect. However, that artifact does not show up for candidate contributions, indicating that the difference in party contribution might be coincidental. The imbalances in whether the candidate ever changed party and in election experience are noteworthy. However, in both cases, the sign of the imbalances is opposite of what would be expected *ex ante*, indicating that they do not represent systematic differences. As all three identification strategies tend to balance the covariates, small imbalances in the unconditional sample does not constitute a major threat to identification.

These balance tests are sensitive to imbalances in the complete estimation window. They could, however, mask notable imbalances in subsets of the window. The identifying assumptions imply that no imbalances occur between any parts of the window. To test this, I will use a test inspired by Caughey and Sekhon (2011). Covariate balance will be tested in disjoint 0.4% wide bins on either side of, and on equal distance from, the RDD cut-off. This is

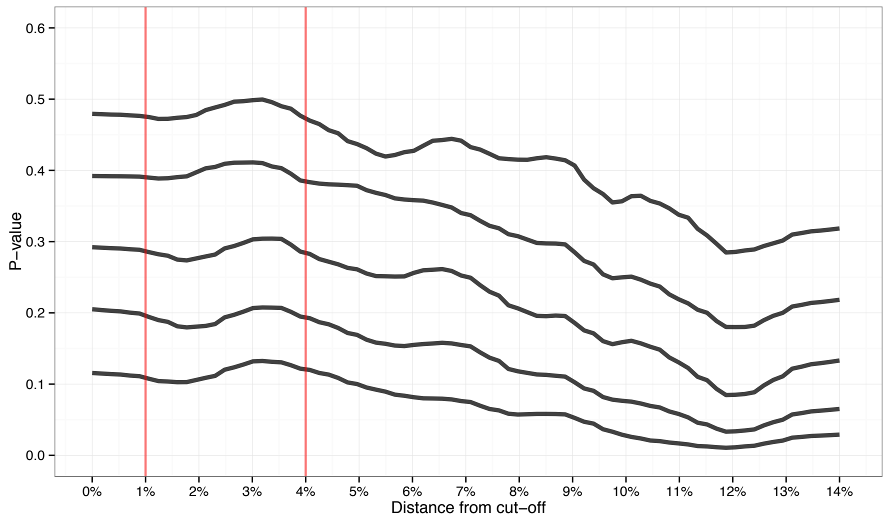

Figure 9. Balances in paired, disjoint bins at equal distance to the RD cut-off. Each line indicates one of the first five deciles of the distribution of p-value from a balance test for each covariate in 0.4% wide disjoint bins at equal distance from the cut-off. The red, vertical lines indicate the limits for the two main estimation windows at 1 and 4% vote margin.

done in 0.1% increments from 0% up to 14% vote margin. This produces a balance tests for each covariate in each bin, thus 6,204 tests in total. In each bin, the first five deciles of the tests' p-values are calculated and presented, with a smoother, in Figure 9. If the current identification strategy is valid, we expect that the smoothed decile trends are flat within the estimation window and positioned at their respective level (i.e., the first decile is at 10% and so on). As we see, this is largely the case. However, outside of the window, the p-value distribution is skewed towards zero. This indicates that the identifying assumptions are not likely to hold outside the window. Figure 20 in the appendix presents the complete distribution using a density plot, and Figure 21 presents the smoothed p-values separately for each covariate.

The last specification test is whether the vote margin is independent of the potential outcomes. If vote share is independent in the studied neighborhood, we expect the average outcome to be constant in that window except for a discontinuity at the RD cut-off. We are forced to condition this analysis on that the party participates in the subsequent election, and as discussed, this might break the causal interpretation. Nevertheless, with the current identification assumptions, we would still expect no trend in the vote margin in the estimation window. Figure 10 plots the proportion of parties that win the election after the RDD election in bins in the neighborhood around the cut-off with two different bin widths. While there is substantial noise, the proportions

Panel A: 0.2% wide bins. *Panel B:* 0.1% wide bins.

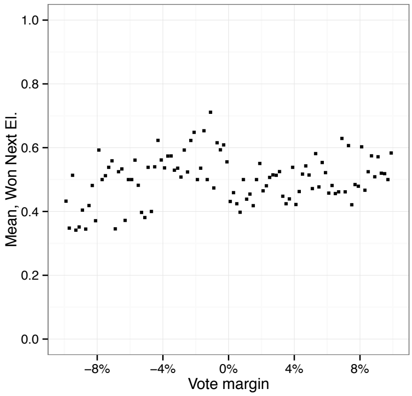

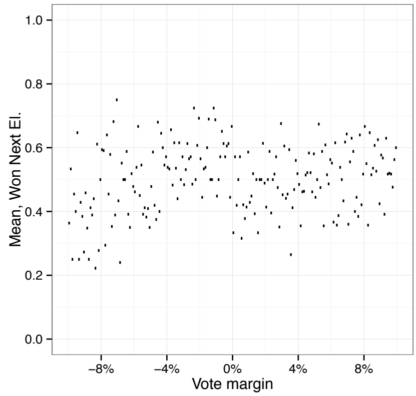

Figure 10. The overall party incumbency effect. The two panels present the proportion of parties that win the election after the RDD election in binned groups, conditional on that they run in that election. The first panel uses a 0.2% bin width, while the second panel uses a 0.1% width.

seem constant in both the 1% and 4% estimation window; there is no clear evidence of a trend.

Nearly all specification tests in this section could also be used to investigate how reasonable the three identification strategies are. For example, if conditioning on observed running status is problematic, this would show up as imbalances in these tests. There are, however, exceedingly many combinations of strategies and tests, and it is therefore not possible to present them all. But, I will present one figure that makes the issue very salient. In Figure 11, the same balance test as in Figure 7 is presented, but for two different samples.

