

Using Regression Discontinuity to Uncover the Personal Incumbency Advantage

Robert S. Erikson¹ and Rocío Titiunik^{2*}

¹*Department of Political Science, Columbia University, 420 W. 118th Street, New York, NY 10027, USA; rse14@columbia.edu*

²*Department of Political Science, University of Michigan, 5700 Haven Hall, 505 South State St, Ann Arbor, MI 48109, USA; titiunik@umich.edu*

ABSTRACT

We study the conditions under which a regression discontinuity (RD) design can be used to recover the personal incumbency advantage, a quantity that has long been of interest to political scientists. We offer an expanded interpretation of the RD design that allows us to back out unbiased estimates of this quantity by focusing on open seats — elections with no incumbent running. Our focus on open seats avoids including in the personal incumbency advantage estimate the spurious advantage that stems from incumbents' higher than average quality — a result of electoral selection. A central result of our model is that the RD design double-counts the personal incumbency advantage because each of the two groups of districts compared in the

*We thank the Editors — Joshua Clinton, Keith Krehbiel and Nolan McCarty — two anonymous reviewers, John Bullock, Matias Cattaneo, Sanford Gordon, Justin Grimmer, Shigeo Hirano, Gregory Huber, Gary King, Kelly Rader, Eric Schickler,

Online Appendix available from:

http://dx.doi.org/10.1561/100.00013137_app

Supplementary Material available from:

http://dx.doi.org/10.1561/100.00013137_supp

MS submitted 26 November 2013; final version received 13 November 2014

ISSN 1554-0626; DOI 10.1561/100.00013137

© 2015 R. S. Erikson and R. Titiunik

RD estimand has an incumbent running for reelection at the time the outcome is measured. We provide a brief empirical illustration of our model that analyzes northern open-seat U.S. House elections between 1968 and 2008, and we also discuss how this setup can be used to study incumbent races, where the required assumptions are more complex. A version of our model in its full generality — and a discussion of the assumptions it requires — is presented in the Supplemental Appendix.

The dramatic increase in U.S. House incumbents' vote margins beginning in the 1960s sparked a large literature on the phenomenon of the incumbency advantage. From its beginnings, this literature has conceptualized the incumbency advantage as a personal electoral advantage that accrues to the incumbent candidate.¹ The personal incumbency advantage can be defined as the votes a candidate gains upon becoming an incumbent due to the direct and indirect benefits of office-holding. The direct benefits include the electoral gains that stem from constituency service, name recognition, and the like.² The indirect benefits arise from the "scareoff" advantage (Cox and Katz, 1996), whereby office-holders' incumbency status deters potentially strong opponents from entering the race.

As we discuss below, there are multiple methodological obstacles to obtaining a valid estimate of the incumbency advantage. Lee (2008)

Jasjeet Sekhon, Robert Van Houweling, Jonathan Wand, seminar participants at Harvard University, New York University, Princeton University, Stanford University, the University of California at Berkeley, and Yale University, and conference participants at the 2011 MIT American Politics Conference, for helpful comments and suggestions. Titiunik gratefully acknowledges financial support from the National Science Foundation (SES 1357561).

¹See Erikson (2015) for a recent and comprehensive review of the incumbency advantage literature. See also Ansolabehere and Snyder Jr. (2002) and references therein.

²In the language of potential outcomes, the personal incumbency advantage can be understood as the difference between the incumbent's observed vote share and the vote share that a non-incumbent candidate of identical quality and party would have received in the same election.

recently proposed using a regression discontinuity (RD) design to address some of these difficulties. This design is built on the assumption that, in closely contested elections, the winning candidate is essentially determined by chance. For example, when this assumption holds, U.S. House districts where the Democratic candidate barely wins can be assumed to be identical to districts where the Republican candidate barely wins, approximating random assignment of the election outcome, and making it possible to recover the effect of the party of the winner — incumbency — upon the partisan vote in the next election. Political scientists have been quick to add the RD design to the toolbox of incumbency advantage methods,³ but have not paid much attention to how the quantity that is estimated by the RD design is related to the personal incumbency advantage — the quantity that has long been of interest to substantive scholars. The goal of our article is precisely to clarify this issue. We present a model under which the personal incumbency advantage can be directly recovered from the RD estimand. In discussing the model’s assumptions, we clarify the conceptual differences between the personal incumbency advantage and the RD estimand, and also between the various components of the personal incumbency advantage. A central result of our model is that the RD estimand double-counts the personal incumbency advantage — an issue that had not been previously noted.

