Quality-Based Explanations of Incumbency Effects[*]

Andrew C. Eggers[†]

June 30, 2016

Abstract

Empirical studies of incumbency effects continue to accumulate, but progress in explaining these findings is modest. I introduce a simple framework that clarifies how differences in candidate characteristics (i.e. quality) could account for empirical findings of incumbency effects based on regression discontinuity designs (RDD). The key observation is that although RDD ensures that marginal winners and losers of close elections are comparable, the candidates who compete in the next election may differ for many reasons, of which "scare-off" (the tendency of incumbents to deter strong challengers) is just one. Using a simple model, I show that quality differences between marginal incumbent-party candidates and their opponents may persist even when incumbency is irrelevant to voters, incumbent retirement is random, and all new candidates are drawn from the same candidate pool. I conclude by showing how quality differences can help explain puzzling findings in prior empirical work.

1 Introduction

Over the past several decades, a huge empirical literature in American politics has shown that incumbents enjoy electoral advantages: for a variety of possible reasons, incumbent candidates are thought to be more successful than comparable candidates who do not already hold the office.[1] In recent years scholars have found evidence of incumbency advantage in other established democracies[2] and, intriguingly, some evidence of incumbency disadvantage in several settings (Uppal, 2009; Klašnja, 2015a; Ariga, 2015; Klašnja and Titiunik, 2016).[3] Although explanations for these findings vary, researchers have tended to see both incumbency advantage and incumbency disadvantage as fundamentally normatively troubling: incumbency advantage suggests that incumbents use their office to insulate themselves from electoral pressure, while incumbency disadvantage suggests that incumbents fail to live up to voters' expectations.

This paper examines one particular explanation for incumbency effects in order to clarify the mechanisms that might explain empirical findings. I focus on quality-based explanations of incumbency effects, by which I mean systematic differences between the characteristics of incumbent-party candidates and their opponents that lead to disproportionate incumbent success or failure. Most other explanations of incumbency effects emphasize the resources, visibility, and power that incumbents gain from officeholding (McKelvey and Riezman, 1992; Bevia and Llavador, 2009; Caselli et al., 2013; Klašnja, 2015c; Fouirnaies and Hall, 2014; Klašnja, 2015b); in short, these explanations emphasize what incumbents *do* (and what their opponents cannot do). Quality-based explanations, in contrast, emphasize who incumbents *are* (and who their opponents are).

My point of departure in defining and exploring quality-based incumbency advantage and disadvantage is the empirical study of party incumbency effects using the regression discontinuity design, i.e. RDD (Lee, 2008). This empirical strategy, which uses close elections to estimate the effect of a party's victory or loss in one election on its success in the next election, has quickly

---

[1]Important contributions include Erikson (1971); Mayhew (1974); Fiorina (1977); Alford and Hibbing (1981); Cain, Ferejohn and Fiorina (1987); Gelman and King (1990); King (1991); Cox and Morgenstern (1993); Cox and Katz (1996); Levitt and Wolfram (1997); Ansolabehere, Snyder and Stewart (2000); Ansolabehere and Snyder Jr (2002); Lee (2008).

[2]See e.g. Katz and King (1999); Hainmueller and Kern (2008); Horiuchi and Leigh (2009); Ariga (2010); Kendall and Rekkas (2012).

[3]De Magalhaes (2015) highlights potential problems in some of the findings of incumbency disadvantage.

become the dominant approach to studying incumbency effects both in the U.S. and in other countries (e.g Hainmueller and Kern, 2008; Butler, 2009; Kendall and Rekkas, 2012; Fouirnaies and Hall, 2014; Eggers et al., 2015; Ariga, 2015; Hall and Snyder, 2015; Fowler, 2015b). Erikson and Titiunik (2015) and Fowler and Hall (2014) explore settings and assumptions under which estimates of the party incumbency effect of Lee (2008) can be used to measure other quantities of interest, and clearly my analysis of incumbency effects applies to those efforts. My analysis also relates, although more indirectly, to RDD-based estimates of incumbency effects focused on individuals rather than parties (Trounstine, 2011; De Magalhaes, 2015). As I explain in Section 5 below, all of the mechanisms that contribute to quality-based incumbency advantage or disadvantage as I define it would also contribute to estimates of individual incumbency effects as in Trounstine (2011) and De Magalhaes (2015).

How can quality differences contribute to findings of incumbency effects? The key idea is that while the marginal winners and losers of "coin-flip" elections should be equal in average quality, these winners and losers may not be the candidates who face each other in the elections that follow close elections; because selection into candidacy likely differs for the winning and losing party, incumbent-party candidates may be systematically stronger or weaker than their opponents. (One indication of how selection into candidacy differs for winning losing parties is that the marginal winner of elections to the U.S. House runs again about 70% of the time while the marginal loser runs again only about 25% of the time, as reported in more detail below.) I show that there are three distinct mechanisms by which the replacement of marginal candidates can yield quality-based incumbency advantage or disadvantage: (1) a difference between the quality of marginal winners who run again and the quality of marginal losers who run again; (2) a difference between the quality of candidates who replace marginal winners and the quality of candidates who replace marginal losers; and (3) a difference between the rate at which marginal winners are replaced and the rate at which marginal losers are replaced. I highlight three ways in which the third of these mechanisms can lead to quality-based incumbency effects, the most surprising of which I call "selection into marginality": if marginal candidates are stronger or weaker on average than the candidate pool (which I show can happen due either to electoral selection or to asymmetries in the distribution of candidate quality, and *must* happen due to

one or the other), then there will be quality-based incumbency advantage or disadvantage as long as marginal winners are more or less likely to run again than marginal losers.[4]

To the extent that the literature on incumbency effects has considered quality-based explanations, it has done so overwhelmingly by focusing on "scare-off", which is the tendency of incumbents to deter strong potential challengers (e.g. Cox and Katz, 1996; Levitt and Wolfram, 1997; Hall and Snyder, 2015). Part of the contribution of this paper is to broaden our understanding of how scare-off can happen while situating scare-off in the context of other quality-based explanations for incumbency effects. The existing literature views scare-off as a mechanism by which the officeholding advantages of incumbents are amplified; for example, if voters support incumbents out of gratitude for constituency service, then strong candidates may prefer not to challenge incumbents, such that incumbents benefit from both voter gratitude and weak opposition. I view scare-off as one way in which quality differences between marginal incumbents and their opponents might arise, but it is certainly not the only way and it may not even be the principal way. I also argue that scare-off could occur not just in response to incumbents' officeholding benefits (as assumed in the existing literature) but also in response to quality differences that arise for other reasons. More broadly, I show that there are quality-based mechanisms that (unlike scare-off) can produce incumbency advantages or disadvantages even when officeholding itself has no electoral impact at all.

My work relates to Ashworth and Bueno de Mesquita (2008), whose "quality difference model" posits that incumbency advantage arises through electoral selection: the logic of that model is simply that incumbents have won before, and the qualities that helped them win are likely to remain appealing to voters.[5] At first glance Ashworth and Bueno de Mesquita (2008)'s analysis may not seem directly relevant to empirical work on incumbency effects because they seek to explain the success of the average incumbent, while empirical researchers typically seek to explain the success of the *marginal* incumbent with an RDD or other control strategy. One of the contributions of my analysis is to show how electoral selection as studied by Ashworth and Bueno de Mesquita (2008) does in fact contribute to findings of incumbency effects even

---

[4]This statement also assumes that all new candidates emerge from the same candidate pool and candidates who re-run are typical of marginal winners and losers.

[5]Ashworth and Bueno de Mesquita (2008)'s model also incorporates a scare-off effect: incumbents with sufficiently high quality run uncontested.

when we focus on marginal incumbents. The key idea is that if incumbents are stronger than average, then marginal winners of elections involving incumbents will also be stronger than average, which means they will have a quality advantage if they subsequently face an average challenger.[6] Thus while RDD and other control strategies may seem to render electoral selection irrelevant for explaining findings of incumbency advantage, I argue that it may in fact play a crucial and under-appreciated role.

2   Defining quality-based incumbency effects

I assume that there is a continuum of voters choosing between a candidate from party a and a candidate from party b. For voter i, the utility of electing the candidate from party a is given by

$$u_i(a) = \theta_a + v_i(a) + \gamma I_a$$

where θ_a is the quality of a's candidate, $v_i(a)$ is i's valuation of a's party, and I_a is an indicator equal to 1 if a's candidate is the incumbent. Given a similar expression for party b, the voter prefers a if

$$\theta_a - \theta_b + v_i(a) - v_i(b) + \gamma(I_a - I_b) > 0.$$

In this paper, a candidate's "quality" is a measure of the extent to which voters find the candidate's intrinsic, time-invariant characteristics attractive. I have in mind characteristics such as appearance, intelligence, and background, i.e. hometown, educational credentials, profession, ethnicity, etc., that are determined before the candidate competes for office and thus not affected by whether the candidate actually wins. To the extent that voters value something that candidates gain from previously winning the office, such as seniority or experience in office, this is best captured by γ in my framework. For simplicity I define quality in a way that excludes the potential value of prior electoral experience or other experiences a candidate gains after entering politics, but the framework I introduce below could easily be extended to illuminate the role of non-office experience.[7]

---

[6]This idea can also be found in Erikson and Titiunik (2015), who recommend focusing on marginal open-seat elections as a result.

[7]For example, suppose candidates with more electoral experience are more effective candidates; then if

5

For simplicity, I assume that voters agree on what quality means and that they observe candidates' quality perfectly. Thus conditional on party, all voters would agree which of two candidates is more attractive. All of my analysis applies if we assume that voters receive unbiased private signals of quality and decide their vote on the basis of these private signals.[8]

Note that $v_i(a) - v_i(b)$ captures's i's preference for a's party relative to b's party; I assume that this party preference is distributed in the electorate uniformly on the interval $[-\frac{1}{2}, \frac{1}{2}]$, such that the electorate is equally balanced between those who prefer party a and those who prefer party b. (In the Appendix I show how the main results extend to the case where the electorate prefers one party or the other.) Supposing that all voters vote (and vote sincerely), the vote share for party a at time t is given by[9]

$$V_t = 1/2 + \theta_{at} - \theta_{bt} + \gamma(I_{at} - I_{bt}), \tag{1}$$

where I have added time subscripts to the quality measures and incumbency indicators.