First, the p-values for the sample when conditioning on *observed* re-running status—i.e., without regard to the unobserved potential outcome—are presented with black points. The test indicates severe imbalances in several important covariates, in particular in the candidates' prior experiences. This indicates that a naive approach, where the personal effect is investigated simply by conditioning on whether the candidate ran in the subsequent election, would be misleading. Second, the red line segments present p-values in the sample constructed with the running-on-observables strategy. No obvious systematic differences between the treatment groups seem to exist in this sample. The balance improvement is, of course, somewhat automatic due the matching; lack of severe imbalances does not provide validation that the method works. However, if not for anything else, the test indicates that the method solves the severe imbalances that otherwise would occur.

	Mean Winners	Mean Losers
Miscellaneous		
Female	0.0676	0.0857
Age	46.5	48.4
Married	0.841	0.786
Same birth state	0.850	0.895
Same birth district	0.430	0.548
Campaign contribution		
Own funds	12477	15320
Private persons	12962	11561
Comparnies	10308	11499
Political org.	4109	3986
Other	471	338
Total	40327	42703
Electoral experience		
Incumbent mayor	0.000	0.181
Party's prev. candidate	0.208	0.314
Any election exp.	0.473	0.610
Mayoral el. exp.	0.295	0.462
Council el. exp.	0.159	0.110
Any office holding	0.164	0.300
Mayoral of. holding	0.000	0.176
Council of. holding	0.1498	0.0952
Ever changed party	0.164	0.219
Education		
Primary or less	0.184	0.200
Secondary	0.319	0.262
University	0.498	0.538
Occupation		
Government	0.0483	0.1000
Professional	0.353	0.276
White collor	0.246	0.210
Public	0.0966	0.1095
Blue collor	0.198	0.238
Other	0.0580	0.0667

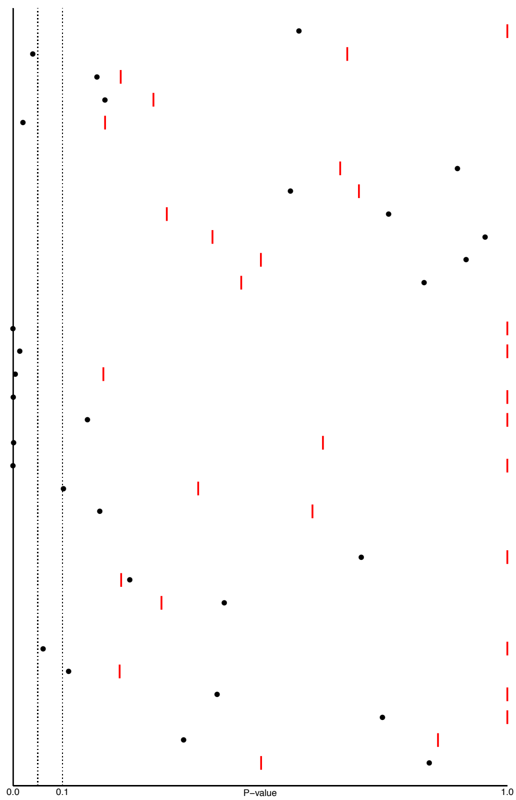

Figure 11. Balance tests for candidate covariates in conditional samples. Each row represents a covariate. The first two columns present the average of the covariate in the treatment and control groups in the subsample of parties with re-running candidates. Circles indicate the p-values from two-sided Fisher's exact tests in this sample. Red line segments indicate the p-values from a paired two-sided Fisher's exact test in the sample constructed by the running-on-observables identification strategy.

6.3 Monotonicity

Two of the identification strategies depend on a monotonicity assumption. The term limit and high party turn-over that exist in the Brazilian setting complicate this assumption considerably. Consider, first, the term limit. As in an electoral setting without a term-limit, we expect candidates running for their first-term to be more likely to run in the subsequent election if they win, i.e., $R_i(1) \geq R_i(0)$; the term-limit does not apply to these candidates. For candidates that are incumbent mayors, this is no longer the case. When the term limit is reached, these candidates are not allowed to take office for another term. The directionality of the monotonicity, thus, depends on whether the candidate runs for a first or a second term.

The effect of the term limit can clearly be seen in Figure 12, where the proportion of re-runners is plotted in bins around the cut-off separately for incumbent mayors (running for their second term) and first-time runners. Among incumbents, hardly any of the winners run in the subsequent elections, as we would expect from the term-limit.[27] First-time runners, on the other hand, seem to run for office to a higher degree when winning, again, as we would expect.

Panel A: Incumbent mayors. *Panel B:* Non-incumbents.

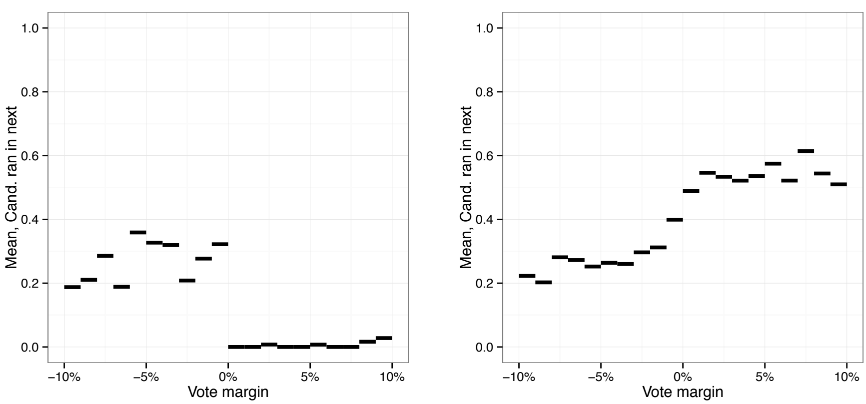

Figure 12. The causal effect of W_i on R_i. The two panels show the propensity of parties' candidates to re-run for office in the election following the RDD election in 1% wide bins around the cut-off. The leftmost panel does this for candidates that are incumbent mayors coming into the RDD election, while the rightmost does the same for non-incumbents.