There are three phenomena that may lead to a positive correlation between incumbency status and electoral success but are *not* part of the personal incumbency advantage. First, the personal incumbency advantage should not be confused with the incumbent’s likely “quality advantage” that stems from the simple fact that, everything else being equal, high quality candidates tend not only to win elections but to win repeatedly (Levitt and Wolfram, 1997; Zaller, 1998). Second, this quality advantage may scare off strong potential opponents, a quality-induced deterrence effect that does not stem from incumbency status and thus should not be counted as part of the personal incumbency advantage. Finally, the personal incumbency advantage is conceptually different from any partisan incumbency advantage that may arise from being the candidate of the incumbent office-holding party in the district. For

³See, e.g., Broockman (2009), Butler (2009), Golden and Picci (2015), and Trounstine (2011).

instance, the incumbent party could build an organizational structure that continues to benefit future candidates independently of these candidates' personal incumbency status.

The as-if randomness of the election outcome in close races makes the RD design an attractive research strategy to circumvent some of these issues. But what exactly this strategy estimates is not fully understood. At its most general level, Lee's implementation of the RD design captures "the overall causal impact of being the current incumbent party in a district on the votes obtained in the district's election" (Lee, 2008, p. 682) — that is, the electoral gain to the incumbent party in a district relative to the vote share that the party would have obtained if it had not won the previous election. Naturally, if the incumbency advantage parameter is defined in this way, the RD design is appropriate and possibly the best available research design to estimate it. But this definition, with its emphasis on a party-level analysis, explicitly ignores (i) whether the candidate who barely wins at election t runs for reelection in the following election at $t + 1$ (the moment when outcomes are measured) and (ii) whether the previously elected incumbent is contesting the close election at t . As we show below, these two issues are crucial to provide a careful interpretation of the incumbency advantage quantity that is estimated with an RD design.

Our goal is to develop and clarify assumptions under which the RD estimand can recover the personal incumbency advantage. With this purpose, we develop a simple and general model that links the quantity estimated by the RD design to the direct and indirect components of the personal incumbency advantage. Our model makes clear that recovering the personal incumbency advantage from the RD estimand in full generality requires strong assumptions that are likely implausible in many electoral contexts. For this reason, we choose to sacrifice generalizability and focus only on open seat races, a strategy that allows us to develop our model under plausible assumptions.

The assumptions necessary for an RD-based estimate of the personal incumbency advantage are plausible for open seat races, but generally not appropriate for incumbent races. When open seat elections are extremely close, each party's winners and losers are, on average, of equal quality. When winners of close open seat contests seek reelection, they (on average) retain their initial quality. The vote in the following election thus reflects two changes — the officeholding premium from incumbency

status (the direct personal incumbency advantage), plus any decline in challenger quality that arises from the incumbent's deterring ability (the indirect scareoff effect). At the RD cutoff, the arbitrary winner has no net advantage due to greater candidate quality (vote appeal) or quality-induced scareoff. The result is that we can readily recover the personal incumbency advantage from the RD effect.

Our model shows that the personal incumbency and the RD design are connected in a way that has previously gone unnoticed: the RD design double-counts the personal incumbency advantage. This occurs because each of the two groups of districts compared in the RD estimand — districts where Democrats barely won and districts where Republicans barely won — has an incumbent running for reelection at $t + 1$. In other words, a party either wins and gains the incumbency advantage in the next election, or it loses — in which case the opposition party gains the incumbency advantage in the next election. The differential between the two party vote shares at $t + 1$ is the observed RD effect, which is twice the personal incumbency advantage. This double-counting feature of the RD design is likely the reason why empirical RD estimates of the incumbency advantage are much larger than those obtained with traditional approaches such as sophomore surge and the Gelman and King (1990) method.

Our main contribution is to clarify the subtle conceptual difficulties that arise when the incumbency advantage is studied with the RD design — difficulties that have previously been unrecognized. Of course, many of the issues we identify (strategic retirement, incumbent's inherent quality advantage, etc.) also complicate traditional strategies for estimating the incumbency advantage, which means that the RD design is still a very promising strategy and may even be superior to its alternatives. However, our message is that the quasi-experimental RD design does not automatically solve all previous difficulties — important conceptual issues remain and must be grappled with.

A secondary contribution of our article is to provide a new RD-based method to estimate the personal incumbency advantage. The RD estimand is a local effect for extremely competitive races that may or may not be informative about the incumbency advantage in lopsided elections. Thus, seen as a strategy for estimating the personal incumbency advantage, the generalizability of our contribution is limited. However, we also note the substantive interest in the size of the incumbency effect

is at its maximum for closely contested elections. One can conjecture that open-seat winners of very safe seats earn little in the way of the incumbency advantage because they have little incentive to bolster their electoral performance. But whether this conjecture is true or not is of little relevance to the dynamics of legislative elections.

After presenting our model for open seats, we illustrate our setup with a brief analysis of northern open-seat U.S. House elections between 1968 and 2008. We also briefly discuss how to use our framework to study incumbent races and explain why the required assumptions in this case are considerably more complex. Our fully general model is presented in the Supplemental Appendix, which also includes a recast of our parametric model in terms of potential outcomes.