The party incumbency effect (Lee, 2008) is defined as the effect for party a of winning a marginal election at time t on party a's vote share at $t + 1$. It is estimated via regression discontinuity design (RDD) as

$$\tau_{RDD} = \lim_{V_t \to .5^+} E\big[V_{t+1}\big|V_t\big] - \lim_{V_t \to .5^-} E\big[V_{t+1}\big|V_t\big], \tag{2}$$

where expectations are taken over districts or time periods. (Because party b's vote share is $1 - V_t$, the party incumbency effect would be the same if we took party b's perspective.)

Substituting equation 1 into equation 2, the incumbency effect can be expressed

$$
\begin{aligned}
\tau_{RDD} =\ & \lim_{V_t \to .5^+} E\big[1/2 + \theta_{a,t+1} - \theta_{b,t+1} + \gamma(I_{a,t+1} - I_{b,t+1})\big|V_t\big] - \\
& \lim_{V_t \to .5^-} E\big[1/2 + \theta_{a,t+1} - \theta_{b,t+1} + \gamma(I_{a,t+1} - I_{b,t+1})\big|V_t\big].
\end{aligned} \tag{3}
$$

Now, define $\bar{\theta}_{I,t+1}$ and $\bar{\theta}_{C,t+1}$ as follows:

$$\bar{\theta}_{I,t+1} \equiv \frac{1}{2}\left(\lim_{V_t \to .5^+} E\big[\theta_{a,t+1}\big|V_t\big] + \lim_{V_t \to .5^-} E\big[\theta_{b,t+1}\big|V_t\big]\right) \tag{4}$$

$$\bar{\theta}_{C,t+1} \equiv \frac{1}{2}\left(\lim_{V_t \to .5^-} E\big[\theta_{a,t+1}\big|V_t\big] + \lim_{V_t \to .5^+} E\big[\theta_{b,t+1}\big|V_t\big]\right). \tag{5}$$

In words, $\bar{\theta}_{I,t+1}$ denotes the quality of marginal incumbent-party candidates at time $t + 1$: it is the average across parties a and b of the quality of candidates who run for the party that marginally won the previous election. Similarly, $\bar{\theta}_{C,t+1}$ is the quality of marginal *challenger-party* candidates: it is the average across parties a and b of the quality of candidates who run for the party that marginally *lost* the previous election. Finally, let p_w denote the proportion of marginally-elected candidates who run for re-election:

$$p_w \equiv \frac{1}{2}\left(\lim_{V_t \to .5^+} E\big[I_{a,t+1} - I_{b,t+1}\big|V_t\big] - \lim_{V_t \to .5^-} E\big[(I_{a,t+1} - I_{b,t+1})\big|V_t\big]\right) \tag{6}$$

Dropping time subscripts, we can then substitute expressions 4, 5, and 6 into equation 3 to obtain

$$\tau_{RDD} = 2\left(\bar{\theta}_I - \bar{\theta}_C + \gamma p_w\right).$$

Thus the RDD estimate of party incumbency advantage reflects both a quality difference (i.e. the difference in average quality between candidates running for the party that marginally won the previous election and candidates running for the party that marginally lost the previous election) and the electoral value of incumbency to voters (multiplied by the probability that the incumbent runs for re-election). We will focus on the part of the overall incumbency effect that is due to quality differences:

Definition Consider $\bar{\theta}_I - \bar{\theta}_C$, which is the difference in average quality between candidates whose party narrowly won the previous election (i.e. marginal incumbent-party candidates) and candidates whose party narrowly lost the previous election (i.e. marginal challenger-party candidates), at the limit as the margin in the previous election goes to zero. Say that there is a *quality-based incumbency advantage* if this is positive, a *quality-based incumbency disadvantage* if this is negative, and *balance on quality* otherwise.

7

3 The three mechanisms for quality-based incumbency effects

To understand how quality-based incumbency advantage and disadvantage can emerge, consider Figure 1. At time t an election takes place; denote by $\bar{\theta}_w$ and $\bar{\theta}_l$ the average quality at time t of the candidates for the winning party and the losing party, respectively, at the limit as the margin between the winner and the loser goes to zero.[10] It should be clear from equation 1 that $\bar{\theta}_w = \bar{\theta}_l$: marginal winners and marginal losers are equal on average.[11] A proportion of marginal winners (denoted by p_w, as above) and losers (p_l) run again at time $t + 1$; denote by $\bar{\theta}_w^r$ and $\bar{\theta}_l^r$ the average quality of these re-running marginal winners and losers, respectively. Marginal winners and losers who do not run again are replaced by new entrants; denote by $\bar{\theta}_w^e$ and $\bar{\theta}_l^r$ the average quality of entrants who replace marginal winners and losers, respectively. Then we can state

$$\bar{\theta}_I - \bar{\theta}_C = \left(p_w \bar{\theta}_w^r + (1 - p_w) \bar{\theta}_w^e \right) - \left(p_l \bar{\theta}_l^r + (1 - p_l) \bar{\theta}_l^e \right). \tag{7}$$

Although there are many ways in which quality-based incumbency advantage and disadvantage can arise, it turns out that in any instance where balance on quality is not met, at least one of three mechanisms must be in play. First suppose the replacement rates are the same ($p_w = p_l = p$). Then the condition in Equation 7 becomes

$$\bar{\theta}_I - \bar{\theta}_C = p\left(\bar{\theta}_w^r - \bar{\theta}_l^r\right) + (1 - p)\left(\bar{\theta}_w^e - \bar{\theta}_l^e\right), \tag{8}$$

which makes it clear that when replacement rates are the same for marginal winners and losers, quality-based incumbency advantage or disadvantage can only arise if at least one of two conditions hold: the average quality of marginal winners who re-run is different from the average quality of marginal losers who re-run ($\bar{\theta}_w^r \neq \bar{\theta}_l^r$, a condition I call "differential selection into re-running") or the average quality of candidates who enter to replace marginal winners is different from the average quality of candidates who enter to replace marginal losers ($\bar{\theta}_w^e \neq \bar{\theta}_l^e$, a

---

[10]Thus the definitions of $\bar{\theta}_w$ and $\bar{\theta}_l$ are the same as the definitions of $\bar{\theta}_{I,t+1}$ and $\bar{\theta}_{C,t+1}$ above, except that they focus on the expected quality at time t rather than $t + 1$.

[11]Sketch of proof: V_t is continuous in both $\bar{\theta}_a$ and $\bar{\theta}_b$, which means the average quality of marginal winners for a (call this $\bar{\theta}_{w,a}$) is equal to the average quality of marginal losers for a (call this $\bar{\theta}_{l,a}$); the same can be said of party b. We have $\bar{\theta}_w = \frac{1}{2}(\bar{\theta}_{w,a} + \bar{\theta}_{w,b})$ and $\bar{\theta}_l = \frac{1}{2}(\bar{\theta}_{l,a} + \bar{\theta}_{l,b})$, and because $\bar{\theta}_{w,a} = \bar{\theta}_{l,a}$ and $\bar{\theta}_{w,b} = \bar{\theta}_{l,b}$, $\bar{\theta}_w = \bar{\theta}_l$.

Figure 1: Channels for quality-based incumbency effects
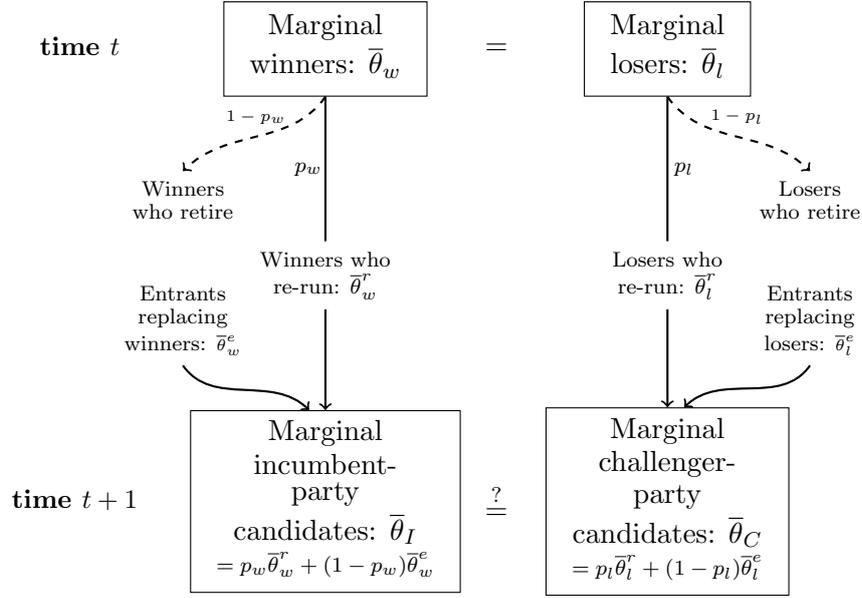
NOTE: Marginal winners and marginal losers have the same average quality ($\bar{\theta}_w = \bar{\theta}_l$), but the candidates who compete in the next election for the incumbent party and the challenger party may differ in quality depending on the proportion of marginal winners and losers who run in the next election (p_w and p_l), the quality of these re-running candidates ($\bar{\theta}_w^r$ and $\bar{\theta}_l^r$), and the quality of entering candidates ($\bar{\theta}_w^e$ and $\bar{\theta}_l^e$).

condition I call "differential selection into entry"). Now suppose neither of these conditions is met (such that $\bar{\theta}_w^r = \bar{\theta}_l^r = \bar{\theta}^r$ and $\bar{\theta}_w^e = \bar{\theta}_l^e = \bar{\theta}^e$). In that case the condition becomes

$$\bar{\theta}_I - \bar{\theta}_C = (p_w - p_l)\big(\bar{\theta}^r - \bar{\theta}^e\big), \tag{9}$$

which makes it clear that when there is neither differential selection into re-running or differential selection into entry, quality-based incumbency advantage or disadvantage can only arise if replacement rates differ for marginal winners and losers ($p_w \neq p_l$) *and* the average quality of candidates who re-run differs from that of replacement candidates ($\bar{\theta}^r \neq \bar{\theta}^e$). I will refer to this third mechanism as "differential replacement rates". In any given setting more than one of these mechanisms may operate, but this paragraph has shown that quality-based incumbency advantage or disadvantage cannot arise unless at least one of them operates. (See the Appendix for a schematic illustration of each mechanism.)

