

Quality-Based Explanations of Incumbency Effects*

Andrew C. Eggers[†]

July 7, 2015

Abstract

Empirical studies of incumbency effects continue to accumulate, but the normative implications of these findings are more ambiguous than has been recognized. Using a simple model, I investigate mechanisms by which quality differences could emerge between incumbents and challengers. I show that in a baseline model where the distribution of candidate types is symmetric and constant over time, incumbency effects (as measured by standard empirical approaches) do not emerge; I then highlight three simple departures from the baseline model that produce quality-based incumbency effects. Each of these mechanisms can result in either incumbency advantage or disadvantage, but the normative implications of these incumbency effects depend on the specific mechanism involved. My analysis clarifies the mechanisms through which incumbency effects can emerge and highlights the difficulty of drawing conclusions about the health of democracy from findings of incumbency advantage or disadvantage.

*First version: February, 2014. This version: July 7, 2015. I thank Dimitri Landa, Leandro de Magalhães, Arthur Spirling, and seminar audiences at the University of Essex, the 2014 meeting of Midwest Political Science Association, the University of Rochester, and the University of Oxford for useful discussions.

[†]University of Oxford and Nuffield College. email: andrew.eggers@nuffield.ox.ac.uk

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 democracies² and, intriguingly, some evidence of incumbency disadvantage in several settings (Uppal, 2009; Klašnja and Titiumik, 2013; Klašnja, 2015a; Ariga, 2015).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 both the mechanisms that explain empirical findings and the normative implications of those findings. I focus on quality-based explanations of incumbency effects, by which I mean systematic differences between the characteristics of incumbents and challengers that lead to disproportionate incumbent success or failure. Ashworth and De Mesquita (2008) formalize the intuitively appealing “quality difference model”, in which incumbency advantage arises simply because incumbents are more appealing than challengers: after all, incumbents have won before, and the qualities that helped them win are likely to remain appealing to voters.⁴ The first contribution of this paper is to show that this account falls short as an explanation for most empirical findings of incumbency effects. This is because scholars studying incumbency effects use techniques that control for past electoral outcomes such as the regression discontinuity design (RDD) (Thistlethwaite and Campbell, 1960; Lee, 2008) or the more parametric approach of Gelman and King (1990). I show that in a simple model in which candidates emerge to contest

¹Important 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).

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

³De Magalhaes (2015) highlights potential problems in some of the findings of incumbency disadvantage; he advocates comparing what happens to winning and losing candidates rather than what happens to their parties (my focus). None of the mechanisms I study would produce an incumbency effect according to his definition, but they all would affect the magnitude of that effect.

⁴Ashworth and De Mesquita (2008)’s model also incorporates a scare-off effect: incumbents with sufficiently high quality run uncontested.

an open seat and the winner subsequently faces a challenger (a simplified version of [Ashworth and De Mesquita \(2008\)](#)), narrowly-elected incumbents have the same expected quality as challengers and thus incumbency effects do not arise.

Although differences between incumbents and (narrowly-elected) challengers do not emerge from the simplest model of incumbency, such differences do emerge under several minor departures from that model. The second contribution of this paper is to highlight three such departures that may help explain findings of incumbency effects. First, if candidate types are distributed asymmetrically (e.g. if most candidates are good but some are bad), then the average narrow winner from one election will be stronger or weaker than the average challenger from the next. Second, if the distribution of candidate types is changing over time (perhaps due to changes in the electorate's tastes), then clearly challengers may be stronger or weaker than candidates who narrowly won in the previous election. Finally, if incumbents' retirement decisions are related to their quality, narrowly-elected incumbents who run for re-election could be stronger or weaker on average than the challengers they face. Fundamentally, all three mechanisms rely on the assumption that the winner of an election (i.e. the incumbent) is more likely to run again than the loser; each mechanism highlights a different way in which the distribution of types among marginal winners of the first election may differ from the distribution of types among new candidates in the second election. One notable feature of these accounts is that, in contrast to most of the existing theoretical work, the same simple mechanisms can produce either incumbency advantage or disadvantage.