1 Recovering the Personal Incumbency Advantage From an RD Design: The Case of Open Seat Races

The original RD setup assumes a two-party system, with the analysis usually focusing on the Democratic share of the two-party vote — the “score” or “running variable” that determines whether a district’s Democratic candidate wins a given election, which we assume occurs at time t . When the Democratic vote share exceeds 50%, the Democratic candidate becomes an incumbent; and when this vote share falls below 50%, the Republican candidate becomes the incumbent. Under appropriate continuity assumptions, this discontinuity at the 50% cutoff can be used to recover a measure of the incumbency advantage at the cutoff (see Hahn *et al.*, 2001; Imbens and Lemieux, 2008; Lee, 2008). The intuitive interpretation is that districts where the Democratic party barely loses the election at t , say obtain a vote share of 49.9%, are on average similar to districts where Democratic candidates barely win the election, say obtain a vote share of 50.1%, in terms of important characteristics such as partisanship, candidate quality, and previous vote shares. Thus, near the 50% cutoff, winning or losing at t can be regarded to be *as if* randomly assigned, and we can compare the vote share of barely winner districts and barely loser districts in the following election, which we assume occurs at $t + 1$, to recover a measure of the incumbency advantage

at the cutoff.⁴ The parameter of interest is the local average treatment effect of a Democratic win at the cutoff — measured as the difference between the average vote shares in districts where the Democratic party barely won the previous election and districts where the Democratic party barely lost the previous election.

This basic setup is our starting point. The treatment group is defined as those districts where the Democratic party barely won the election at t , and the control group as those districts where the Democratic party barely lost the election at t . This leads naturally to the definition of the usual RD estimand: the average vote share at $t + 1$ in treatment districts minus the average vote share at $t + 1$ in control districts.

Our basic model considers only those districts where there is an open seat race at the time election t , that is, districts where no incumbent is running at t . Our unit of observation is the congressional district, which we index by i , and V_{it} and V_{it+1} are the Democratic share of the two-party vote at elections t and $t + 1$, respectively, in district i . The Republican share of the vote is $1 - V_{it+1}$ so it suffices to focus on V_{it+1} .

We model V_{it+1} as

$$V_{it+1} = \text{Par}_{it+1} + \theta I_{it+1} + (D_{it+1} - R_{it+1}) + e_{it+1} \quad (1)$$

where we introduce the concept of *Par* — a measure of the baseline vote for the Democratic party in district i , given the district's partisanship, the election year's partisan trend, no incumbent candidate, and Democratic and Republican candidates of average quality (Erikson and Palfrey, 1998).⁵ In Equation (1), Par_{it+1} is the district's *Par* at $t + 1$; D_{it+1} and R_{it+1} are the *added* quality of the Democratic and Republican candidates, respectively, running at $t + 1$ in district i above the quality of the average open seat candidate in their respective parties. In other words, candidate quality is normed to be conditional on *Par*

⁴Despite the fact that this is a common and intuitive interpretation of the RD estimand, stronger assumptions than continuity are required to ensure that such an as-if random interpretation is valid. We ignore this distinction here, but see Cattaneo *et al.* (2015) for details.

⁵This conceptualization of *Par* is similar to that of Erikson and Palfrey (1998), except we subsume candidate spending under candidate quality. We introduce *Par* to capture the baseline vote in our model, but under the assumptions we introduce below it, we need not estimate it to recover the personal incumbency advantage.

so that, for any value of *Par*, each party's expected candidate quality (conditional on it being an open seat) is zero. Thus, in expectation, both candidates are of average quality (conditional on an open seat and *Par*). The quantity $(D_{it+1} - R_{it+1})$ is the "quality differential" between the Democratic and Republican candidates, which on average equals zero for open seats. Finally, I_{it+1} equals 1 if the newly elected time t winner is a Democrat who seeks reelection at $t + 1$, -1 if the newly elected time t winner is a Republican seeking reelection at $t+1$, and 0 if there is no incumbent running at $t + 1$; e_{it+1} is a residual. For simplicity of exposition, our model does not include a possible edge for the retiring incumbent's party — the partisan incumbency advantage discussed above — but we discuss the consequences of this simplification below.

We only include cases where the time t winner seeks reelection at $t + 1$ — implying that I_{it+1} is never zero. By excluding districts that are open seats at $t + 1$, we may introduce bias from selective mortality, a possibility here if anticipation of the $t + 1$ outcome affects retirement decisions. As we show briefly below, this is not a restrictive assumption in our empirical illustration, since almost all U.S. House freshman incumbents seek reelection. The more general version of our model that we present in the Supplemental Appendix accommodates the possibility of open seats at $t + 1$ under appropriate local assumptions,⁶ which leads to a slightly modified estimand.