I now consider each of the three mechanisms in turn with the goal of clarifying how each mechanism might arise and therefore how empirical researchers might detect it.

Mechanism 1: Differential selection into re-running

Differential selection into re-running occurs if marginal winners and losers are not equal in average quality conditional on choosing to run again ($\bar{\theta}_w^r \neq \bar{\theta}_l^r$). Why might this happen?

The key difference between marginal winners and marginal losers is that marginal winners hold office and marginal losers do not. It may be reasonable for all sitting incumbents to consider trying to retain their seats, regardless of their abilities or attractiveness as candidates; it could be, for example, that marginal winners' decisions to re-run are unrelated to their quality. Meanwhile, marginal losers presumably have been pursuing other professional avenues, and it may be that only those with limited outside options (and lower quality) consider running again. This difference in the process of selection into re-running between marginal winners and losers would tend to produce quality-based incumbency advantage.

Differential selection into re-running may also arise due to a perceived electoral advantage to incumbency. Suppose that voters prefer incumbents for some reason (i.e. $\gamma > 0$). Then stronger marginal losers may choose to pursue other options rather than face an uphill climb against an incumbent in the next election. (Weaker marginal losers may have no other option.) The resulting quality-based incumbency advantage would contribute along with γ to the overall incumbency advantage. This is the mechanism conventionally referred to as "scare-off" as it applies to the decision of marginal winners and losers to re-run or not. (We will see shortly that "scare-off" can also apply to entry decisions.)

Differential selection into re-running could also produce quality-based incumbency disadvantage. It could be, for example, that the best marginal winners use office as a springboard to better opportunities – higher office, or lucrative jobs in the private sector – while the weakest marginal losers are discouraged from running again, with the result that re-running marginal winners are weaker on average than re-running marginal losers. Or, if voters prefer non-incumbents for some reason (i.e. $\gamma < 0$), then perhaps the best marginal incumbents would choose to look elsewhere rather than face an uphill battle as an incumbent, which would again produce quality-based incumbency disadvantage through what might be called "reverse scare-off".

Mechanism 2: Differential selection into entry

Differential selection into entry occurs when candidates who replace marginal winners and losers are not equal in average quality ($\overline{\theta}_w^e \neq \overline{\theta}_l^e$). Why might this happen?

One reason is that potential replacements for marginal losers may be affected by the perceived benefits or costs of incumbency; this is "scare-off" as it applies to entry decisions. That is, strong potential challengers may believe that voters prefer incumbents ($\gamma > 0$) and thus decide to pursue other options rather than enter to fight an uphill battle against an incumbent; this would tend to produce a quality-based incumbency advantage. As with selection into re-running, it could also work in reverse: if voters dislike incumbents, strong potential challengers could be attracted by the possibility of running against an incumbent, which would tend to produce a quality-based incumbency disadvantage.

Differential selection into entry could also be a response to differential selection into re-running. Suppose, for example, that only the best marginal winners re-run while only the worst marginal losers do; suppose also that potential candidates have only a noisy signal about their own type before they decide to enter. Then it could be that strong replacement candidates choose to enter for the winning party (given that their opponents are expected to be weak) while strong replacement candidates choose not to enter for the losing party (given that their opponents are expected to be strong). Thus there could be scare-off that responds not to voter preferences or incumbent actions but rather to other forms of quality-based incumbency advantage or disadvantage.

Mechanism 3: Differential replacement rates

Replacement rates differ when the proportion of marginal winning candidates who re-run is not the same as the proportion of marginal losing candidates who re-run ($p_w \neq p_l$). Why would this occur?

The clearest reason why the replacement rate may differ for marginal winners and losers is that marginal winners hold the office while marginal losers do not. Many marginal losers may pursue other professional avenues after losing, and by the time of the next election they may prefer to continue on this alternative path rather than interrupt it to campaign for office;

Table 1: Re-running by marginal winners and losers in U.S. elections

Dataset	Winner runs, loser does not	Loser runs, winner does not	Both run	Neither run
US House, 1948-2012	0.47	0.04	0.19	0.30
US State Legislatures, 1968-2010	0.44	0.05	0.12	0.38
US Statewide Offices, 1970-2012	0.53	0.05	0.06	0.36

NOTE: Analysis based on replication data from Hall and Snyder (2015); marginal elections are those where the margin in the previous race was within 5%.

for marginal winners, by contrast, it may be convenient to campaign and desirable to continue in the same office. Depending on the electoral context, officeholders may also be the party's nominee by default while non-officeholders may need to compete for the chance to run again. For both of these reasons we would expect a higher rate of re-running among winners than losers. In the U.S. (where incumbency effects have been studied more than anywhere else), marginal winners are in fact much more likely to run again than marginal losers. Table 1 shows that in almost half of all cases where the previous election was decided by less than 5% the winner ran again while the loser did not. (By contrast, the reverse occurred in no more than 5% of cases, and both ran in 20% of cases in the U.S. House but less frequently in other offices.)

In some settings, marginal losers may actually be more likely to run again than marginal winners. Some offices could be a springboard to other opportunities, such that most winners go on to other things while most losers try again. Also, where term limits are used it may be impossible for (some) winners to run again, while all losers are free to do so. (In Mexico, for example, incumbents are prohibited from running for re-election to most offices.) We will see below that this could be an important explanation for some findings of incumbency disadvantage.

If differential replacement rates are to produce quality-based incumbency advantage or disadvantage in the absence of either differential selection into re-running or differential selection into entry, it must also be the case that re-running candidates are stronger or weaker on average than replacement candidates. (See Equation 9.) How would this happen? I highlight three sub-mechanisms.

3a: Differential replacement rates and selection into re-running

Suppose that marginal winners and losers have average quality $\bar{\theta}$, and that all replacement candidates also have average quality $\bar{\theta}$ (i.e. $\bar{\theta}_w = \bar{\theta}_l = \bar{\theta}_w^e = \bar{\theta}_l^e = \bar{\theta}$), but that marginal winners

and losers who re-run have average quality $\overline{\theta}' > \overline{\theta}$. Then if the rate of re-running is higher for winners than for losers ($p_w > p_l$), there will be quality-based incumbency advantage: re-running candidates are stronger than replacement candidates, and the winning party makes more use of re-running candidates. Conversely, there would be quality-based incumbency disadvantage if candidates who re-run are weaker than average (i.e. if $\overline{\theta}' < \overline{\theta}$) or if the rate of re-running is higher for losers ($p_w < p_l$).

For example, it may be that only marginal winners and losers who are above average consider running again, and that all of the marginal winners in this category run again while only half of the marginal losers in this category (selected at random) do so. If replacement candidates have the same average quality as marginal winners and losers, then there will be a quality-based incumbency advantage: re-running candidates are stronger than replacement candidates, but the incumbent party uses a higher proportion of re-running candidates and thus its candidates are stronger on average.

3b: Differential replacement rates and changes in the candidate pool

Marginal candidates who re-run could also be stronger or weaker than replacement candidates due to changes in the candidate pool. Suppose that at the time of the marginal election (time t), marginal winners and losers are typical of the pool of candidates who enter elections in that period. (In the next section I examine conditions under which this is true.) A random subset of these marginal winners and losers re-run at time $t + 1$. Between time t and time $t + 1$, however, the pool of candidates who enter elections gets weaker or stronger, such that replacement candidates are stronger or weaker than replacement candidates; if there is also a difference in the proportion of winners and losers who re-run, there will be quality-based incumbency advantage or disadvantage.

Why might there be changes in the quality of the candidate pool from one election to the next? It may be that politics becomes a more or less attractive pursuit for talented people over time, resulting in a rise or fall in the quality of the candidate pool; alternatively, the process by which candidates are selected may become better or worse at identifying good types (e.g. because the "smoke-filled room" is replaced by a primary election). If changes of this kind

take place over many election cycles it could explain a finding of incumbency advantage or disadvantage.

A slightly different story is that the electorate's taste in candidates may change from one election to the next, and the candidate pool might change in response. Suppose, for example, that between the election at time t and the election at time $t + 1$ an economic crisis occurs, such that voters prioritize economic management ability (a trait to which they had previously not paid much attention); in anticipation of this change in tastes, a number of candidates with stronger economic management ability enter the candidate pool at time $t + 1$. From the perspective of voters' tastes at time $t+1$, the pool of candidates at time $t+1$ is stronger than the marginal winners and losers from time t who return at time $t + 1$; if marginal winners are more likely to run again, this will lead to quality-based incumbency disadvantage. Changes in taste may therefore reduce incumbency advantage overall or even produce incumbency disadvantage, particularly during periods when political priorities shift dramatically.

3c: Differential replacement rates and selection into marginality

Marginal candidates who re-run could also be stronger or weaker than replacement candidates due to a phenomenon I call "selection into marginality". Suppose that replacement candidates in each period are drawn from the same time-invariant candidate pool. The average quality of re-running candidates could be different from that of replacement candidates (even without selection into re-running or changes in the candidate pool) if marginal winners and losers differ in average quality from the candidate pool. Along with differential replacement rates this would yield quality-based incumbency advantage or disadvantage.

Why would marginal winners and losers be stronger or weaker on average than the candidate pool? This is a subtle point, so I investigate it in more depth in the next section. The key is that close elections are not a random sample of all elections; rather, given my assumptions about voter preferences they are elections when two candidates with similar quality face each other. If both candidates are drawn from the same candidate pool (as in an open-seat election), then types that are common in the candidate pool will be even more common in close elections, which could mean that marginal candidates are stronger or weaker than the candidate pool. If

one candidate is drawn from the candidate pool and the other is stronger than the candidate pool (e.g. because she is an incumbent), then close elections will disproportionately feature candidates who are stronger than the candidate pool. Finally, the conditions for balance on quality cannot simultaneously be met for open-seat elections and elections involving incumbents. I prove and generalize all of these points in the next section.