[27] There are a few winning incumbent mayors that run for a third term (in total 12, or 0.4%), seemingly contrary to the elections rules. There are three possible explanations for this. First, there could be a matching error where two different candidates erroneously been given the same ID number. Second, a candidate could possibly *run* for office even if he or she was prohibited to *take* office. Third, there could be, to me unknown, exceptions made to this rule.

Now, consider the high party turn-over. In Brazil's unstable party system, the election outcome can affect not only whether the candidate runs in the subsequent election, but also whether the party runs. Monotonicity must then be extended to include party participation. For first-time runners, this is arguably unproblematic: winning an election is likely to increase the likelihood of running for both candidates and parties. However, for incumbent mayors, this is not the case. Due to the term limit, winning the election surely lowers the probability that the candidate re-runs. The term-limit does, however, not limit the parties to run with another candidate; it is conceivable that winning an election sometimes increases the likelihood that the party runs in the subsequent election even if its candidate cannot run.

Let P_i indicate whether party i participate in the subsequent election, and as before, let R_i indicate whether the candidate does. The monotonicity assumption will differ depending on whether the term-limit is binding. First, consider first-time runners. The term-limit does not apply to these candidates, and as discussed above, winning an election will, if anything, cause both the candidate and the party to run in next election: $P_i(1) \geq P_i(0)$ and $R_i(1) \geq R_i(0)$. This implies that the "joint" monotonicity assumption is $R_i(1)P_i(1) \geq R_i(0)P_i(0)$.

Now, consider parties with incumbent candidates. Due to the term limit, these candidates will have $R_i(0) \geq R_i(1) = 0$. The relationship between $P_i(1)$ and $P_i(0)$ is less clear. We can imagine that there are two types of parties. On the one hand, as the party system is weak in Brazil, it is not uncommon that parties depend greatly on specific candidates; this type of party will participate only if their candidates do. For these parties, when winning the election, the candidate cannot run in the next election, and as a result, the party does not run either: we have $P_i(0) \geq P_i(1) = 0$. On the other hand, there are surely some parties that are established and do not depend on whether their candidates run. This second type of parties might still decide not to re-run, but then for other reasons. For these parties, we expect that winning only can cause the party to run: $P_i(0) \leq P_i(1)$. While we could try to identify the two different types of parties, it is not needed for identification of the effects. In particular, all parties with incumbent candidates will, under the current assumptions, satisfy $P_i(1) \geq P_i(0)[1 - R_i(0)]$ and $R_i(0) \geq R_i(1) = 0$ which is sufficient for identification. The monotonicity assumptions for parties with first-time runners and incumbents are illustrated in Figure 13.

The running-on-observable strategy requires us to identify units that are always- or never-runners. In a stable party system without a term-limit—i.e., where $R_i(1) \geq R_i(0)$ holds for all parties—we know that all losing parties with re-running candidates (i.e., $W_i = 0$ and $R_i = 1$) are always-runner. In the Brazilian setting, this is not the case; the alternative monotonicity assumptions will complicate the analysis. We can, however, identify always- and never-runners in the subsamples depending on incumbency status.

When accounting for party participation, always-runners are defined as $R_i(1) = R_i(0) = P_i(1) = P_i(0) = 1$. Consider parties with non-incumbent

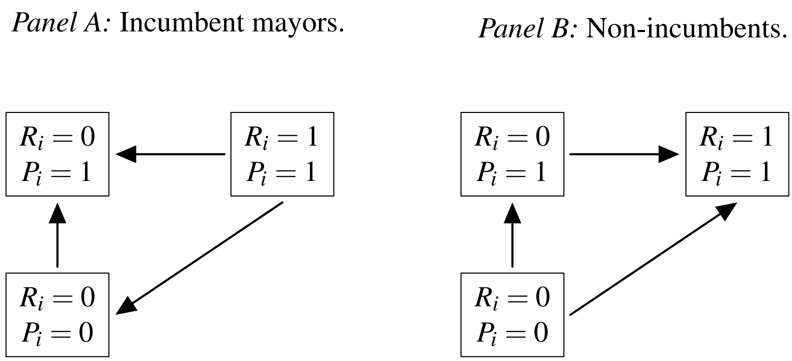

Figure 13. The monotonicity assumption in each subsample. Each box represented a set of observed re-running statuses for the party and candidate. The arrows indicate the assumed unidirectional flows caused by winning an election. For example, the leftmost arrow in the first panel indicate that, in this subsample, all parties with $P_i(0) = R_i(0) = 0$ have $R_i(1) = 0$ and $P_i(1) \geq 0$.

candidates with $W_i = 0$, $P_i = 1$ and $R_i = 1$. That is, parties that lost the preceding election, but where both the party and the candidate participate in the next election. The revealed potential outcomes for these parties are $R_i(0) = P_i(0) = 1$. As the parties had non-incumbent candidates, the relevant monotonicity assumption is $R_i(1)P_i(1) \geq R_i(0)P_i(0)$. Thus, for these parties, we have $R_i(1)P_i(1) = R_i(0)P_i(0) = 1$, which implies that they are always-runners.