Henceforth, we use V_{it+1}^w and V_{it+1}^l to denote, informally, the expected value of the Democratic vote share in Democrats' barely-winner and barely-loser districts, respectively, a notation we also use to denote analogous expectations for the other variables in our model.⁷ Taking

⁶In applications where these local assumptions are implausible, researchers may bound the estimate by retaining the observations where incumbents retire at $t + 1$ and considering how the estimated effect changes when different values are imputed to the missing outcome (i.e., to the unobserved incumbent vote share) for these observations. See Manski (2007) for a comprehensive approach to nonparametric bounds.

⁷More formally, for any variable y , the w and l subscripts indicate, respectively, the right and left limits of $E(y|V_{it} = v)$ when v approaches 50% from above and below, i.e. $y^w \equiv \lim_{v \rightarrow \frac{1}{2}^+} E(y|V_{it} = v)$ and $y^l \equiv \lim_{v \rightarrow \frac{1}{2}^-} E(y|V_{it} = v)$. We adopt the more intuitive interpretive interpretation of expectations within a narrow window for simplicity of exposition.

expectations for barely-winner and barely-loser districts in Equation (1) yields

$$\begin{aligned} V_{t+1}^w &= \text{Par}_{t+1}^w + \theta \cdot I_{t+1}^w + (D_{t+1}^w - R_{t+1}^w) \\ V_{t+1}^l &= \text{Par}_{t+1}^l + \theta \cdot I_{t+1}^l + (D_{t+1}^l - R_{t+1}^l) \end{aligned}$$

Combining these expressions, the RD estimand becomes

$$\begin{aligned} \tau^{\text{RD}} &= V_{t+1}^w - V_{t+1}^l \\ &= (\text{Par}_{t+1}^w - \text{Par}_{t+1}^l) + \theta \cdot (I_{t+1}^w - I_{t+1}^l) \\ &\quad + (D_{t+1}^w - R_{t+1}^w) - (D_{t+1}^l - R_{t+1}^l) \end{aligned} \quad (2)$$

We are interested in learning about θ but, as shown in Equation (2), this parameter is not immediately available: τ^{RD} recovers the personal incumbency advantage plus the difference in Par ($\text{Par}_{t+1}^w - \text{Par}_{t+1}^l$) and the candidate quality differentials $(D_{t+1}^w - R_{t+1}^w) - (D_{t+1}^l - R_{t+1}^l)$. Note that, given our assumption of no retirement at $t + 1$, $I_{t+1}^w = 1$ and $I_{t+1}^l = -1$.

By virtue of the RD assumptions, barely-winner and barely-loser districts' Par at t are on average equal to each other, $\text{Par}_t^w = \text{Par}_t^l$. We assume any average change in Par between t and $t + 1$ affects barely winner and barely loser districts similarly, so that the equality still holds in the following election, $\text{Par}_{t+1}^w = \text{Par}_{t+1}^l$. Under this simplifying assumption, it follows that

$$\tau^{\text{RD}} = \theta \cdot (I_{t+1}^w - I_{t+1}^l) + \text{QD}_{t+1}^w + \text{QD}_{t+1}^l \quad (3)$$

where we have defined the *quality differential* terms $\text{QD}_{t+1}^w \equiv (D_{t+1}^w - R_{t+1}^w)$ and $\text{QD}_{t+1}^l \equiv (R_{t+1}^l - D_{t+1}^l)$.

Equation (3) shows the RD estimand as the sum of three terms. The first term is the direct personal incumbency advantage, θ , multiplied by $(I_{t+1}^w - I_{t+1}^l) = 2$. The second term, QD_{t+1}^w , is the average difference in quality between the Democratic and Republican candidates at $t + 1$ in districts where the Democrat barely wins. The third term, QD_{t+1}^l , is the average difference in quality between the Republican and Democratic candidates at $t + 1$ in districts where the Democrat barely loses.

To study the quality differential terms, we first consider what happens at election t . The as-if randomness of close races in the RD design

guarantees that, at t , Democratic candidates in barely-winner districts are of equal average quality to Democratic candidates in barely-loser districts ($D_t^w = D_t^l$), and Republican candidates in barely-winner districts are of equal average quality than Republican candidates in barely-loser districts ($R_t^w = R_t^l$). Moreover, although generally in any given race the winner candidate will tend to be of higher quality than the loser candidate, we assume that in very close open seat races winner and loser will tend to be of the same average quality. Since we define quality relative to Par, this implies that in open seat races near the 50% cutoff, $D_{it}^w = R_{it}^w = D_{it}^l = R_{it}^l = 0$.