4 Selection into marginality

The previous section highlighted the three mechanisms that can produce quality-based incumbency advantage or disadvantage. In this section I elaborate on the last mechanism considered: marginal incumbent-party candidates can be stronger or weaker on average than their opponents if marginal winners and losers are replaced at different rates, and if marginal winners and losers are stronger or weaker on average than the candidate pool – a phenomenon I call "selection into marginality". I show that selection into marginality contributes to incumbency effects under very general conditions.

Building on the setup of the previous analysis, I begin with assumptions that shut down other possible sources of incumbency advantage or disadvantage. I assume that incumbency status *per se* is irrelevant to election outcomes, such that party a's vote share is given by $V_t = \frac{1}{2} + \theta_{at} - \theta_{bt}$. I assume that marginal candidates who re-run (for both the winning and the losing side) are no weaker or stronger than the pool of marginal winners and losers, i.e. $\overline{\theta}_w^r = \overline{\theta}_l^r = \overline{\theta}_w$. I assume that there is a time-invariant candidate pool from which all entering candidates are drawn at random; thus $\overline{\theta}_w^e = \overline{\theta}_l^e = \overline{\theta}$, where $\overline{\theta}$ indicates the average quality of the candidate pool. (Thus incumbency has no direct electoral relevance and there is no selection into re-running or differential selection into entry.) Finally, I assume that $p_w \neq p_l$, i.e. that replacement rates differ for winners and losers.[12] Under these assumptions, $\overline{\theta}_I - \overline{\theta}_C = (p_w - p_l)(\overline{\theta}_w - \overline{\theta})$; thus the necessary and sufficient condition for balance on quality is that $\overline{\theta}_w = \overline{\theta}$, i.e. marginal winners (and losers) have the same average quality as the candidate pool.

I begin by characterizing the distribution of quality among marginal winners and losers,

---

[12]One situation in which these conditions would be met is where all marginal winners run for re-election and all marginal losers are replaced, i.e. $p_w = 1, p_l = 0$.

given the distribution of quality among candidates:

Lemma 1 *Quality of marginal candidates: Suppose that candidates for party a and b have quality measures θ_a and θ_b independently drawn from $f_{\theta_a}(x)$ and $f_{\theta_b}(x)$ respectively. The posterior density of quality among winners and losers evaluated where party a wins half the votes is proportional to $f_{\theta_a}(x)f_{\theta_b}(x)$.*

Proof Denote by $m \equiv V_t - 1/2$ the vote margin between a and b, and denote by $f_{\theta_a|m}(x|y)$ the density of θ_a evaluated at x conditional on the margin $m = y$, with $f_{\theta_b|m}(x|y)$ similarly defined. Using Bayes Rule, we have

$$f_{\theta_a|m}(x|y) = \frac{f_{m|\theta_a}(y|x)f_{\theta_a}(x)}{C}. \tag{10}$$

Then, dropping the constant and evaluating this where $y = 0$, we have

$$f_{\theta_a|m}(x|0) \quad \propto \quad f_{m|\theta_a}(0|x)f_{\theta_a}(x) \tag{11}$$

$$= \quad f_{\theta_b|\theta_a}(x|x)f_{\theta_a}(x) \tag{12}$$

$$= \quad f_{\theta_b}(x)f_{\theta_a}(x) \tag{13}$$

where going from line 11 to line 12 we use the assumption that $m = \theta_a - \theta_b$, so that evaluating the density at $m = 0$ is equivalent to evaluating it at $\theta_a = \theta_b = x$; going from line 12 to line 13 we use the assumed independence of θ_a and θ_b. By the same argument we can work out that $f_{\theta_b|m}(x|0) \propto f_{\theta_a}(x)f_{\theta_b}(x)$. Thus the distribution of quality among marginal winners and losers is proportional to $f_{\theta_a}(x)f_{\theta_b}(x)$. ∎

It should make intuitive sense that the distribution of quality among marginal winners and losers is proportional to the product of the *ex ante* quality distributions for the two candidates: under our assumptions a close election occurs when two candidates have the same type, so the probability of observing a marginal winner or loser of a given type is proportional to the probability of drawing that type from both distributions, which given independence is just the product of the probabilities.

Now consider the situation where both candidates in the first election are drawn from the same candidate pool, which I will refer to as an open-seat contest. Under what conditions will

there be balance on quality?

Proposition 1 *Open-seat contests: Suppose that the distribution of quality for all entering candidates is described by $g(\theta)$ and that the replacement rate differs for marginal winners and losers. A sufficient (but not necessary) condition for balance on quality following open-seat contests is symmetry in $g(\theta)$.*

Proof By Lemma 1, the distribution of quality among marginal winners of open-seat contests is proportional to $g^2(\theta)$. Note that the square of a symmetric function with point of symmetry μ is also a symmetric function with point of symmetry μ. (Symmetry about μ implies that $g(\mu + c) = g(\mu - c)$ for all c; thus $g^2(\mu + c) = g^2(\mu - c)$ for all c, meaning that $g^2(\theta)$ is also symmetric about μ.) Also note that the expectation of a random variable with a distribution that is symmetric around point μ is μ. Thus if $g(\theta)$ is symmetric around point μ, then the average quality of marginal winners and losers of open-seat contests is μ, and thus the same as the average quality in the candidate pool.

To show that symmetry is not necessary, it suffices to provide a case in which quality is distributed asymmetrically but marginal candidates are equally strong as the candidate pool on average. Suppose that quality takes on three values, -1, 0, and 1, that the probability of the low and high type are given by p and q, and that a close election happens only when two candidates of the same type face each other. Then the average quality of the candidate pool is given by $q - p$ and (using Lemma 1) the average quality of marginal candidates is given by $\frac{q^2 - p^2}{p^2 + (1-p-q)^2 + q^2}$. The two are equal when $p = q$ (indicating a symmetric distribution of candidate quality), but there are also asymmetric solutions that can be recovered numerically. Figure 6 in the Appendix shows all solutions, both symmetric and asymmetric. ∎

Figure 2 shows an example in which asymmetry in the distribution of quality in the candidate pool leads marginal winners and losers to be weaker than the candidate pool on average. In the distribution of quality in the candidate pool (labeled $g(\theta)$), weaker types are more prevalent than stronger types. This skew is accentuated in the distribution of quality among marginal winners of open-seat elections (labeled $g^2(\theta)$). Given this distribution of candidate types, then, marginal incumbents elected in open-seat contests will be weaker than the challengers they face,

Figure 2: The distribution of candidate types among marginal winners and losers of open-seat elections, given an asymmetric distribution of candidate types
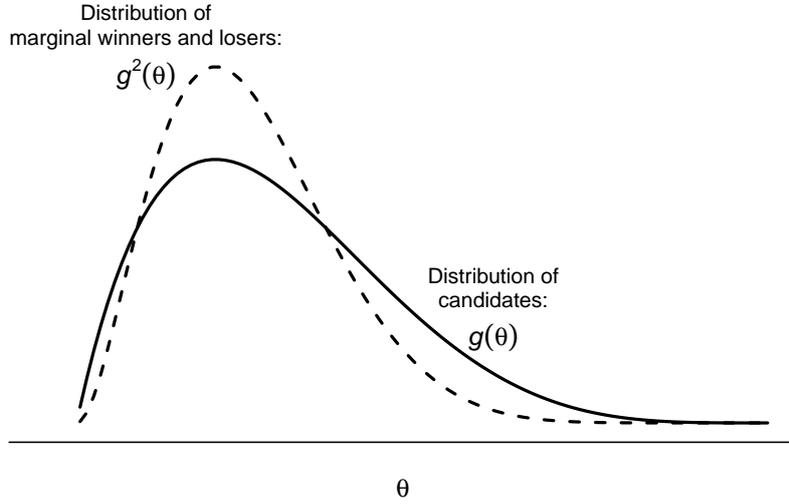
NOTE: The solid curve indicates the distribution of quality in the candidate pool; the dashed curve indicates the corresponding distribution of quality in the pool of candidates who marginally win or lose open-seat elections.

producing quality-based incumbency disadvantage.

Now we consider elections in which one candidate is an incumbent who previously won an open-seat election:

Proposition 2 *Balance on quality and electoral selection: Suppose all entering candidates are drawn randomly and independently from $g(\theta)$, with corresponding cumulative density function $G(\theta)$. Then for any non-degenerate $g(\theta)$ there can be balance on quality following open-seat elections or following elections involving winners of open-seat elections but not both.*