Never-runners, in this context, are parties that run in the next election independently of the election result, but their candidates do not; that is, they have $P_i(1) = P_i(0) = 1$ and $R_i(1) = R_i(0) = 0$. Consider parties with incumbent candidates that lost the election, run in the next election but their candidates do not: $W_i = 0$, $P_i = 1$ and $R_i = 0$. The revealed potential outcomes for these parties are $R_i(0) = 0$ and $P_i(0) = 1$. As they have incumbent candidates, the relevant monotonicity assumption is $P_i(1) \geq P_i(0)[1 - R_i(0)]$. This implies that $P_i(1) = P_i(0)[1 - R_i(0)] = 1$. Due to the term limit, we know that these parties have $R_i(1) = 0$. That is, these parties are never-runners.

Monotonicity is always a strong assumption, and the modified version especially so. To use it, we must ensure that it applies to the studied elections. It, for example, precludes that winning elections cause candidates to be able to run for higher-ranking offices in the subsequent elections. While such behavior is more common in local and regional elections, most candidates in the Brazilian setting tend to be at the end rather than in the beginning for their careers, or not be professional politicians at all (Titiunik, 2011). While this increases our confidence in the assumption, we must keep in mind that the results depend on it.

6.4 Personal incumbency effect

Turning to the results, I first present the estimates of the personal incumbency effect in Table 1. The first panel contains the effect on the propensity to win the election following the RDD election for the three strategies, and the second panel contains the effect on the vote share in the subsequent election.

Starting with always-runner stratification strategy, in a 4% estimation window, every potential always-runner (parties with a first-time running candidate with observed $R_i = 1$) is matched to its three ($k = 3$) closest neighbors, in both treatment and control, based on their covariate distances. If all six parties also have candidates that ran in the subsequent elections, the party is included in the studied sample. This produces a sample of 39 units out of the 1,452 parties that had non-incumbent candidates that re-ran.[28] The point estimates indicate a slight negative effect on the propensity to win, and essentially no effect on vote share. In both cases, the null hypothesis of no effect is not rejected. The always-runner stratification estimate is considerably higher than the estimates with the other two strategies. This difference could be due both to the estimands' localness—the strategies could simply refer to different effects—or the high degree of uncertainty with the current estimate.

Table 1. *Personal incumbency effects.*

Strategy	Mean losers	M. winners	Effect	P-value	Obs.
Panel A: Victory propensity					
AWS	0.667	0.600	-0.0667	0.7420	39
NCS	0.625	0.463	-0.1620	0.1233	110
ROO	0.645	0.512	-0.1337	0.0163	344
Panel B: Vote share					
AWS	0.501	0.502	0.000487	0.985	39
NCS	0.472	0.455	-0.016224	0.601	110
ROO	0.490	0.487	-0.003016	0.818	344

Note: The two panels present the estimates of the personal incumbency effect for two outcomes. Each row represent a different identification strategy where AWS indicates always-runner stratification, NCS indicates non-complier stratification and ROO the running-on-observables strategy.

Continuing with non-complier stratification, the match tolerance is set to $\varepsilon = 0.05$, which produces a sample of 110 observations. The estimated effect on victory propensity now decreases to -16.2 percentage points with a p-value

[28] If the monotonicity assumption holds, we can estimate the number of always-runners by doubling the number of re-runners among losing parties, which indicates that the total is 1,148. Always-runner stratification is thus a very local effect in this case.

just shy of the 0.1 mark. The estimated effect on vote share remains close to zero, and it would not be a remarkable observation under the null.

Last, the running-on-observables strategy allows us to estimate the effect for all non-incumbent always-runners. The 172 parties in the 1% vote margin window with non-incumbent candidates that re-ran for office despite losing the election (which under monotonicity are always-runners) are matched to their closest neighbor among winning parties with re-running candidates. This yields a sample of 344 observations. The personal incumbency effect is estimated by the difference in outcomes between these two groups. The effect on propensity to win the subsequent election is -13.4 percentage points. The null hypothesis is rejected, with a p-value of 0.016. The effect on vote share is, however, close to zero also with this strategy.

Panel A: Victory propensity. *Panel B:* Vote share.

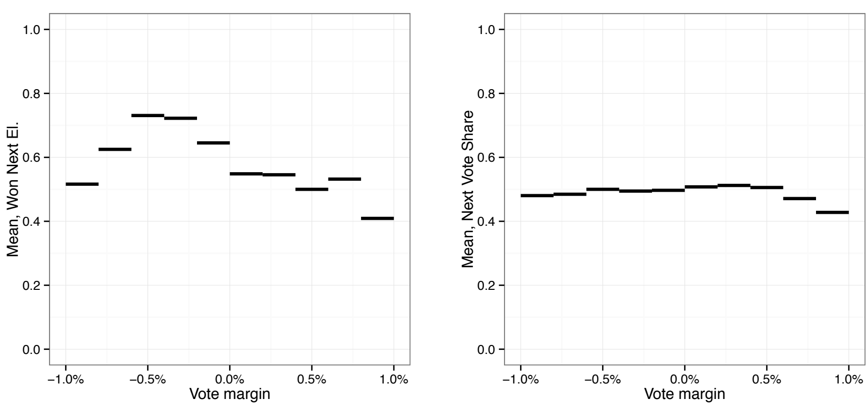

Figure 14. Personal incumbency effect with the running-on-observables strategy. The two panels show the average outcome in 0.2% wide bins around the cut-off. The leftmost panel presents the propensity to win the election following the RDD election, and the rightmost panel the average vote share in that election.