Next we turn to $t + 1$, when the open seat winners become first-term incumbents. In order to characterize the quality differential in this election, we assume that, *on average*, the open seat winners at time t maintain their original quality from time t , so that $D_{t+1}^w = D_t^w = 0$ and $R_{t+1}^l = R_t^l = 0$. This leaves only D_{t+1}^l and R_{t+1}^w from Equation (3), which we assume are negative under a scare off argument. At time $t + 1$, the barely-losing parties (at time t) must select new candidates from their candidate pools. But instead of average quality candidates as at time t , they may be forced to pick worse candidates due to a scare off effect: potential high-quality challengers, knowing that incumbents have access to resources they can exploit for electoral advantage, may be discouraged from entering the race (Cox and Katz, 1996). In other words, as the freshman incumbents gain their new incumbency advantage, they may adversely affect the quality of their opponents. This strategic behavior results in $t + 1$ challengers of lower average quality than the typical open seat candidate, thus leading to $D_{t+1}^l < 0$ and $R_{t+1}^w < 0$ — recall that D_{t+1}^l (R_{t+1}^w) represents the quality of the Democratic (Republican) challenger in districts where there is a Republican (Democratic) incumbent at $t + 1$.

As mentioned above, the scareoff may itself be decomposed into two sources. First, challengers may be deterred from entering the race simply because their opponent is an incumbent who has access to perquisites of office which are believed to translate into an electoral advantage. We call this the *incumbency scareoff*, as it arises directly from the incumbency status of the incumbent candidate. Second, challengers may also be deterred from entering the race due to the incumbent's initial quality, apart from the personal incumbency advantage. We call this the *quality scareoff*. Importantly, in extremely close open seat races,

our assumptions guarantee that the expected quality scareoff is zero, since the winner and the loser are assumed to be of equal average quality. In these races, the only component of the scareoff is the incumbency scareoff.⁸ Hereafter, we use the symbol σ to refer to the incumbency scareoff.

The combination of these conditions (restriction to open seats at t , definition of quality relative to Par, winner and loser of average quality, scare off effect) leads to $QD_{t+1}^w = (D_{t+1}^w - R_{t+1}^w) = -R_{t+1}^w > 0$ and $QD_{t+1}^l = (R_{t+1}^l - D_{t+1}^l) = -D_{t+1}^l > 0$, that is, to a quality differential that is positive and equal to the incumbency scareoff effect in both barely-winning and barely-losing districts. For simplicity, we assume that the incumbency scareoff is equal for both parties: $-D_{t+1}^l = -R_{t+1}^w \equiv \sigma$.

Under these conditions, the RD estimand in Equation (3) simplifies to

$$\begin{aligned}\tau^{\text{RD}} &= 2\theta + (-R_{t+1}^w - D_{t+1}^l) \\ &= 2\theta + 2\sigma,\end{aligned}$$

which shows that *the usual RD estimate is twice the personal incumbency advantage*.⁹ Thus, we can recover the personal incumbency advantage from the RD design simply by dividing by 2:

$$\frac{\tau^{\text{RD}}}{2} = \theta + \sigma$$

Why does this double-counting phenomenon arise? In *barely-winner* districts, the Democratic vote share *increases* by the personal incum-

⁸We are not claiming that a quality scareoff does not exist in general, only that it does not exist at the RD cutoff. Given a valid RD design, open-seat winners and losers must *on average* be of the same quality at t at the cutoff. Then, at $t + 1$ the winner's *expected* quality will also be of average quality. Apart from the incumbency factor, the prospective challenger at $t + 1$ faces the same circumstances as her party's loser at time t , an expectation of an "average" candidate. Thus, at the cutoff, there is no quality scareoff. A generalization of our model that allows for both types of scareoff is discussed in the Supplemental Appendix.

⁹Note that the assumption that the direct personal incumbency advantage θ and the incumbency scareoff σ are the same for the two parties is not at all required to obtain this double-counting result. It is simple to show that if these effects vary by party, the RD estimand equals twice the *average* personal incumbency advantage across both parties, i.e., $\tau^{\text{RD}} = 2(\frac{\theta^d + \theta^r}{2} + \frac{\sigma^d + \sigma^r}{2})$ where superscripts "d" and "r" refer to the Democratic and Republican parties, respectively.

bency advantage because the newly elected incumbent is a Democrat — with both direct access to the benefits of office and the ability to scare off strong Republican opponents. Conversely, in *barely-loser* districts, the Democratic vote share *decreases* by the personal incumbency advantage because the newly elected incumbent is a Republican with access to the benefits of office and the ability to deter strong Democratic challengers from entering the race. Thus, for the subset of districts where incumbents are running for reelection at $t + 1$, when we subtract the $t + 1$ Democratic party's vote share in barely-loser districts from the Democratic vote share in barely-winner districts, we count the personal incumbency twice.