Proof Recall that the necessary condition for balance on quality (assuming differential replacement rates, no selection into re-running, and no differential selection into entry) is that marginal winners and losers must have the same average quality as the candidate pool. If balance on quality is to be achieved both following open-seat contests (case 1) and following elections involving the winners of open-seat contests (case 2), this means that the average quality of marginal winners/losers in case 1 must be the same as that in case 2. Lemma 1 proved that the distribution of quality among marginal winners of open-seat contests (case 1) is proportional to $g^2(\theta)$. Using Bayes' Rule, the *ex ante* quality of winners of open-seat elections is proportional to $G(\theta)g(\theta)$, so
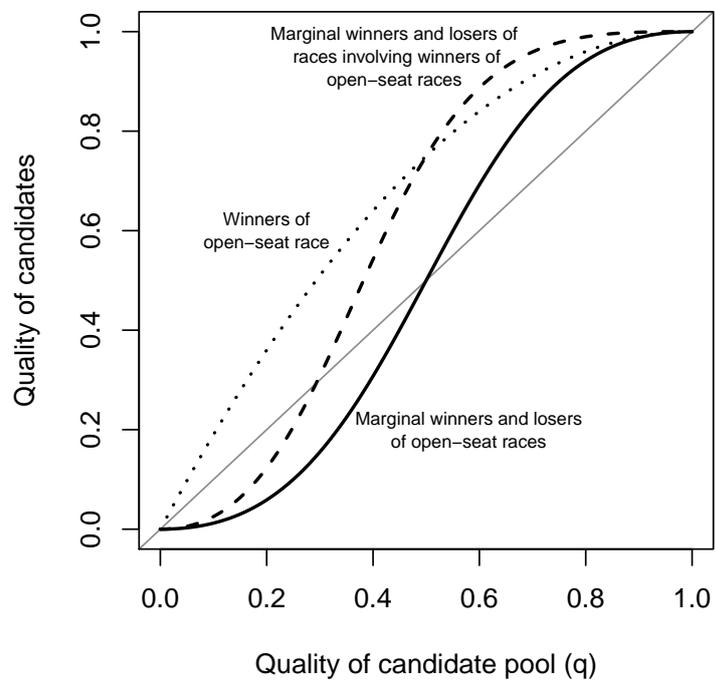
using Lemma 1 the quality of marginal winners of elections involving the winners of open-seat elections (case 2) is proportional to $G(\theta)g^2(\theta)$. By definition $G(\theta)$ is weakly increasing, which means that the distribution of quality in case 2 places a higher weight on higher values of θ and a lower weight on lower values of θ relative to the distribution of quality in case 1, which implies that average quality is lower in case 1 than in case 2, which means that the necessary condition for balance on quality cannot be simultaneously met in both cases. ∎

To illustrate the foregoing points about selection into marginality I focus on the binary case, where candidates are either strong or weak, with q denoting the probability of a strong candidate being drawn from the candidate pool; I assume that close elections happen only when two candidates of the same type face each other. The solid black curve in Figure 3 shows the quality of marginal winners and losers of open-seat elections as a function of q. (This can be calculated as $q^2/\left(q^2 + (1-q)^2\right)$ using the logic of Lemma 1.) When the candidate pool is evenly balanced between good types and bad types, marginal winners and losers of open-seat elections are also evenly balanced between good types and bad types (and thus the solid black curve intersects the 45 degree line at $q = 1/2$). When good types are more common than bad types, marginal winners and losers of open-seat elections are stronger than the candidate pool; the reverse is true when good types are less common than bad types. The figure shows that when quality is binary, symmetry (i.e. $q = 1/2$) is not just a sufficient condition but also a necessary condition for balance on quality following an open-seat election.

The dotted curve in Figure 3 (labeled "Winners of open-seat races") shows the distribution of quality for *all* winners of open-seat elections (not just marginal winners); due to electoral selection, these candidates are stronger on average than the candidate pool for any q.[13] By conditioning on a close election we eliminate *electoral selection*, i.e. the tendency of stronger candidates to win elections, as a cause of quality differences between open-seat winners and the candidate pool. But unless $q = 1/2$ (or $q \in \{0, 1\}$, which are degenerate cases), quality differences remain between marginal candidates and the candidate pool due to selection into marginality.

---

[13]Given quality q of candidates L and R, the *ex ante* quality of the winner (not conditional on a close election) is $2q\left(1 - \frac{q}{2}\right)$.

Figure 3: The quality of winning candidates as a function of the quality of the candidate pool and the type of contest (binary case)

Finally, the dashed curve in Figure 3 shows the quality of marginal winners of elections involving the winners of open-seat contests; note that these candidates are stronger *ex ante* for any value of q than the marginal winners of open-seat contests. Thus if $q = 1/2$, such that there is balance on quality following open-seat races, there will be a quality-based incumbency advantage whenever the marginal race involves an incumbent; there is a lower q such that we have balance on quality for the latter case,[14] but then there is quality-based incumbency *disadvantage* following open-seat races.

5   Discussion and implications

This paper has examined mechanisms by which differences in candidate quality could contribute to findings of incumbency advantage or disadvantage. In this section I consider how this analysis relates to existing empirical work on incumbency advantage, and how some of the mechanisms discussed above could be further tested.

Testing quality-based explanations

If quality can be observed, then there is a straightforward way to test quality-based explanations for findings of incumbency effects: compare the quality of marginal incumbent-party candidates and their opponents. Put differently, one can test these explanations by testing the effect of a victory by party a at time t on the quality difference between party a's candidate and party b's candidate at time $t + 1$.

Hall and Snyder (2015) use an RDD to conduct just such a test in the context of U.S. elections. Their measure of a candidate's quality is the candidate's prior officeholding experience; thus for candidates for U.S. Congress the measure is whether the candidate has held office at the state level or in the U.S. Senate. As it turns out, they do not find a statistically significant effect in any of the settings they study, though the point estimates are all positive (i.e. suggesting that marginal incumbent-party candidates are stronger than their opponents) and similar in magnitude. These results are presented specifically as a test of scare-off, but it should be clear that they offer a test of any quality-based explanation.

---

[14]Using Lemma 1, it is easily shown that this is the q such that the quality of open-seat winners is 1/2.

21

One conclusion to draw from Hall and Snyder (2015) is that quality differences do not explain incumbency effects in the U.S. settings they study. Another interpretation is that prior officeholding is a poor measure of quality: at best it is a noisy proxy for many of the characteristics voters care about, especially given that the average CEO who runs for Congress is probably more appealing to voters than the average state senator who runs for Congress. Future work could build on Hall and Snyder (2015) by applying the same kind of test using more fine-grained measures of candidate quality, perhaps following Stone et al. (2010), and studying additional settings.

Future work could also test the theoretical finding that electoral selection contributes to incumbency advantage through selection into marginality. Proposition 2 implies that incumbency effects should be larger (all else equal) in cases where the marginal election involves an incumbent than in cases where the marginal election is an open-seat election. The difference in the two incumbency effects should be smaller when incumbency is valuable (i.e. $\gamma > 0$) and/or incumbency deters strong challengers, because this would reduce the average quality of incumbents; to the extent that retirement patterns differ following close open-seat elections and other close elections this also complicates the comparison.

Electoral selection as an explanation for cross-national variation in incumbency advantage

Lemma 1 implies that, when we consider elections involving incumbents, the quality of marginal winners is increasing in the quality of incumbents; this implies that quality-based incumbency advantage will depend on how strong incumbents are relative to the candidate pool. Throughout the analysis above I assumed that the better candidate always won (because there is always partisan balance and quality is perfectly observed), but one could easily extend the model to incorporate partisan imbalance, partisan tides, or random shocks to observed quality, such that the better candidate sometimes loses;[15] the result would be that the average quality of incumbents would drop.[16] This suggests that in settings where elections are more effective at

---

[15]The Appendix extends the main results to the case of partisan imbalance.

[16]For example, assuming binary types and no noise or partisan imbalance a weak incumbent can only be elected if both candidates are weak; with noise/partisan imbalance a weak incumbent must be elected in that case but might also be elected when a weak type faces a strong type.

selecting the better candidate we should see higher incumbency advantage.

Viewed in this light, a quality-based account might help explain why incumbency advantage appears to be larger in developed democracies than in developing democracies, based on the recent studies cited in the introduction: if elections in developed democracies are more effective at selecting the stronger candidate, whether because of better information, stricter media scrutiny, or lower rates of fraud, this could explain why we see larger incumbency advantage in these elections.

By the same token, a quality-based account may also help explain why incumbency advantage varies in developed democracies over time and across countries. All else equal, elections are weaker instruments of selection (and thus quality-based incumbency effects should be smaller) when fewer voters are nearly indifferent between the candidates on partisan grounds. Consistent with this, Eggers and Spirling (2015) show larger incumbency effects in U.K. election contests when the two main parties are more ideologically proximate; Jacobson (2015) also shows that incumbency advantage in the U.S. tracked the degree of partisan attachment in the electorate over the post-World War II period.

Incumbency disadvantage and term limits

Klašnja and Titiunik (2016) report findings of incumbency disadvantage from Brazil and several other Latin American countries where term limits apply. They interpret these findings as a result of "unfulfilled accountability" in systems with weak parties: term-limited incumbents shirk from their responsibilities (their parties being unable to prevent them from doing so), and voters punish the incumbent's party in the next election (even though the incumbent is not able to run). Similarly, Fowler and Hall (2014) find suggestive evidence of a party incumbency disadvantage in cases where the incumbents in U.S. state legislatures do not run again due to term limits.[17]

The analysis in this paper offers a more straightforward explanation for incumbency disadvantage in the presence of term limits: candidates who replace term-limited incumbents are

---

[17]Fowler and Hall (2014)'s analysis uses previous close elections involving incumbents as a source of exogenous variation in whether term limits bind in a given election or not. Fowler and Hall (2014, p. 506) speculate that the partisan incumbency effect is negative because voters seek to balance power, because they dislike parties, or because they prefer change.

weaker than their opponents. When a close election occurs at time t and the winner is prevented from running again, the rate of re-running for the winning party is 0. If some marginal losers run again, and all replacement candidates are drawn from the same candidate pool, then the incumbent party will be at a disadvantage if re-running losers are stronger than the candidate pool, which would be the case due to selection into re-running or selection into marginality. For example, it could be that marginal candidates who re-run are stronger than the average marginal winners and losers, and because only the losing party has re-running candidates they are at an advantage. Or, perhaps close elections at time t disproportionately involve strong candidates (due to selection into marginality), such that re-running marginal losers are stronger than the candidate pool and the losing party is at an advantage. This simple explanation could account for both the findings of incumbency disadvantage in term-limited contexts in Latin America (as reviewed by Klašnja and Titiunik (2016)) and for the finding of negative partisan incumbency advantage in Fowler and Hall (2014).

Explaining variation in incumbency effects in Romania

Klašnja (2015b) presents evidence of incumbency disadvantage in Romanian mayoral elections. Intriguingly, he shows that this disadvantage is larger in smaller municipalities, where the salary of the mayor is statutorily set at a lower level and (he argues) the quality of the candidate pool is lower. Although he explains this with reference to an accountability model, the analysis above suggests at least two quality-based accounts that are simpler. One explanation is differential selection into re-running: when the mayor's pay is lower, the incentives for high-quality incumbents to seek better opportunities are higher; this could lead to the best incumbents leaving (perhaps for higher office), and the incumbents who re-run having lower average quality than the candidate pool, resulting in incumbency disadvantage. Another explanation is selection into marginality: when the candidate pool is skewed towards low-quality candidates, marginal winners may be weaker than the average candidate, resulting in incumbency disadvantage.