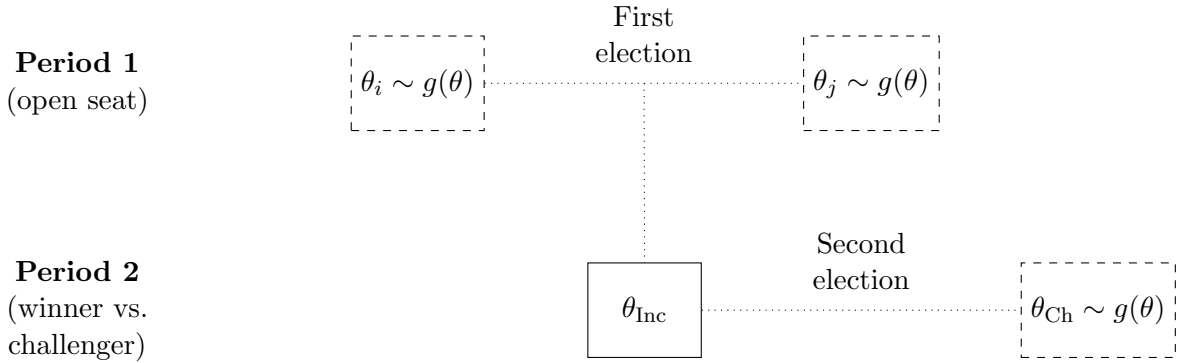
Building on this analysis, the third contribution of this paper is to emphasize that the normative implications of a finding of incumbency advantage or disadvantage are murkier than has been appreciated. Consider a finding of incumbency advantage. According to the conventional interpretation, such a finding is normatively troubling because it suggests that incumbents use the benefits of office to insulate themselves from electoral competition. This paper offers three mechanisms that produce incumbency advantage but that have nothing to do with incumbent insulation: incumbents are advantaged in these accounts simply because they are stronger on average than the candidates they face. Furthermore, the mechanisms I consider imply divergent conclusions about what incumbency advantage or disadvantage means for voter welfare.

If incumbency advantage is caused by a decline in the quality of candidates, then it suggests that voters would be *better off* if this quality decline were reversed and incumbency advantage disappeared; on the other hand, if incumbency advantage is caused by the disproportionate retirement of low-quality incumbents, then it suggests that voters would be *worse off* if this retirement pattern were stopped and incumbency advantage disappeared. Thus a finding of incumbency advantage (or disadvantage) cannot be definitively viewed as normatively more problematic than a counterfactual situation in which there are no incumbency effects.

Broadly, this paper is part of a recent strain of literature that seeks to provide a more satisfying theoretical explanation for empirical findings of incumbency effects. Early empirical studies in the U.S. tended to take for granted that the advantages of incumbents could be explained by the things incumbents did for voters or the extra resources they obtained as officeholders (e.g. Erikson, 1971; Cain, Ferejohn and Fiorina, 1987; Cox and Morgenstern, 1993); similarly, Uppal (2009) concludes that incumbents are disadvantaged in Indian state politics because they fail to do *enough* for voters. What is missing from these accounts is an explanation of why rational voters would systematically reward (or punish) incumbents for their past actions: what makes voters think that challengers would perform any differently? Muthoo and Shepsle (2014) (building on McKelvey and Riezman, 1992) attribute incumbency advantage to the seniority system in legislatures, which induces a systematic voter preference for incumbents because incumbents are more senior and can therefore acquire more local benefits.⁵ Caselli et al. (2013) links incumbency advantage to the enhanced ability of incumbents to signal their type to voters. Ashworth and De Mesquita (2008)'s answer, as noted above, is that incumbents on average are stronger than challengers due to electoral selection. This paper contributes to this literature by clarifying the role of electoral selection in explaining incumbency effects as estimated by empirical political scientists and highlighting three simple quality-based mechanisms that could explain why rational voters would systematically prefer incumbents (or challengers).

⁵See also Klačnja (2015b), in which incumbency disadvantage occurs in part because more experienced legislators are more adept at stealing.

Figure 1: A two-period framework for studying incumbency effects



NOTE: In the first period, two candidates meet, with their quality measures θ_i and θ_j drawn independently from the same distribution $g(\theta)$. In the second period, the winner of the first election stands for re-election against a challenger, whose quality θ_{Ch} is also drawn from $g(\theta)$.

1 Baseline model: no incumbency effects

Consider a sequence of elections, one at time $t = 1$ and one at time $t = 2$. (See the schematic in Figure 1.) In the first period, two parties independently pick candidates from a pool of candidates with a quality measure characterized by symmetric distribution $g(\theta)$. The types are observed and an election takes place. In the second period, the winner of the first-period election (whose quality is characterized by θ_{Inc}) competes against a challenger whose quality θ_{Ch} is drawn from the distribution $g(\theta)$. This is the same two-period structure used by [Ashworth and De Mesquita \(2008\)](#) to study the role of candidate characteristics in incumbency effects.