Before turning to our empirical illustration, we note that our model makes the assumption that there is no partisan incumbency advantage. A party incumbency advantage would result if voters gave a premium to the incumbent party apart from the incumbent party's candidate. This party advantage could arise in an open seat if, for instance, voters took into account which party has held the seat previously. The retiree's party could hold an advantage because voters think it is generally a good idea to vote for the party they elected last time. Or it could be that retiring incumbents bequeath powerful campaign organizations to their parties' replacement candidates that their rivals cannot match. If this were the case, we would see a larger RD "effect" for open seats when the incumbent party at $t - 1$ (that is, the retiree's party) loses than when it wins. This is because the differential when the incumbent party loses would include both a loss of any party effect by one party and the gain of the party effect by the other.

Our model could be easily generalized to include a party incumbency effect in addition to the personal incumbency advantage. In this case, assuming symmetry between the parties and additive separability, the RD estimand would equal twice the sum of the personal incumbency advantage plus the partisan incumbency advantage. Naturally, recovering each of these effects individually (as opposed to their sum) would require additional assumptions. But this extension to our model would not alter its two main conceptual insights, namely, that a focus on open seats allows us to avoid the spurious advantage created by incumbents' higher-than-average quality, and that the usual RD estimand double counts the quantities of substantive interest.

2 Illustration: RD Estimates of the Personal Incumbency Advantage in the U.S. House

We now illustrate our theoretical findings with a brief analysis of first-year U.S. House incumbents following their open seat victories. We compiled election returns for all U.S. House elections in the period 1968–2008 from the *CQ* Voting and Elections Collection. Our analysis excludes South and Border states.¹⁰ The unit of observation is a congressional district, the running variable or score is the Democratic margin of victory at t and the outcome variable is the Democratic share of the total vote at $t + 1$. There are 399 cases. Excluded are the 14 instances when an open seat winner did not contest the next election. Excluding these cases does not threaten our inferences because the problem of strategic retirement vanishes completely in very competitive races: in the period we study, only one open seat winner within the 48–52% vote window retired at election $t + 1$. Also excluded are instances where redistricting occurs between the open-seat election and the freshman election.

Table 1 shows results of parametric and nonparametric estimation of the RD effect. In the parametric analysis, we estimate a linear regression of the outcome on the score and the incumbency variable (Dem. Win t), which takes on values -1 if the time t winner (and thus freshman incumbent) is a Republican and $+1$ if a Democrat — effectively undoing the double counting. In the nonparametric results, we estimate the difference in the intercepts of two local linear regressions of the outcome on the score at either side of the cutoff, using mean-squared-error optimal bandwidth and robust confidence intervals (Calonico *et al.*, 2014b).

As shown in the first column of Panel (A), the estimated personal incumbency advantage estimated with a simple parametric linear regression is 6.80 percentage points. The second column's equation adds

¹⁰We start in 1968 because of the well-documented growth of the incumbency advantage in the 1960s. We exclude South and Border state districts because — at least early in the period — it is difficult to estimate their partisan loyalties from presidential voting, a matter that takes on importance in the analysis below. The states excluded are Alabama, Arkansas, Florida, Georgia, Kentucky, Louisiana, Maryland, Mississippi, North Carolina, Oklahoma, South Carolina, Tennessee, Texas, Virginia and West Virginia.

Table 1: Linear regression RD estimates of incumbency advantage in U.S. House elections All open seat contests at t where freshman incumbents run again at $t + 1$.

Panel (A). Parametric estimation: linear regression			
Dem. Win t (Personal incumbency advantage)	6.80 (0.60)	7.66 (0.54)	7.85 (0.53)
Dem. Vote t	0.74 (0.05)	0.76 (0.04)	0.54 (0.06)
Presidential Vote t	—	—	0.28 (0.06)
RMSE	8.09	7.02	6.83
Year effects	No	Yes	Yes
Panel (B): Non-parametric local linear regression			
RD effect (τ^{RD})	15.12 [9.83, 22.16]		
Personal incumbency advantage ($\tau^{\text{RD}}/2$)	7.56 [4.91, 11.08]		

Note: There are 399 observations, in both panels; sample excludes southern and border states. In Panel (A), results are from parametric linear regression with standard errors in parenthesis; for midterm years, presidential vote equals previous presidential election in the district. In Panel (B), results are from local linear regression with triangular kernel using the robust confidence intervals and CCT implementation of mean-squared-error optimal bandwidth selector developed by Calonico *et al.* (2014b). Estimation implemented in Stata package `rdrobust` by Calonico *et al.* (2014a). Estimated main and pilot bandwidths are 11.23 and 20.02, respectively; bias estimated with quadratic polynomial. 95% robust confidence intervals are in brackets.

year dummy variables, which add precision as seen by the reduction in the RMSE. Now the estimated incumbency effect rises almost a point beyond the estimate in the first column, with a slightly tighter standard error. The third column adds the presidential vote as a further control, giving the tightest prediction of all and a slightly higher estimated incumbency advantage of 7.85.