Implications for estimating the personal incumbency advantage

My analysis of selection into marginality is related to Erikson and Titiunik (2015), who establish conditions under which estimates of party incumbency effects can be used to uncover the personal incumbency advantage, i.e. the advantage to a candidate of running as the incumbent rather than as a non-incumbent. They argue that the marginal winners and losers of open-seat elections are likely to be of average quality, in which case we can rule out electoral selection as an explanation for the incumbent party's subsequent success, but they note that the marginal winners of elections involving incumbents could be stronger or weaker than the average candidate, which makes analysis complicated. They therefore focus on marginal open-seat elections, discarding other marginal elections in their main analysis. The analysis above is related in two main ways. First, Proposition 1 shows that symmetry in the distribution of candidate quality is a sufficient condition for their assumption that the marginal winners of open-seat contests have the same average quality as the candidate pool. Second, Proposition 2 shows why, if their assumption is met for open-seat elections, it is unlikely to be met for other types of marginal contests. The analysis becomes more complicated when we allow for incumbency to be electorally valuable ($\gamma > 0$) or allow for scare-off in the challenger pool; both adjustments would tend to reduce the quality of winners of marginal elections that involve incumbents. In future work my framework could be extended to examine those issues more carefully.

My analysis is also related to Fowler and Hall (2014), who use close elections and term limits in state legislatures in order to decompose party incumbency advantage into personal and partisan components (i.e. the benefit to an individual of being the incumbent candidate, and the benefit to a party of being the incumbent party). One of the assumptions behind their analysis is that incumbents who do not run for election are replaced by candidates of equal average quality; in terms of my framework, this corresponds to the assumption that $\overline{\theta}_I = \overline{\theta}_w$. (My analysis has explored conditions under which this assumption may not be met; the online appendix of Fowler and Hall (2014) considers violations of this assumption and argues that the consequences are likely to be small.) Supposing that this assumption holds, any quality-based incumbency effects arise due to differences between the average quality of marginal losers ($\overline{\theta}_l$) and marginal challenger-party candidates ($\overline{\theta}_C$). The conventional interpretation is that

differences between $\overline{\theta}_l$ and $\overline{\theta}_C$ arise due to scare-off and thus belong in the personal incumbency advantage: the candidate for the non-incumbent party is weak because the incumbent is running. My analysis has shown how $\overline{\theta}_l$ and $\overline{\theta}_C$ can differ due to selection into re-running, changes in the candidate pool, and selection into marginality; this is true even when incumbency has no value. In that sense part of the difference between $\overline{\theta}_l$ and $\overline{\theta}_C$ may belong in the partisan incumbency advantage: it is an advantage (or disadvantage) that accrues to the winning party, but not due to the incumbency status of that party's candidate.

Quality-based explanations of candidate-level incumbency effects

As noted in the introduction, in addition to the party-based approach of Lee (2008), one can use RDD to study incumbency effects at the candidate level as advocated by De Magalhaes (2015) (see also Trounstine, 2011). This approach asks how an individual's narrow electoral success or failure in an election at time t affects his or her probability of running for office at time $t + 1$ and/or winning office at time $t + 1$. Because the candidate-level approach focuses on individuals rather than parties, it may seem that the quality-based explanations I examine do not apply: after all, the mechanisms I explore depend on marginal winners and losers being replaced by other candidates, while the candidate-level approach keeps its focus on the marginal winners and losers themselves. In fact all of the processes I examine would affect candidate-level estimates. Suppose, for example, that all marginal winners run for re-election, but half of marginal losers are replaced. Then selection into re-running and selection into entry will both affect the candidate-level estimates: if it is the strong marginal losers who run again, this will tend to reduce winners' success rates and increase losers' success rates and thus reduce the candidate-level incumbency effect (just as it would reduce the party incumbency effect); if the candidates who replace departing marginal losers are weaker than average (perhaps because of scare-off), then this would tend to increase winners' success rates and thus increase the candidate-level incumbency effect. Similar arguments could be made if we consider a situation where some marginal winners are replaced. The clearest difference between candidate-level and party-level effects is that the rate of re-running directly enters the estimate for candidate-level effects, rather than simply weighting the re-running and entering candidates as in the party-

level version. Otherwise, all of the mechanisms I described (differential selection into re-running, differential selection into entry, changes in the candidate pool, selection into re-running, selection into marginality) have just as much scope for explaining candidate-level estimates as they do party-level estimates.

6   Conclusion

This paper has highlighted mechanisms by which quality differences between marginal incumbent-party candidates and their opponents could explain findings of incumbency advantage and disadvantage. When some of the winners and losers of marginal elections do not run for re-election, there is scope for such differences to arise through three distinct mechanisms. In some settings, these mechanisms suggest promising explanations for findings of incumbency advantage or disadvantage; in others, quality-based explanations may play a very small role. In addition to suggesting explanations for empirical findings, my analysis may also further clarify what assumptions are necessary to interpret findings of incumbency effects as reflecting officeholding benefits, "scare-off", or some other specific mechanism.

Along with other recent theoretical work (e.g. Meirowitz, 2008; Ashworth and Bueno de Mesquita, 2008; Caselli et al., 2013; Muthoo and Shepsle, 2014), this paper has added to the list of possible mechanisms that could contribute to incumbency effects. To an empiricist, the growing list of explanations for incumbency can be viewed as a mixed blessing. On the one hand, theoretical work that clarifies these mechanisms suggests possible explanations for puzzling findings and points toward new tests to assess those explanations. On the other hand, this work highlights the difficulty of drawing policy implications from a simple finding of incumbency advantage or disadvantage. As explored in this paper, a finding of large incumbency advantage could be an indication that electoral selection is working well; a finding of large incumbency disadvantage could be an indication that parties and voters successfully deter weak losers from trying again. In other words, this paper and related theoretical work on incumbency effects suggests that we usually cannot learn very much simply by estimating the electoral effects of incumbency; it is only in conjunction with additional analysis that we can draw conclusions about how electoral competition works and what it means for voter welfare.

References

Alford, John R and John R Hibbing. 1981. "Increased incumbency advantage in the House." *The Journal of Politics* 43(04):1042–1061.

Ansolabehere, Stephen, James M Snyder and Charles Stewart. 2000. "Old voters, new voters, and the personal vote: Using redistricting to measure the incumbency advantage." *American Journal of Political Science* 44(1):17–34.

Ansolabehere, Stephen and James 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.

Ariga, Kenichi. 2010. Entrenched Incumbents, Irresponsible Parties? Comparative Analysis of Incumbency Advantage Across Different Electoral Systems PhD thesis University of Michigan.

Ariga, Kenichi. 2015. "Incumbency Disadvantage under Electoral Rules with Intraparty Competition: Evidence from Japan." *Journal of Politics* 77(3):xxx–XXX.

Ashworth, Scott and Ethan Bueno de Mesquita. 2008. "Electoral selection, strategic challenger entry, and the incumbency advantage." *Journal of Politics* 70(4):1006–1025.

Bevia, Carmen and Humberto Llavador. 2009. "The informational value of incumbency." *Journal of Public Economic Theory* 11(5):773–796.

Butler, Daniel Mark. 2009. "A regression discontinuity design analysis of the incumbency advantage and tenure in the US House." *Electoral Studies* 28(1):123–128.

Cain, Bruce E, John A Ferejohn and Morris P Fiorina. 1987. "The personal vote: Constituency service and electoral independence.".

Caselli, Francesco, Tom Cunningham, Massimo Morelli and Inés Moreno Barreda. 2013. "The Incumbency Effects of Signalling." *Economica* .

Cox, Gary W and Jonathan N Katz. 1996. "Why did the incumbency advantage in US House elections grow?" *American Journal of Political Science* pp. 478–497.

Cox, Gary W and Scott Morgenstern. 1993. "The increasing advantage of incumbency in the US states." *Legislative Studies Quarterly* pp. 495–514.

De Magalhaes, Leandro. 2015. "Incumbency effects in a comparative perspective: Evidence from Brazilian mayoral elections." *Political Analysis* 23(1):113–126.

Eggers, Andrew C, Anthony Fowler, Jens Hainmueller, Andrew B Hall and James M Snyder. 2015. "On the validity of the regression discontinuity design for estimating electoral effects: New evidence from over 40,000 close races." *American Journal of Political Science* 59(1):259–274.

Eggers, Andrew C. and Arhur Spirling. 2015. "Incumbency Effects and the Strength of Party Preferences: Evidence from Multiparty Elections in the United Kingdom?".

Erikson, Robert S. 1971. "The advantage of incumbency in congressional elections." *Polity* 3(3):395–405.

Erikson, Robert S and Rocío Titiunik. 2015. "Using Regression Discontinuity to Uncover the Personal Incumbency Advantage." *Quarterly Journal of Political Science* 10(1):101–119.

Fiorina, Morris P. 1977. "The case of the vanishing marginals: The bureaucracy did it." *The American Political Science Review* 71(1):177–181.

Fouirnaies, Alexander and Andrew B Hall. 2014. "The Financial Incumbency Advantage: Causes and Consequences." *The Journal of Politics* 76(03):711–724.

Fowler, Anthony. 2015*a*. A Bayesian Explanation for Incumbency Advantage. In *111th APSA Annual Conference, San Francisco.*

Fowler, Anthony. 2015*b*. "Do Elections Select Better Representatives?".

Fowler, Anthony and Andrew B Hall. 2014. "Disentangling the Personal and Partisan Incumbency Advantages: Evidence from Close Elections and Term Limits." *Quarterly Journal of Political Science* 9(4):501–531.

Gelman, Andrew and Gary King. 1990. "Estimating incumbency advantage without bias." *American Journal of Political Science* pp. 1142–1164.