Whether incumbents are advantaged or disadvantaged in this model depends of course on how voters respond to candidate quality. I assume a continuum of citizens, each of whom has partisan preferences but also values politicians with higher values of θ . Given candidates i and j , each citizen k prefers candidate i if

$$\theta_i + \zeta_k > \theta_j,$$

where ζ_k measures the citizen's partisan preference in favor of i 's party. Denote by ζ_m the median partisan preference; $\zeta_m = 0$ implies an evenly balanced electorate, while $\zeta_m > 0$ implies an electorate that tends to prefer i 's party. The electorate is thus perfectly divided between the

two candidates when $\theta_i + \zeta_m = \theta_j$. To lay out the main points most simply, I assume in what immediately follows that $\zeta_m = 0$, and therefore that the electorate is balanced between the two parties; in the Appendix I prove the main results for the more general case.

There are many possible approaches to modeling the process by which voters draw inferences about candidates' quality: previous papers have variously assumed that voters learn by observing the incumbent's performance in office (Caselli et al., 2013), by observing whether a qualified challenger emerges (Gordon, Huber and Landa, 2007; Gordon and Landa, 2009), or by observing a noisy public signal of quality about both candidates (Ashworth and De Mesquita, 2008). In order to highlight my points in the simplest possible way, I assume that voters perfectly observe each candidate's quality prior to the election, with the implication that the election ends in a tie when $\theta_i = \theta_j$ (or, in the more general case, when $\theta_i + \zeta_m = \theta_j$). All of my conclusions would hold if voters instead received a noisy (unbiased) public signal of each candidate's quality in advance of the election, as in Ashworth and De Mesquita (2008).⁶ The analysis would also proceed in the same way, though again with added complexity, if we assumed instead that voters voted according to (noisy) private signals about candidate quality that they received in the first election and we asked what inferences voters would make about the quality of incumbents in the second period conditional on the results of the first election.⁷

In line with most of the empirical literature (e.g. Gelman and King, 1990; Lee, 2008), I define incumbency effects in the context of this model as follows:

Definition The incumbency effect for a given party is the effect of marginally winning the election at $t = 1$ on the probability of winning the election at $t = 2$.

Given the symmetry between the parties in my model, the incumbency effect for either party is given by $\Pr(\theta_{\text{Inc}} > \theta_{\text{Ch}} | m = 0) - .5$, where m refers to the margin in the first election. Thus incumbency effects occur when the marginal incumbent is stronger or weaker on average than the challenger she faces.

⁶To see this, note that Lemma 1 can be rewritten to apply to *perceived* quality of the candidates rather than actual quality: thus voters' posterior belief about the quality of the marginal incumbent is proportional to the square of the prior distribution of perceived quality. If the noise in the voters' signal is symmetric, this prior distribution will also be symmetric, which implies that Proposition 1 holds and thus that there is no incumbency effect.

⁷In this approach, the inference that we make in Lemma 1 and Proposition 1 would instead be made by the voter. Note that there is no private information in Ashworth and De Mesquita (2008) because there is just one representative voter.

I focus on marginal incumbents (rather than e.g. average incumbents) for the same reason empirical researchers using regression discontinuity methods (e.g. [Lee, 2008](#)) do: we want to isolate the effect of incumbency in the model from other differences that explain which candidates become incumbents and which do not. The same goal is accomplished in a more parametric way by [Gelman and King \(1990\)](#); the purpose of controlling for previous election outcomes in that model is to focus on the effect of incumbency after adjusting for other differences between settings where a party won and lost the previous election. By conditioning on a tie in the first election, we similarly control for electoral context in a way that maintains a close link to empirical work and thus aids in the interpretation of empirical findings.⁸

We can characterize the quality of marginal incumbents (i.e. those who won with a first-period margin $m = 0$) as follows:

Lemma 1 *Assume $\zeta_m = 0$. Then $p(\theta_{\text{Inc}}|m = 0) \propto g(\theta)^2$.*

Proof Given a balanced electorate ($\zeta_m = 0$), the first-period election will be tied if and only if $\theta_i = \theta_j$. We can characterize candidate i 's type, given a tied election, as follows:

$$\begin{aligned} p(\theta_i|m = 0) &\propto p(m = 0|\theta_i)g(\theta_i) \\ &= p(\theta_j = \theta_i|\theta_i)g(\theta_i) \\ &= g(\theta_i)g(\theta_i). \end{aligned}$$

Because candidates i and j are drawn from the same distribution of types, candidate j 's type (conditional on a tied election) could be characterized the same way. Thus whichever party wins the tied election, the incumbent's type is drawn from a distribution that is proportional to $g(\theta)^2$. ■

The lemma should make intuitive sense: for any type of candidate we can consider, that type of candidate will be found in a tied election only if two candidates of that same type emerged in the election; thus the probability of observing a given type of marginal *incumbent* is proportional to the square of the probability of observing that type of *candidate*.