The negative effect on victory propensity might seem puzzling considering the absence of an effect on the vote share. The results are, however, consistent with an explanation where the electorate gains additional information about candidates when they win elections. In that situation, desired candidates would (credibly) reveal their type to the electorate and thereby enjoy increased vote margins when they win. Undesirable candidates can no longer hide their type and suffer decreased vote margins. The two effects can offset each other when vote shares are the outcome, leading to an average effect close to zero. It will, however, increase the spread of the vote share distribution. If this distribution is centered above the RDD cut-off (which is likely since all parties won the preceding election), there will be a negative effect on victory propensity. In particular, an increase in vote margin for desirable candidates would not increase their propensity to win as much as a decrease for undesirable candidates would increase their propensity to lose. Under this explanation, the negative personal incumbency effect is mainly driven by that undesirable candidates being voted out of office. The results can, however, not rule out alternative explanations.

Neither of these methods allow for good plots of the average outcome in bins around the cut-off, as in the usual RDD. This is partly because the plots are restricted to the estimation window, due to the matching, and partly because of the low sample sizes. Despite these caveats, in Figure 14, the average outcome in the running-on-observable sample is plotted using 0.2% wide bins. Note the difference in scale compared to previous graphs; the bin-width only is one fifth of, for example, Figure 12. This explains the increased bin variability.

6.5 Direct party incumbency effect

With the direct party incumbency effect, we must identify a set of parties with candidates that do not run in the subsequent election: never-runners. Due to the high party turn-over, we must also ensure that the parties participate in the next election. In other words, we are looking for parties with $R_i(0) = R_i(1) = 0$ and $P_i(0) = P_i(1) = 1$.

Table 2. *Direct party incumbency effects.*

Strategy	Mean losers	M. winners	Effect	P-value	Obs.
Panel A: Victory propensity					
NCS	0.333	0.389	0.0556	0.7817	57
ROO	0.396	0.188	-0.2083	0.0305	96
Panel B: Vote share					
NCS	0.394	0.448	0.0541	0.1882	57
ROO	0.425	0.386	-0.0388	0.0678	96

Note: The two panels present the estimates of the personal incumbency effect for two outcomes. Each row represent a different identification strategy where AWS indicates always-runner stratification, NCS indicates non-complier stratification and ROO the running-on-observables strategy.

With the always-runner stratification strategy (or, in this case, never-runner stratification), the monotonicity assumption is not needed, and we can include all potential never-runners. Like when estimating the personal incumbency effect, we match all potential never-runners (parties with $R_i = 0$ and $P_i = 1$) with their three closest neighbors in both treatment and control. Again, if all matches have $R_i = 0$ and $P_i = 1$, we include the unit in the estimation sample.

Among 869 potential never-runners within the 4% estimation window, only a single unit is estimated to be an actual never-runner. While we could increase the estimation window, or lower k, to increase the number of estimated never-runners, such changes would not lead to credible results; current choices already push the limit. Instead, I will forgo any attempt to estimate the direct party effect with this strategy.

Panel A: Victory propensity. *Panel B:* Vote share.

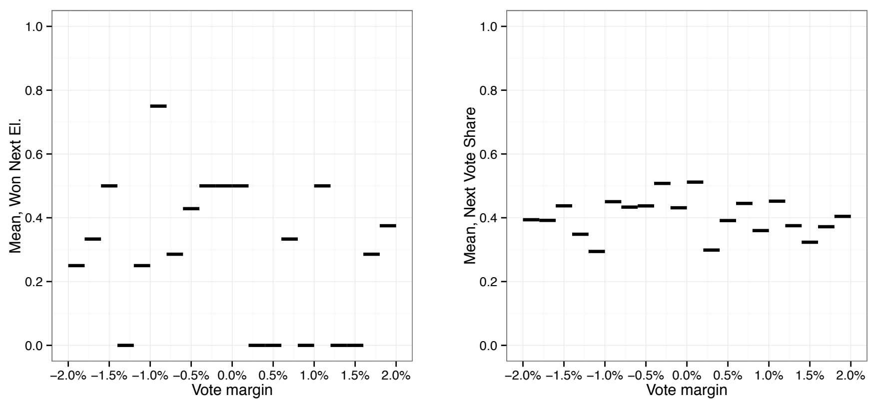

Figure 15. Direct party incumbency effect with the running-on-observables strategy. The two panels show the average outcome in 0.2% wide bins around the cut-off. The leftmost panel presents the propensity to win the election following the RDD election, and the rightmost panel the average vote share in that election.

With non-complier stratification, the monotonicity assumption is needed, thus we restrict our attention to parties with incumbent candidates as discussed in Section 6.3. With a tolerance at $\varepsilon = 0.05$, this results in a sample of 57 parties. The estimates from this sample are presented in the first rows in Table 2. Contrary to the personal effect, the estimates are here positive for both the propensity to win and vote share. However, in neither case can the null hypothesis of no effect be rejected.

For the running-on-observables strategy, due to the sparsity of observations with this conditioning set, I extend the estimation window to 2% compared to the 1% window used for the personal effect. This leads to 48 parties that are revealed never-runners under monotonicity (i.e., losing parties with incumbent candidates where the party ran in the subsequent election, but the candidate did not). These are matched with winning parties with incumbent candidates, producing a sample of 96 parties. The estimated effect is now negative and strongly so, with a 20.8 percentage point decrease in the propensity to win, and 3.9 percentage point decrease in vote share. In both cases, the hypothesis tests reject the sharp null of no effects.