Hainmueller, J. and HL Kern. 2008. "Incumbency as a source of spillover effects in mixed electoral systems: Evidence from a regression-discontinuity design." *Electoral Studies* 27(2):213–227.

Hall, Andrew B and James M Snyder. 2015. "How Much of the Incumbency Advantage is Due to Scare-Off?" *Political Science Research and Methods* 3(03):493–514.

Horiuchi, Yusaku and Andrew Leigh. 2009. "Estimating Incumbency Advantage: Evidence from Multiple Natural Experiments." *Unpublished paper* .

Jacobson, Gary C. 2015. "Its Nothing Personal: The Decline of the Incumbency Advantage in US House Elections." *The Journal of Politics* 77(3):861–873.

Katz, Jonathan N and Gary King. 1999. "A statistical model for multiparty electoral data." *American Political Science Review* pp. 15–32.

Kendall, Chad and Marie Rekkas. 2012. "Incumbency advantages in the Canadian Parliament." *Canadian Journal of Economics/Revue canadienne d'économique* 45(4):1560–1585.

King, Gary. 1991. "Constituency service and incumbency advantage." *British Journal of Political Science* 21(01):119–128.

Klašnja, Marko. 2015*a*. "Corruption and the Incumbency Disadvantage: Theory and Evidence.".

Klašnja, Marko. 2015*b*. "Corruption and the Incumbency Disadvantage: Theory and Evidence." *The Journal of Politics* 77(4):928–942.

Klašnja, Marko. 2015*c*. "Increasing Rents and Incumbency Disadvantage." *Journal of Theoretical Politics* .

Klašnja, Marko and Rocıo Titiunik. 2016. "The Incumbency Curse: Weak Parties, Term Limits, and Unfulfilled Accountability." *American Political Science Review* .

Lee, David S. 2008. "Randomized Experiments From Non-Random Selection in U.S. House Elections." *Journal of Econometrics* 142(2):675–697.

Levitt, Steven D and Catherine D Wolfram. 1997. "Decomposing the sources of incumbency advantage in the US House." *Legislative Studies Quarterly* pp. 45–60.

Mayhew, David R. 1974. "Congressional elections: The case of the vanishing marginals." *Polity* pp. 295–317.

McKelvey, Richard D and Raymond Riezman. 1992. "Seniority in legislatures." *American Political Science Review* 86(04):951–965.

Meirowitz, Adam. 2008. "Electoral contests, incumbency advantages, and campaign finance." *The Journal of Politics* 70(03):681–699.

Muthoo, Abhinay and Kenneth A Shepsle. 2014. "Seniority and incumbency in legislatures." *Economics & Politics* 26(1):13–37.

Stone, Walter J, Sarah A Fulton, Cherie D Maestas and L Sandy Maisel. 2010. "Incumbency reconsidered: prospects, strategic retirement, and incumbent quality in US House elections." *The Journal of Politics* 72(01):178–190.

Trounstine, Jessica. 2011. "Evidence of a local incumbency advantage." *Legislative Studies Quarterly* 36(2):255–280.

Uppal, Yogesh. 2009. "The disadvantaged incumbents: estimating incumbency effects in Indian state legislatures." *Public Choice* 138(1-2):9–27.

A Schematic illustration for the three mechanisms of quality-based incumbency effects

Figure 4 illustrates the three mechanisms of quality-based incumbency effects. In each diagram, the average quality of marginal winners and losers who re-run ($\overline{\theta}_w^r$ and $\overline{\theta}_l^r$) is depicted on the right vertical axis, the average quality of candidates who replace departing marginal winners and losers ($\overline{\theta}_w^e$ and $\overline{\theta}_l^e$) is depicted on the left vertical axis, and the proportion of marginal winners and losers who re-run (p_w and p_l) is depicted on the horizontal axis; the average quality of marginal incumbent-party candidates is given by height of the black dot located at a horizontal distance of p_w along the line connecting $\overline{\theta}_w^r$ and $\overline{\theta}_w^e$, and the average quality of marginal incumbent-party candidates is given by the height of the gray dot at a horizontal distance of p_l along the line connecting $\overline{\theta}_l^r$ and $\overline{\theta}_l^e$. If there is quality-based incumbency advantage the black dot is higher than the gray dot; if there is quality-based incumbency disadvantage the reverse is true. In "differential selection into re-running" the dots differ in height because $\overline{\theta}_w^r \neq \overline{\theta}_l^r$ (while $\overline{\theta}_w^e = \overline{\theta}_l^e$ and $p_w = p_l$); in "differential selection into entry" the dots differ in height because $\overline{\theta}_w^e \neq \overline{\theta}_l^e$ (while $\overline{\theta}_w^r = \overline{\theta}_l^r$ and $p_w = p_l$); in "differential replacement rates" the dots differ in height because $p_w \neq p_l$ and $\overline{\theta}_w^r \neq \overline{\theta}_w^e$ (while $\overline{\theta}_w^r = \overline{\theta}_l^r$ and $\overline{\theta}_w^e = \overline{\theta}_l^e$). One can confirm geometrically that quality-based incumbency advantage or disadvantage cannot occur unless at least one of these mechanisms operates.

The diagram in Figure 5 shows how the three mechanisms might combine to produce quality-based incumbency advantage. (The arguments could all be reversed to yield quality-based incumbency disadvantage.) In the setting described by the diagram, marginal winners are more likely to run for re-election than losers ($p_w > p_l$). Marginal winners who run for re-election are stronger on average than marginal winners who run for re-election ($\overline{\theta}_w^r > \overline{\theta}_l^r$), perhaps because serving in office weeds out the worst marginal winners and the best marginal losers find better opportunities in the private sector. Candidates who replace marginal losers are also somewhat weaker than candidates who replace marginal winners ($\overline{\theta}_l^e > \overline{\theta}_w^e$), perhaps because of scare-off that can be traced to the advantages of incumbency ($\gamma > 0$) or to differential selection into re-running (i.e. the fact that $\overline{\theta}_w^r > \overline{\theta}_l^r$). The net effect of these differences is that marginal

Figure 4: The three mechanisms for quality-based incumbency effects

Differential selection into re-running	Differential selection into entry	Differential rate of replacement
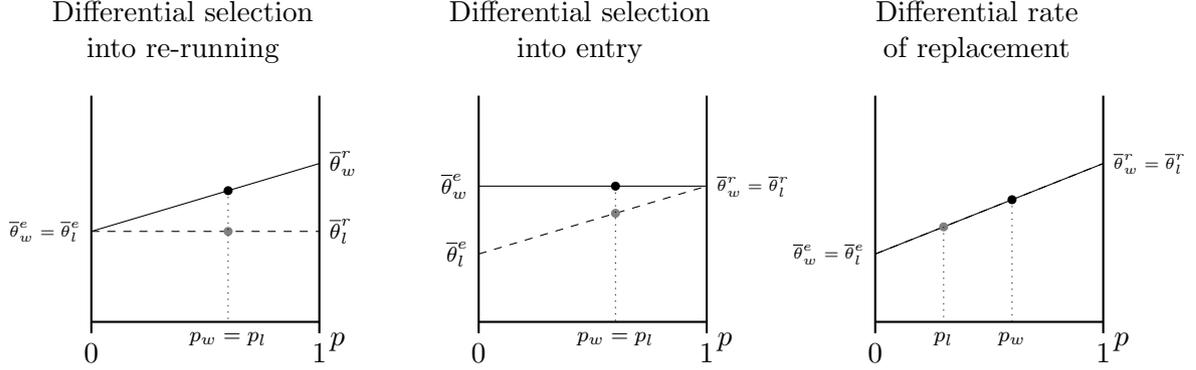
NOTE: The average quality of re-running candidates for the winning and losing parties ($\overline{\theta}_w^r$ and $\overline{\theta}_l^r$, respectively) is depicted on the right vertical axis; the average quality of entering (replacement) candidates for the winning and losing parties ($\overline{\theta}_w^e$ and $\overline{\theta}_l^e$, respectively) is depicted on the left vertical axis. The re-running rate for the winning and losing parties (p_w and p_l, respectively) is depicted on the horizontal axis. The average quality of marginal incumbent-party candidates and marginal challenger-party candidates is shown with a black dot and a gray dot, respectively.

Figure 5: An example that combines the three mechanisms
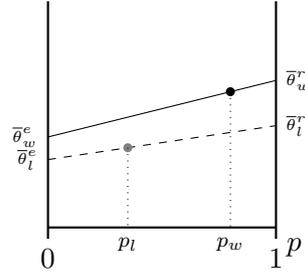
NOTE: See Figure 4.

incumbent-party candidates are stronger than marginal challenger-party candidates, i.e. there is a quality-based incumbency advantage.

B More general analysis with partisan imbalance

In the main analysis I assumed that the electorate was balanced between parties a and b, with partisan preferences $v_i(a) - v_i(b)$ distributed in the society uniformly on the interval $[-1/2, 1/2]$. Here I extend the main points for the more general case where partisan preferences $v_i(a) - v_i(b)$ in the electorate are distributed uniformly on the interval $[-1/2 + c, 1/2 + c]$.

Under this assumption, the vote share for party a is given by

$$V_t = 1/2 + c + \theta_{at} - \theta_{bt} + \gamma\big(I_{at} - I_{bt}\big). \tag{14}$$

Because party a receives the boost of c regardless of incumbency status, the analysis proceeds exactly as in the case with party balance, with the result that

$$\tau_{RDD} = 2\big(\overline{\theta}_I - \overline{\theta}_C + p_w\gamma\big). \tag{15}$$

Thus partisan imbalance does not change the way RDD-based incumbency effects relate to the quality of marginal incumbency-party candidates and their opponents.