⁸In the same spirit, both [Fowler and Hall \(2014, online appendix\)](#) and [Erikson and Titiunik \(2015, online appendix\)](#) decompose empirical findings of incumbency effects with formal analysis that conditions on a close election.

Using Lemma 1, we can show that incumbency effects do not arise in the baseline model:

Proposition 1 *Given symmetric $g(\theta)$, $\Pr(\theta_{Inc} > \theta_{Ch}|m = 0) = .5$.*

Proof Lemma 1 established that the posterior density of the quality of an incumbent who wins a tied election is proportional to $g(\theta)^2$. 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 $g(\mu + a) = g(\mu - a)$ for all a ; thus $g(\mu + a)^2 = g(\mu - a)^2$ for all a , meaning that $g(\theta)^2$ is also symmetric about μ .) Thus if $g(\theta)$ is symmetric around point μ , $g(\theta)^2$ is also symmetric around point μ . This implies that $\Pr(\theta_{Inc} > \theta_{Ch}|m = 0) = \Pr(\theta_{Inc} < \theta_{Ch}|m = 0)$, which is only possible if $\Pr(\theta_{Inc} > \theta_{Ch}|m = 0) = .5$. ■

Although the average incumbent is stronger than the average challenger, the *marginal* incumbent is just as strong as the average challenger when the distribution of candidate types is symmetric. Thus under our definition of incumbency effects (which reflects the definition used in most empirical work), no incumbency effects arise in the baseline model.

2 Three sources of quality-based incumbency effects

I have shown above that under the assumptions used in Ashworth and De Mesquita (2008), the marginal incumbent and the average challenger are equally strong and thus quality-based incumbency effects do not emerge. Small deviations from those assumptions do, however, produce incumbency advantages or disadvantages.

2.1 Asymmetry in the distribution of candidate types

We assumed above (following Ashworth and De Mesquita (2008)) that the distribution of candidate types $g(\theta)$ is symmetric. If $g(\theta)$ is not symmetric, however, the marginal incumbent could be stronger or weaker than the average candidate and thus incumbency advantages or disadvantages could emerge.

Intuitively, the reason why asymmetry in the distribution of candidate types produces incumbency effects is that, following from the analysis above (Lemma 1), close elections tend to

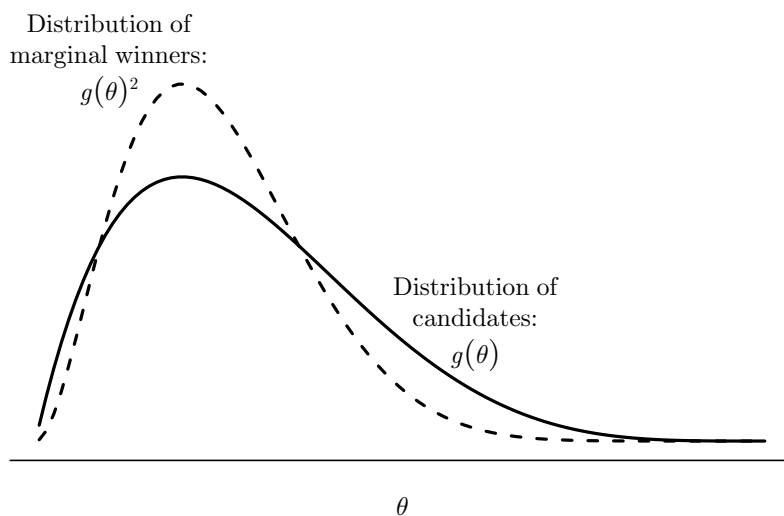
disproportionately feature candidates who are *typical* or *modal*; when the distribution of candidate types is asymmetric, these *typical* candidates may be systematically weaker or stronger than the *average* candidate. The key idea is that close elections tend to be elections in which the candidates are *similar* in quality. This means that types that are rare in the candidate pool will be even more rare in the pool of close winners; types that are common in the candidate pool will be even more common in the pool of close winners. If most candidates are good but a few are bad, the pool of close winners will be overwhelmingly good rather than bad because most of the bad candidates will have been beaten handily by good candidates; this implies that close winners will be systematically stronger than the average candidate, resulting in incumbency advantage.