The results are plotted in Figure 15. However, due to the very small sample sizes (each bin contain on average only 4.8 observations), the bins exhibit a lot of variation.

Comparing the direct party and personal effects, we see that the direct party effect is lower than the personal effect. Going back to the two discussed explanations for the negative overall party effect, this indicate that the punitive, rather than the preventive, mechanism are more consistent with these results—in line with the discussion in Titiunik (2011). This conclusion, however, rest on an assumption that the effects are the same in both the studied sub-populations. This is a strong assumption which, in general, does not hold. While providing some indication that the punitive mechanism might be more relevant, the illustration in this section does not provide enough support for any definite conclusions.

7 Concluding remarks

In this chapter, I have presented a causal model with which several previously discussed incumbency effects can be defined. The model assumes manipulation of both whether the party won the preceding election, and whether the party's candidate re-runs for office. Holding one of these variables constant, while varying the other, yields the definitions of four different effects. One of these effects, the legislator incumbency effect, corresponds exactly to a past definition by Gelman and King (1990). Two of the effects, the personal and direct party incumbency effects, are not new concepts but have, to my knowledge, never been formally defined.

The definitions allow us to investigate how previous methods in the literature are related. This reveals that the party incumbency effect, as investigated with the standard RDD strategy, can be decomposed into the effects defined in this study. While the prospects of estimating these parts directly are slim, the decomposition helps us interpret the effect and could act as a basis for an explanation of why the party effect differs between electoral settings. A similar exercise was conducted for other methods used in the previous literature, and it reveals that these methods mainly focus on the legislator effect.

Motivated by the lack of prior investigations of the personal and direct party effect, three identification strategies of these effects were discussed. Using various assumptions, the effects are shown to be identified for in three sub-populations of varying sizes. The usefulness of the strategies, both in terms of the severeness of the assumptions and the localness of the estimands, are highly dependent on the specifics of the election setting.

To illustrate these strategies, the incumbency effects were investigated in the setting of Brazilian mayoral elections. This analysis highlights several difficulties when investigating the effect in an unstable party system with a term-limit, as is the case in Brazil. The results indicate that both the personal and direct party effects are negative. The effects are consistent with an explanation where the electorate punishes parties for previously bad performance, but are less consistent with an explanation where the electorate has preferences against second-term (lame-duck) mayors and therefore disfavor candidates seeking reelection.

References

Ansolabehere, Stephen, James M. Snyder Jr., and Charles Stewart (2000) "Old voters, new voters, and the personal vote: Using redistricting to measure the incumbency advantage," *American Journal of Political Science*, Vol. 44, No. 1, pp. 17–34.

Cattaneo, Matias D., Brigham Frandsen, and Rocío Titiunik (2013) "Randomization inference in the regression discontinuity design: An application to party advantages in the U.S. Senate."

Caughey, Devin and Jasjeet S. Sekhon (2011) "Elections and the regression discontinuity design: Lessons from close us house races, 1942–2008," *Political Analysis*, Vol. 19, No. 4, pp. 385–408.

Cox, Gary W. and Jonathan N. Katz (1996) "Why did the incumbency advantage in US House elections grow?," *American Journal of Political Science*, Vol. 40, No. 2, pp. 478–497.

——— (2002) *Elbridge Gerry's salamander: The electoral consequences of the reapportionment revolution*: Cambridge University Press.

Cummings, Milton C. Jr. (1966) *Congressmen and the electorate*: The Free Press.

Diamond, Alexis and Jasjeet S. Sekhon (2012) "Genetic matching for estimating causal effects: A general multivariate matching method for achieving balance in observational studies," *Review of Economics and Statistics*, Vol. 95, No. 3, pp. 932–945.

Erikson, Robert S. (1971) "The advantage of incumbency in congressional elections," *Polity*, Vol. 3, No. 3, pp. 395–405.

Erikson, Robert S. and Rocío Titiunik (2013) "Using regression discontinuity to uncover the personal incumbency advantage," Unpublished manuscript.

Ferraz, Claudio and Frederico Finan (2011) "Electoral accountability and corruption: Evidence from the audits of local governments.," *American Economic Review*, Vol. 101, pp. 1274–1311.

Frangakis, Constantine E. and Donald B. Rubin (2002) "Principal stratification in causal inference," *Biometrics*, Vol. 58, No. 1, pp. 21–29.

Gelman, Andrew and Gary King (1990) "Estimating incumbency advantage without bias," *American Journal of Political Science*, Vol. 34, No. 4, pp. 1142–1164.

Hahn, Jinyong, Petra Todd, and Wilbert van der Klaauw (2001) "Identification and estimation of treatment effects with a regression-discontinuity design," *Econometrica*, Vol. 69, No. 1, pp. 201–209.

Imbens, Guido W. and Joshua D. Angrist (1994) "Identification and estimation of local average treatment effects," *Econometrica*, Vol. 62, No. 2, pp. 467–475.

Lee, David S. (2001) "The electoral advantage to incumbency and voters' valuation of politicians' experience: A regression discontinuity analysis of elections to the U.S. House,"Technical report, NBER Working Paper 8441.

——— (2008) "Randomized experiments from non-random selection in U.S. House elections," *Journal of Econometrics*, Vol. 142, No. 2, pp. 675–697.

Levitt, Steven D. and Catherine D. Wolfram (1997) “Decomposing the sources of incumbency advantage in the US House,” *Legislative Studies Quarterly*, Vol. 22, No. 1, pp. 45–60.

McCrary, Justin (2008) “Manipulation of the running variable in the regression discontinuity design: A density test,” *Journal of Econometrics*, Vol. 142, No. 2, pp. 698–714.