We can now consider the main results for selection into marginality in the more general case where there may be partisan imbalance:

Lemma B.1 *Quality of marginal candidates (partisan imbalance): Suppose that candidates for party a and b have quality measures θ_a and θ_b independently drawn from $f_{\theta_a}(x)$ and $f_{\theta_b}(x)$ respectively. The posterior density of quality among winners and losers evaluated where party a wins half the votes is proportional to $f_{\theta_a}(x+c)f_{\theta_b}(x) + f_{\theta_a}(x)f_{\theta_b}(x-c)$.*

Proof Denote by $m \equiv V_t - 1/2$ the vote margin between a and b, and denote by $f_{\theta_a|m}(x|y)$ the density of θ_a evaluated at x conditional on the margin $m = y$, with $f_{\theta_b|m}(x|y)$ similarly defined. Using Bayes Rule, we have

$$f_{\theta_a|m}(x|y) = \frac{f_{m|\theta_a}(y|x)f_{\theta_a}(x)}{C}. \tag{16}$$

Then, dropping the constant and evaluating this where $y = 0$, we have

$$f_{\theta_a|m}(x|0) \quad \propto \quad f_{m|\theta_a}(0|x)f_{\theta_a}(x) \tag{17}$$

$$= \quad f_{\theta_b|\theta_a}(x+c|x)f_{\theta_a}(x) \tag{18}$$

$$= \quad f_{\theta_b}(x+c)f_{\theta_a}(x) \tag{19}$$

where going from line 17 to line 18 we use the assumption that $m = \theta_a + c - \theta_b$, so that evaluating the density at $m = 0$ is equivalent to evaluating it at $\theta_b = \theta_a + c$; going from line 18 to line 19

we use the assumed independence of θ_a and θ_b. By the same argument we can work out that $f_{\theta_b|m}(x|0) \propto f_{\theta_a}(x)f_{\theta_b}(x-c)$. The distribution of quality among marginal winners and losers is an even mixture of these two distributions. ∎

Now we consider the situation where both candidates in the first election are drawn from the same candidate pool. Under what conditions will there be balance on quality?

Proposition B.1 *Open-seat contests (partisan imbalance): Suppose that the distribution of quality for all entering candidates is described by $g(\theta)$ and that the replacement rate differs for marginal winners and losers. A sufficient condition for balance on quality following open-seat contests is symmetry in $g(\theta)$.*

Proof By Lemma B.1, the distribution of quality among marginal winners of open-seat contests is proportional to $g(\theta)\Big(g(\theta + c) + g(\theta - c)\Big)$. If $g(\theta)$ is a symmetric function with point of symmetry μ, then $g(\theta + c)$ is a symmetric function with point of symmetry $\mu + c$ and $g(\theta - c)$ is a symmetric function with point of symmetry $\mu - c$; the sum of the two is a symmetric function with point of symmetry μ. Finally, note that the product of two symmetric functions with common point of symmetry μ is another symmetric function with point of symmetry μ. (Given two functions $h_1(x)$ and $h_2(x)$, each with point of symmetry μ, symmetry implies that $h_1(\mu + d) = h_1(\mu - d)$ and $h_2(\mu + d) = h_2(\mu - d)$ for all d, and thus that $h_1(\mu - d)h_2(\mu - d) = h_1(\mu+d)h_2(\mu+d)$ for all d, which implies that $h_1(x)h_2(x)$ is symmetric with point of symmetry μ.) Finally note that the expectation of a random variable with a distribution that is symmetric around point μ is μ. Thus if $g(\theta)$ is symmetric around point μ, then the average quality of marginal winners and losers of open-seat contests is μ, which is the average quality in the candidate pool. ∎

Now we consider elections in which one candidate is an incumbent who previously won an open-seat election:

Proposition B.2 *Balance on quality and electoral selection: Suppose all entering candidates are drawn randomly and independently from $g(\theta)$, with corresponding cumulative density function $G(\theta)$. Then for any non-degenerate $g(\theta)$ there can be balance on quality following open-seat elections or following elections involving winners of open-seat elections but not both.*

Proof Recall that the necessary condition for balance on quality (assuming differential replacement rates, no selection into re-running, and no differential selection into entry) is that marginal winners and losers must have the same average quality as the candidate pool. If balance on quality is to be achieved both following open-seat contests (case 1) and following elections involving the winners of open-seat contests (case 2), this means that the average quality of marginal winners/losers in case 1 must be the same as that in case 2. Lemma B.1 proved that the distribution of quality among marginal winners of open-seat contests (case 1) is proportional to $g(\theta)\Big(g(\theta + c) + g(\theta - c)\Big)$. Using Bayes' Rule, the *ex ante* quality of winners of open-seat elections is proportional to $g(\theta)\Big(G(\theta - c) + G(\theta + c)\Big)$, and using this result and Lemma B.1 the quality of marginal winners of elections involving the winners of open-seat elections (case 2) is proportional to $g(\theta)\Big(G(\theta + c)g(\theta + c) + G(\theta - c)g(\theta - c)\Big)$. By definition $G(\cdot)$ is weakly increasing, which means that the distribution of quality in case 2 places a higher weight on higher values of θ and a lower weight on lower values of θ relative to the distribution of quality in case 1, which implies that average quality is lower in case 1 than in case 2, which means that the necessary condition for balance on quality cannot be simultaneously met in both cases. ∎

C   Non-necessity of symmetry

Given three types $\{-1, 0, 1\}$ that occur with probabilty $\{p, 1 - p - q, q\}$ in the candidate pool, and assuming that close elections occur when two candidates of the same type face each other, there is balance on quality following open-seat elections when

$$q - p = \frac{q^2 - p^2}{p^2 + (1 - p - q)^2 + q^2}, \tag{20}$$

which equates the average quality in the candidate pool (left) and the average quality of marginal winners of open-seat elections calculated following Lemma 1 (right).

Figure 6 shows all solutions to this equation, with p on the horizontal axis and q on the vertical axis. Points along the 45-degree line indicate symmetric solutions; points along the intersecting curve indicate asymmetric solutions. The four zones created by the intersection of the curve and the 45-degree line indicate distributions where marginal winners and losers

36

would be stronger or weaker than the candidate pool, resulting in quality-based incumbency advantage or disadvantage if the rate of re-running differs between marginal winners and losers.

Figure 6: Three types: symmetric and asymmetric distributions of quality that yield balance on quality following open-seat contests
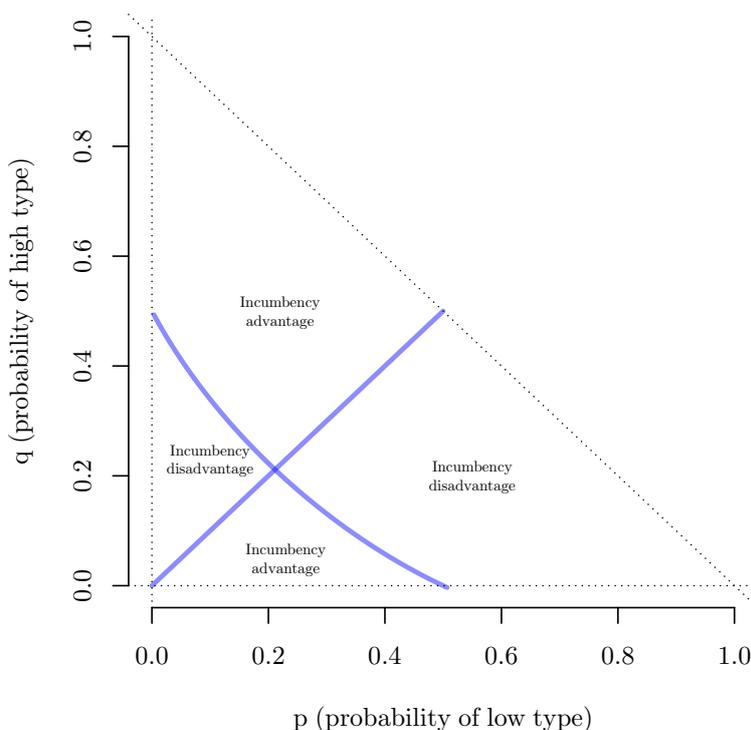
NOTE: Given three types $\{-1, 0, 1\}$ occurring in the candidate pool with probabilities $\{p, 1-p-q, q\}$, the thick curves show the combinations of p and q such that marginal candidates are no stronger or weaker than the candidate pool.

The intuition for the asymmetric solutions is as follows. Start from a situation where the distribution of quality over the three types is $\{1/2, 1/2, 0\}$, i.e. $p = 1/2$ and $q = 0$ and the candidate pool is evenly divided between low and medium types. (On Figure 6, this is the point halfway along the horizontal axis from the origin.) From the sufficiency part of Proposition 1 we know that there is balance on quality, as we effectively have a symmetric binary distribution of types. Now we consider moving some mass from the low type to the medium type while keeping the probability of the high type at zero (so that the distribution of types is characterized by $\{1/2 - \epsilon, 1/2 + \epsilon, 0\}$; on Figure 6 this change is represented by a leftward move from the point identified above). In the new distribution there will be quality-based incumbency advantage,

as the more common type (the medium type) will be over-represented in the pool of marginal winners and losers compared to the candidate pool, and thus marginal winners and losers will be stronger than the candidate pool. (This can be seen from Figure 3.) Going back to the original distribution, we now consider moving mass from the low type to the high type (so that the distribution of types is characterized by $\{1/2 - \epsilon, 1/2, \epsilon\}$; on Figure 6 this change is represented by an upward shift from the point identified above). The average quality of marginal incumbents can be calculated from Lemma 1 as $\frac{\epsilon^2 - \left(\frac{1}{2} - \epsilon\right)^2}{\epsilon^2 + \left(\frac{1}{2}\right)^2 + \left(\frac{1}{2} - \epsilon\right)^2}$; comparing this to the average quality of the candidate pool (which is $2\epsilon - \frac{1}{2}$), we confirm that the difference is zero when ϵ is zero but the derivative of this difference with respect to ϵ evaluated where $\epsilon = 0$ is negative, indicating that moving a small amount of mass from the low to the high type would result in quality-based incumbency disadvantage. Thus moving mass from the low type to the medium type makes marginal winners stronger than the candidate pool, while moving mass from the low type to the high type makes marginal winners weaker than the candidate pool. This implies that there is a k such we could move $k\epsilon$ mass to the low type and $(1 - k)\epsilon$ mass to the high type while preserving balance on quality. Graphically, this corresponds to moving in the northwesterly direction along the curve in Figure 6.