To see the point more clearly, suppose there are just two types of politicians, low and high (θ_L and θ_H), and suppose that a close election only takes place when two same-type politicians face each other. If we denote by q the probability of a high-type candidate, the probability of a high-type incumbent (conditional on the election being close) is $\frac{q^2}{q^2+(1-q)^2}$: this is the fraction of same-type races that involve two high-type candidates.⁹ If $q = 1/2$, meaning that the distribution of types is symmetric, then conditioning on a close election the probability of a high-type incumbent is $1/2$; consistent with Proposition 1, there are no incumbency effects because the probability of a high-type challenger is also $1/2$. For other values of $q \in (0, 1)$, however, narrowly-elected incumbents will be stronger (where $q > 1/2$) or weaker (where $q < 1/2$) than candidates in general. Again, the reason is that common types are over-represented in close races, meaning that if high types are common ($q > 1/2$) close winners will be stronger than average candidates (producing incumbency advantage) and if low types are common ($q < 1/2$) close winners will be weaker than average candidates (producing incumbency disadvantage).

For an example with continuous types, consider Figure 2, in which the solid line depicts a hypothetical right-skewed distribution of candidate types $g(\theta)$ and the dashed line depicts the (normalized) posterior distribution $g(\theta)^2$ of candidate types conditional on a close election. Because the prior distribution of candidate types is right-skewed, stronger candidates are underrepresented in the pool of marginal winners and thus the average incumbent would be

⁹Note that this is just a special case of Lemma 1.

Figure 2: The distribution of candidate types among marginal winners, given an asymmetric prior distribution of candidate types



NOTE: The solid curve indicates the distribution of candidate types; the dashed curve indicates the resulting distribution of candidate types conditional on a tied election.

stronger than the average challenger. The same argument could of course be made in reverse for a left-skewed distribution, in which case weaker candidates would be underrepresented in the pool of marginal winners and thus the average incumbent would be weaker than the average challenger.

In practice, is the distribution of candidate quality symmetric or asymmetric? Of course it is difficult to say; candidate quality is not easy to define, let alone measure. In some cases it may be reasonable to assume that the qualities voters care about are distributed more or less normally, as are many things in nature, but in other cases symmetry seems unlikely to hold. If voters perceive candidates in a binary way (e.g. qualified/unqualified or corrupt/uncorrupt), it would be surprising if exactly half of all candidates ended up in each category. Asymmetries in the distribution of candidate types could also emerge for non-binary measures due to the process of candidate selection. If the process of candidate selection eliminates candidates who fall below a certain threshold of physical attractiveness, for example, then the candidate pool may end up being asymmetric due to the absence of unattractive candidates; similarly, asymmetries emerge if the process of candidate selection is designed to produce candidates with high levels of honesty

but fails to keep out some corrupt ones. Although it is difficult to judge the importance of these asymmetries in practice, we should at least be aware of the role they might play in explaining empirical findings.

2.2 Changes in candidates' characteristics over time

A more obvious point is that incumbency effects could arise if the distribution of candidate types changes over time, i.e. if the distribution of candidate types in the first period $g_1(\theta)$ is different from the distribution in the second period $g_2(\theta)$. The analysis above established that, assuming a symmetric distribution of candidate types, marginal incumbents will not be systematically stronger or weaker than the pool of first-period candidates. It follows that if the pool of candidates in the second period is weaker on average than the pool in the first period, there will be incumbency advantage; if the pool of candidates gets stronger, there will be incumbency disadvantage.

The distribution of candidate characteristics could change from one election to the next for a variety of reasons. For example, politics may become more or less attractive compared to other pursuits, which would affect the pool of potential candidates; the selection of candidates may also become more or less effective at yielding candidates the electorate will find attractive. Such changes could produce incumbency advantage or disadvantage, depending on whether candidates become weaker or stronger. Slightly extending the model, the pool of candidates may also change in response to changes in the electorates' preferences. To take a stark example, suppose voters prefer candidates with political experience in the first election, when the economy is at peace, but they prefer candidates with military experience in the second election, when the country is at war. Based on the analysis above, the marginal incumbent who emerges from the first round will be typical of the first-period candidates, who (given voter preferences) are likely to have political experience. In the second election, we might expect soldier-candidates to emerge and to have a systematic advantage over the politician-incumbents who are attempting to defend their seats. To the extent that candidates have fixed characteristics (whether they are types, ideological positions, or voting records) while the pool of challengers can respond to changes in voter preferences, we would thus expect incumbents to be at a systematic