Finally, the nonparametric results in Panel (B) show that the standard RD effect in our northern open-seat sample is 15.12: when the Democratic party barely wins election t , its vote share in election $t + 1$ is about 15 percentage points higher than it would have been if it had barely lost. In order to recover the personal incumbency effect from

this estimate, we divide this number by 2. Our estimate of the personal incumbency advantage is therefore 7.56 percentage points, with a 95% confidence interval ranging from 4.91 to 11.08. This is very similar to the results shown in Panel (A).

3 RD and Incumbent-Contested Seats

With our RD-based design, inference for incumbent-held seats poses a more serious challenge than for open seats. When discussing open seats with close races, it was safe to assume that the winning candidates' performance at $t+1$ is due to the direct and indirect effects of incumbency. This is because, in every district, the determination of the time t winner is a virtual coin flip between two candidates who in expectation are of average quality (given the district and the year), which justifies our assumption that the $t+1$ challenger will be of lower than average quality due to the scareoff effect. For incumbent races, we can make no such assumption. Embattled incumbents at the cusp of losing tend to be below-average candidates. Their challengers who take them to the 50% cutoff might be better than average candidates, although they may still be scared off by their opponent's incumbency. The math for open seats no longer applies, implying that, without additional assumptions, the model we introduced above cannot be used to estimate the personal incumbency advantage in incumbent races.

In our open seats analysis, two elements were needed to simplify the quality differential terms $QD_{t+1}^w \equiv (D_{t+1}^w - R_{t+1}^w)$ and $QD_{t+1}^l \equiv (R_{t+1}^l - D_{t+1}^l)$: (i) defining candidate quality conditional on Par and adopting the normalization $D_{it}^w = R_{it}^w = D_{it}^l = R_{it}^l = 0$ and (ii) assuming that, *on average*, the open seat winners at time t maintain their original quality from time t , i.e., $D_{t+1}^w = D_t^w = 0$ and $R_{t+1}^l = R_t^l = 0$, and that their opponents are of lower-than-average quality due to the scareoff effect. Under these assumptions, we obtained $QD_{t+1}^w = -R_{t+1}^w > 0$ and $QD_{t+1}^l = -D_{t+1}^l > 0$, which allowed us to show that each quality differential term was positive and exactly equal to the incumbency scareoff effect.

The assumptions for open seats are certainly strong, but we believe they can be justified. In contrast, offering a similar justification for the expected sign of the quality differentials in incumbent seats is much more challenging. If incumbents who barely survive are relatively weak,

they might be challenged by candidates of stronger-than-average quality, which would make the quality differential term negative and may offset the direct personal incumbency effect. But since even in this case the incumbency-induced scareoff effect would not cease to operate — everything else equal, incumbency is still a deterrent — ascribing a sign to the quality differential terms would require very strong assumptions regarding the relative qualities of incumbent and challengers. Depending on the particular electoral context under investigation, scholars may or may not feel comfortable making these assumptions.

4 Conclusion

The main virtue of the RD design as Lee (2008) applied it to the study of incumbency was to shift the focus from the incumbent candidate to the incumbent party, avoiding the need to make specific assumptions about how the strategic entry and exit of candidates into the race affects the estimates. The RD incumbency effect, focusing on the “overall advantage to the party”, provided an estimate of the incumbency advantage that entirely sidestepped the methodological difficulties that had been at the center of the incumbency advantage literature for decades. The methodological virtues of the design, however, came at the price of a vague conceptualization of the incumbency advantage, an issue that has gone largely unnoticed despite the growing popularity of the RD design among incumbency advantage scholars. Our paper focused on this issue, and studied the conditions under which the incumbency advantage as traditionally understood by the political science literature can be recovered from an RD design.

We show how the RD design can be used to identify the personal incumbency advantage — the specific advantage that incumbents obtain as incumbents that they did not have in their initial, non-incumbent race. We show that the necessary ingredients are (a) that in districts with elections near the 50% vote threshold at time t , winners and losers are of average quality and high-quality challengers are deterred from contesting the seat at $t + 1$, and (b) that there is little or no strategic retirement affecting reelection decisions at $t + 1$. Given these conditions, a simple additive model shows that the RD effect double counts the incumbency advantage, summing the personal vote gains from winners

in each party. We briefly illustrated our framework with an analysis of closely contested U.S. House open seats in northern races during the 1968–2008 period, where we believe these conditions are approximately met. For such districts, we estimate the gain from incumbency to be about seven percentage points. We also discussed how much stronger the assumptions must be to perform an analogous analysis of incumbent races.