Rubin, Donald B. (1974) “Estimating causal effects of treatments in randomized and nonrandomized studies.,” *Journal of Educational Psychology*, Vol. 66, No. 5, pp. 688–701.

Splawa-Neyman, Jerzy, Dorota M. Dabrowska, and Terence P. Speed (1923/1990) “On the application of probability theory to agricultural experiments. Essay on principles. Section 9,” *Statistical Science*, Vol. 5, No. 4, pp. 465–472.

Titiunik, Rocío (2011) “Incumbency advantage in brazil: Evidence from municipal mayor elections.”

Uppal, Yogesh (2009) “The disadvantaged incumbents: estimating incumbency effects in Indian state legislatures.,” *Public Choice*, Vol. 138, pp. 9–27.

VanderWeele, Tyler and Stijn Vansteelandt (2009) “Conceptual issues concerning mediation, interventions and composition,” *Statistics and Its Interface*, Vol. 2, pp. 457–468.

Appendices

A Additional graphs

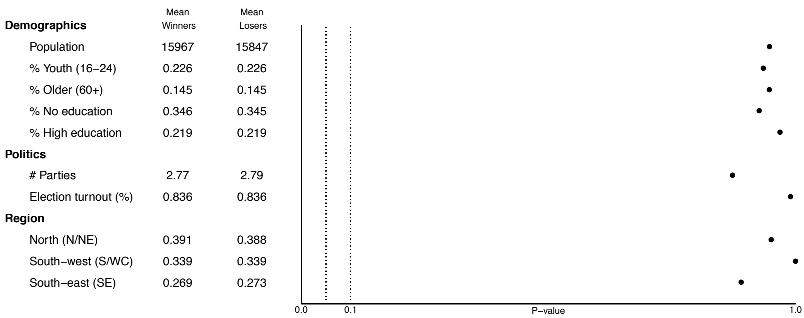

Figure 16. Balance tests for district covariates in the 1% estimation window. Each row represents a covariate. The first two columns present the average of the covariate in the treatment and control groups. Circles indicate the p-values from two-sided Fisher's exact tests.

Miscellaneous	Mean Winners	Mean Losers
Female	0.085	0.103
Age	48.2	48.7
Married	0.799	0.779
Same birth state	0.850	0.846
Same birth district	0.422	0.436
Campaign contribution		
Own funds	10356	10628
Private persons	14243	11608
Comparnies	15897	12388
Political org.	6633	4699
Other	600	512
Total	47730	39835
Electoral experience		
Incumbent mayor	0.232	0.195
Party's prev. candidate	0.296	0.289
Any election exp.	0.581	0.555
Mayoral el. exp.	0.445	0.417
Council el. exp.	0.107	0.114
Any office holding	0.349	0.310
Mayoral of. holding	0.232	0.195
Council of. holding	0.0996	0.1018
Ever changed party	0.218	0.199
Education		
Primary or less	0.198	0.212
Secondary	0.306	0.306
University	0.496	0.482
Occupation		
Government	0.135	0.123
Professional	0.302	0.302
White collor	0.195	0.205
Public	0.0882	0.0894
Blue collor	0.193	0.182
Other	0.0873	0.0983

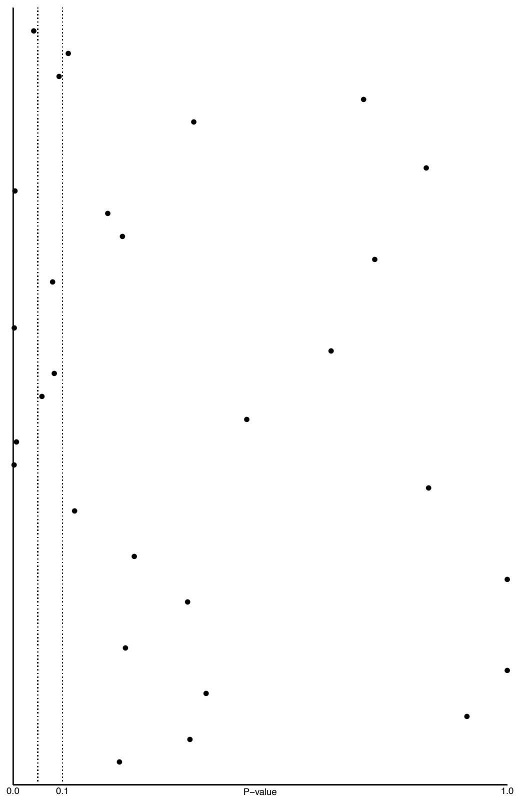

Figure 17. Balance tests for candidate covariates in the 4% estimation window. Each row represents a covariate. The first two columns present the average of the covariate in the treatment and control groups. Circles indicate the p-values from two-sided Fisher's exact tests.