disadvantage.¹⁰

Again, the difficulty of measuring candidate quality makes it difficult to assess whether changes in candidate quality are contributing to incumbency effects, but we can at least point to some intuitive implications for empirical researchers. If these changes are indeed an important aspect of incumbency effects in a given setting, we should expect to see that incumbency effects are concentrated in periods where the process of candidate selection and/or the tastes of voters are changing most substantially. For example, we might expect to see a spike in incumbency disadvantage in elections following major changes in voter priorities, such as a war or an economic crisis. During a period of persistent decline or increase in the quality of candidates, we would expect to see evidence of incumbency advantage or disadvantage, which would disappear when the system stabilized. In both cases we should expect to see that the incumbency effects coincide with changes in observable candidate characteristics such as educational or professional backgrounds.

2.3 Non-random retirement by marginal incumbents

In the baseline model above, we assumed that all incumbent candidates stand for re-election. If narrowly-elected incumbent-party candidates sometimes choose not to seek re-election, and if the decision to retire is related to candidate quality, then the incumbent-party candidates and challengers could systematically differ in ways that would produce incumbency advantage or disadvantage.

To see how non-random retirement could produce incumbency effects, recall that the distribution of ability in the set of marginal incumbents is proportional to $g(\theta)^2$, where $g(\theta)$ is the distribution of types in the candidate pool. Now suppose that incumbents with quality below some threshold k choose not to run for re-election, perhaps because they realize that it is more likely that they will lose. If μ is the expected value of the distribution of candidate types, the expected value of θ among the incumbent candidates who run again will be above μ ; if we assume that retiring incumbents are replaced by candidates drawn from $g(\theta)$, the expected value

¹⁰Kramer (1977) presents a model in which incumbent candidates must defend their positions in multidimensional policy space against challengers who choose vote-maximizing alternatives. By the usual “chaos” results, the incumbent always loses. Thus to the extent that incumbents have fixed positions in a political space without a Condorcet winner, they may be at a systematic disadvantage even in the absence of changing preferences.

of θ among all incumbent-party candidates will also be above μ . Thus the disproportionate retirement of weak incumbents produces incumbency advantage. The reverse is of course true if *strong* incumbents disproportionately retire, for example because they have greater opportunities to seek higher office or make money in the private sector: returning incumbents will be weaker than average, as will incumbent-party candidates on average, which will produce incumbency disadvantage.

The idea that incumbency effects could arise from strategic retirement decisions is related to a long-standing methodological debate about how to define and estimate incumbency effects. The [Gelman and King \(1990\)](#) approach to estimating incumbency effects measures the difference in a party's success when its candidate is the incumbent as compared to otherwise, controlling for prior outcomes. As [Gelman and King \(1990\)](#) recognized, this approach may overstate the party's benefits from fielding an incumbent candidate if incumbents strategically choose to retire when they face unfavorable circumstances.¹¹ It is partly because of this potential bias that researchers following [Lee \(2008\)](#) have focused instead on the the difference in a party's success when it won the previous election in the constituency as compared to otherwise (regardless of whether the incumbent herself stands for election). This analysis shows that strategic retirement may still matter when we move our focus to party incumbency effects, though in a different way: whereas strategic retirement was a source of potential bias in [Gelman and King \(1990\)](#)'s approach (because it indicated an omitted variable in the model), in the new approach strategic retirement is one of many mechanisms that may produce incumbency effects. Thus findings of party incumbency advantage may arise from the benefits of officeholding, for example, but also from systematic quality differences between incumbent-party candidates and their opponents that emerge due to strategic retirement decisions.

This discussion of strategic retirement may also help us to understand why there have been few findings of incumbency disadvantage, and why these findings have been concentrated in systems with weak parties. To the extent that candidates' quality becomes known through campaigning and/or serving in office, strategic parties should make efforts to force out incumbents with quality below that of the average challenger and replace them with newly-chosen

¹¹[Gelman and King \(1990\)](#) conclude that this bias is minimal in their case because U.S. House incumbents choose to run again even when involved in a substantial scandal.

candidates. If parties are successful in doing this, the incumbent party should generally be at an advantage regardless of asymmetries or changes over time in the distribution of candidate characteristics. This suggests that quality-based incumbency disadvantage should only arise when parties are unable to replace low-quality incumbents, for example because they do not control the candidate selection process.¹²

3 Normative implications

In this section we consider the normative implications of quality-based incumbency effects. In particular, we ask: If incumbency effects in a given setting are occurring because of one of the mechanisms introduced above, would voters be better off or worse off if that mechanism were “shut down” and we returned to the baseline model in which there are no incumbency effects? We define voter welfare as an increasing function of the quality of the candidate who wins the second election; the implications are the same if voter welfare also depends on the quality of the first-period winner.