In discussing the assumptions under which the personal incumbency advantage can be recovered from an RD design, our paper also illustrates a more general methodological point regarding the use of natural experiments or quasi-experimental designs to study substantive questions. The use of quasi-experimental designs is on the rise because these designs offer sources of exogenous variation in contexts where experimentation is typically infeasible. In our context, the RD offers as-if randomization of winner and loser candidates, which makes both groups comparable and opens many possibilities for causal inference. However, the main feature of natural experiments is that researchers are not in control of the treatment assignment, which often means that the quantity that is readily available from the design is related but not exactly equivalent to the quantity of substantive interest (Sekhon and Titiunik, 2012).

In our example, even under the assumption that winners and losers are as-if randomly assigned, the incumbency advantage estimated under the RD design differs from the personal incumbency advantage that is of central interest, and only after we define a set of precise assumptions can we recover the latter from the quasi-experimental estimate. These assumptions require careful theoretical understanding of the phenomenon of incumbency advantage. This is a recurring theme with the use of quasi-experimental designs: their full promise can only be realized with careful theoretical understanding of the substantive question under study, something that is independent of whether the quasi-experiment is a valid source of exogenous variation.

References

- Ansolabehere, S. and J. M. Snyder Jr. 2002. “The Incumbency Advantage in US Elections: An Analysis of State and Federal Offices, 1942–2000”. *Election Law Journal*. 1(3): 315–338.

- Broockman, D. E. 2009. "Do Congressional Candidates Have Reverse Coattails? Evidence from a Regression Discontinuity Design". *Political Analysis*. 17(4): 418–434.
- Butler, D. M. 2009. "A Regression Discontinuity Design Analysis of the Incumbency Advantage and Tenure in the U.S. House". *Electoral Studies*. 28(1): 123–128.
- Calonico, S., M. Cattaneo, and R. Titiunik. 2014a. "Robust Data-Driven Inference in the Regression-Discontinuity Design". *Stata Journal*. 14(4): 909–946.
- Calonico, S., M. Cattaneo, and R. Titiunik. 2014b. "Robust Nonparametric Confidence Intervals for Regression-Discontinuity Designs". *Econometrica*. 82(6): 2295–2326.
- Cattaneo, M., B. Frandsen, and R. Titiunik. 2015. "Randomization Inference in the Regression Discontinuity Design: An Application to the Study of Party Advantages in the U.S. Senate". *Journal of Causal Inference*. 3(1): 1–24.
- Cox, G. W. and J. N. Katz. 1996. "Why Did the Incumbency Advantage in U.S. House Elections Grow?" *American Journal of Political Science*. 40(2): 478–497.
- Erikson, R. S. 2015. "The Congressional Incumbency Advantage Over Sixty Years: Measurement, Trends, and Implications". In: *Governing in a Polarized Age: Elections, Parties, and Political Representation*. Ed. by A. Gerber and E. Hendry. forthcoming. New York: Cambridge University Press.
- Erikson, R. S. and T. R. Palfrey. 1998. "Campaign Spending and Incumbency: An Alternative Simultaneous Equations Approach". *Journal of Politics*. 60(May): 355–373.
- Gelman, A. and G. King. 1990. "Estimating Incumbency Advantage without Bias". *American Journal of Political Science*. 34(4): 1142–1164.
- Golden, M. A. and L. Picci. 2015. "Redistribution and Reelection Under Proportional Representation: The Postwar Italian Chamber of Deputies". Forthcoming, *Legislative Studies Quarterly*.
- Hahn, J., P. Todd, and W. Van der Klaauw. 2001. "Identification and Estimation of Treatment Effects with a Regression Discontinuity Design". *Econometrica*. 69(1): 201–209.
- Imbens, G. W. and T. Lemieux. 2008. "Regression Discontinuity Designs: A Guide to Practice". *Journal of Econometrics*. 142(2): 615–635.

- Lee, D. S. 2008. "Randomized Experiments from Non-random Selection in U.S. House Elections". *Journal of Econometrics*. 142(2): 675–697.
- Levitt, S. D. and C. D. Wolfram. 1997. "Decomposing the Sources of Incumbency Advantage in the U.S. House". *Legislative Studies Quarterly*. 22(1): 45–60.
- Manski, C. F. 2007. *Identification for Prediction and Decision*. Harvard University Press.
- Sekhon, J. S. and R. Titiunik. 2012. "When Natural Experiments Are Neither Natural Nor Experiments". *American Political Science Review*. 106(1): 35–57.
- Trounstine, J. 2011. "Evidence of a Local Incumbency Advantage". *Legislative Studies Quarterly*. 36(2): 255–280.
- Zaller, J. 1998. "Politicians as Prize Fighters: Electoral Selection and Incumbency Advantage". In: *Party Politics and Politicians*. Ed. by J. Geer. Baltimore: Johns Hopkins University Press.