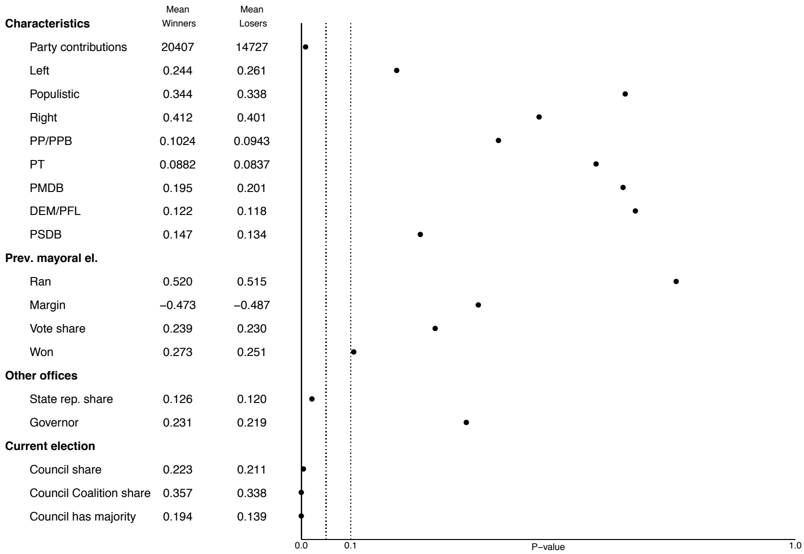

Figure 18. Balance tests for party covariates in the 4% estimation window. Each row represents a covariate. The first two columns present the average of the covariate in the treatment and control groups. Circles indicate the p-values from two-sided Fisher's exact tests.

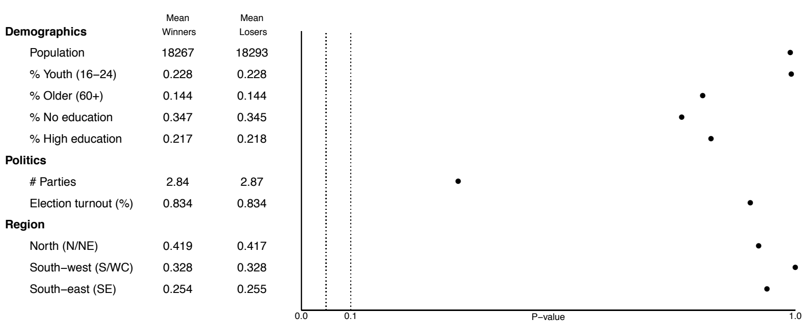

Figure 19. Balance tests for district covariates in the 4% estimation window. Each row represents a covariate. The first two columns present the average of the covariate in the treatment and control groups. Circles indicate the p-values from two-sided Fisher's exact tests.

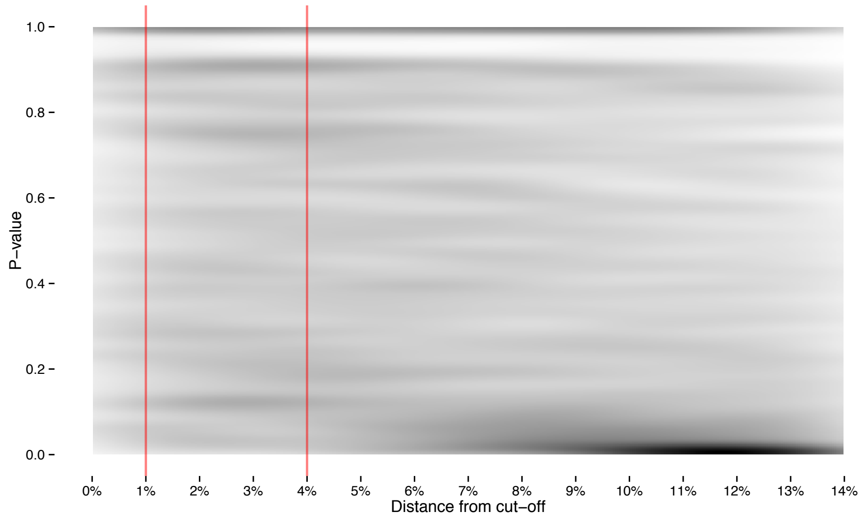

Figure 20. Density of p-value from balance test in bins around the RDD cut-off. The graph plots the density of the p-value from balance tests for each of 44 covariates in 0.4% wide disjoint bins at equal distance from the cut-off. Darker areas indicate more densely populated regions.

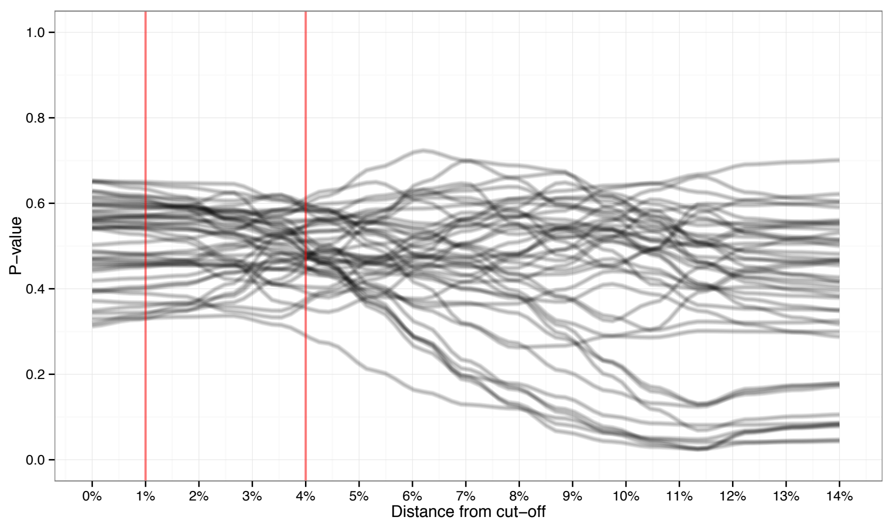

Figure 21. Balances in bins at equal distance to the RDD cut-off for separate covariates. Each line represents the smoothed p-value of one of the 44 covariates from a balance test in 0.4% wide disjoint bins at equal distance from the cut-off. The red, vertical, lines indicate the limits for the two main estimation windows at 1 and 4% vote margin.