3.1 Asymmetric candidate types

Compare two scenarios. In the baseline scenario, candidate types have symmetric density $g(\theta)$ with mean μ in both the first and second period, such that we find no incumbency effect by Proposition 1. In the alternative scenario, the distribution of candidate types in both periods is a mixture of density $g(\theta)$ (i.e. the same density as in the first scenario) and a point mass at some point x . By the logic of Section 2.1, there will be incumbency advantage or disadvantage depending on whether x is greater than or less than μ , i.e. depending on whether the asymmetry comes in the form of a dose of above-average candidates or below-average candidates. In which scenario are voters better off – the baseline scenario with no incumbency effects or the alternative scenario with incumbency advantage or disadvantage?

It is clear that voters in the second period are better off in the alternative scenario than in the baseline scenario if $x > 0$, i.e. if there is an added dose of above-average candidates.

¹²Even if a party is capable of replacing low-quality incumbents it may be that it benefits from not doing so: it could be that a party can improve its pool of prospective candidates by offering them the *de facto* right to stand for re-election regardless of the quality they turn out to have.

Table 1: Are incumbency effects normatively troubling? Comparison of voter welfare to baseline scenario with no incumbency effects

Source of incumbency effect	Voter welfare, compared to the baseline case	
	If incumbency advantage	If incumbency disadvantage
Asymmetric distribution of candidate types	Worse	Better
Change in distribution of candidate types from one election to the next	Worse	Better
Non-random retirement	Better	Worse

NOTE: The normative implications of quality-based incumbency effects are ambiguous. Compared to the baseline model with no incumbency effects, incumbency advantage is associated with higher voter welfare if it arises from the non-random retirement of marginal incumbents but lower voter welfare if it arises from asymmetry in the distribution of candidate types or changes in the distribution of candidate types over time. The reverse is true for incumbency disadvantage.

By contrast, voters are worse off in the alternative scenario with $x < 0$. In this comparison, then, incumbency disadvantage is associated with a welfare gain for voters while incumbency advantage is associated with a welfare loss. (Table 1 summarizes these conclusions along with the others in this section.)

3.2 Changing candidate types over time

Again compare two scenarios. The baseline scenario remains the one in which candidate types have symmetric density $g(\theta)$ with mean μ in both the first and second period, such that we find no incumbency effect by Proposition 1. In the alternative scenario the distribution of candidate types is shifted by a constant z in the second period; there will thus be incumbency advantage if $z < 0$ and incumbency disadvantage if $z > 0$, as discussed in Section 2.2. In which scenario are voters better off – the baseline scenario with no incumbency effects or the alternative scenario with incumbency advantage or disadvantage?

Here again the answer is quite simple: voters are better off than in baseline if $z > 0$ (i.e. the quality of candidates is improving over time) and worse off than baseline if $z < 0$ (i.e. the quality of candidates is getting worse). As above, incumbency disadvantage is associated with a welfare gain for voters compared to the baseline scenario and incumbency advantage is associated with

a welfare loss.

3.3 Non-random retirement by marginal incumbents

Our final comparison is between the baseline scenario (symmetric and time-invariant distribution of candidate types) and an alternative scenario in which marginal incumbents with quality above or below a fixed value choose to retire, at which point their party replaces them with a new draw from the candidate pool. As discussed in Section 2.3, if weak marginal incumbents retire we see incumbency advantage (because the average marginal incumbent-party candidate is better than the average candidate) and if strong marginal incumbents retire we see incumbency disadvantage (because the average marginal incumbent-party candidate is worse than the average candidate). In which scenario are voters better off – the baseline scenario with no incumbency effects or the alternative scenario with incumbency advantage or disadvantage?

If voter welfare is an increasing function of the quality of the second-period incumbent, voters in this simple model benefit from the retirement of weak incumbents and are hurt by the retirement of strong ones. Thus in this case (in contrast to the cases above) incumbency advantage is associated with a welfare gain for voters compared to the baseline scenario and incumbency disadvantage is associated with a welfare loss.

4 Conclusion

This paper has studied a simple model of electoral selection in order to clarify how systematic quality differences could emerge between marginal incumbents and challengers. Such differences do not emerge when the distribution of candidate types is symmetric and constant over time and incumbent retirement is independent of type; this suggests that empirical researchers can discard quality-based explanations of incumbency effects if those conditions are met in a given setting. Differences do emerge (and thus incumbency effects should be detected) if the distribution of candidate types is asymmetric, if that distribution changes over time, or if incumbent retirement is systematically related to quality. All three mechanisms could produce incumbency advantage or disadvantage, and the normative implications of each type of effect depend on the mechanism. Crucially, all of these conclusions emerge from studying a model of electoral selection using a

definition of incumbency effects that conditions on past electoral outcomes, as is standard in empirical work based on RDD (Lee, 2008) and the more parametric approach of Gelman and King (1990).

Along with other recent theoretical work (e.g. Meirowitz, 2008; Ashworth and 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 interpreting a simple finding of incumbency advantage or disadvantage, particularly when (as this paper shows) the normative implications depend on which mechanism is responsible. 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.

Appendix: More general case with partisan imbalance

Now consider a more general situation in which the electorate is not necessarily balanced:

Lemma 2 *If the winner of a tied election is chosen at random, then $p(\theta_{\text{Inc}}|m = 0) \propto \left(g(\theta + \zeta_m) + g(\theta - \zeta_m) \right) g(\theta)$.*

In words, the lemma states that the posterior probability of observing a given type of candidate emerging from a tied election is proportional to the prior probability of observing that type of candidate times the probability of seeing a candidate who is weaker or stronger than that candidate by precisely the degree of partisan imbalance.

Proof Given an advantage to i 's party of ζ_m , and assuming that the winner of a tied election is chosen at random, the posterior density of θ_{Inc} following a tied election can be characterized as follows:

$$\begin{aligned} p(\theta_{\text{Inc}}|m = 0) &\propto p(m = 0|\theta_{\text{Inc}})p(\theta_{\text{Inc}}) \\ &\propto p(\theta_j = \theta_i + \zeta_m|\theta_i)p(\theta_i) + p(\theta_i = \theta_j - \zeta_m|\theta_j)p(\theta_j) \\ &= g(\theta + \zeta_m)g(\theta) + g(\theta - \zeta_m)g(\theta). \end{aligned}$$

Rearranging gives the expression in Lemma 2. ■

Note that Lemma 1 is just a special case of Lemma 2 in which $\zeta_m = 0$.¹³

We can now state Proposition 1 more generally:

Proposition 2 *If $g(\theta)$ is a symmetric distribution with expectation μ , then $\Pr(\theta_{\text{Inc}} > \theta_{\text{Ch}}) = .5$.*

Proof Lemma 2 showed that $p(\theta_{\text{Inc}}|m = 0) \propto \left(g(\theta + \zeta_m) + g(\theta - \zeta_m) \right) g(\theta)$. If $g(\theta)$ is a symmetric function with point of symmetry μ , then $g(\theta + \zeta_m)$ is a symmetric function with point of symmetry $\mu + \zeta_m$ and $g(\theta - \zeta_m)$ is a symmetric function with point of symmetry $\mu - \zeta_m$; the sum of the two is a symmetric function with point of symmetry μ . Finally, note that the

¹³To see how the second line of the proof follows from the first, observe that the election is tied if there is a quality difference of ζ_m between the candidates; if a candidate of type θ wins a tie-breaker, this can occur either because she is the weaker candidate (who benefits from partisan imbalance) or because she is the stronger candidate (who is hurt by partisan imbalance). To get to the third line, we only need the fact that θ_i and θ_j are drawn independently from $g(\theta)$.

product of two symmetric functions with a 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 $h_1(\mu + a) = h_1(\mu - a)$ and $h_2(\mu + a) = h_2(\mu - a)$ for all a , and thus that $h_1(\mu - a)h_2(\mu - a) = h_1(\mu + a)h_2(\mu + a)$ for all a , which implies that $h_1(x)h_2(x)$ is symmetric with point of symmetry μ .) Thus the posterior density $p(\theta_{Inc}|m = 0)$ is, like $g(\theta)$, a symmetric density with point of symmetry μ . This symmetry implies that $\Pr(\theta_{Inc} > \theta_{Ch}|m = 0) = \Pr(\theta_{Inc} < \theta_{Ch}|m = 0)$, which is only possible if $\Pr(\theta_{Inc} > \theta_{Ch}|m = 0) = .5$. ■

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.
- 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.

- Erikson, Robert S. 1971. "The advantage of incumbency in congressional elections." *Polity* 3(3):395–405.
- Erikson, Robert S and Rocio 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.
- 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.
- Gordon, Sanford C and Dimitri Landa. 2009. "Do the advantages of incumbency advantage incumbents?" *The Journal of Politics* 71(04):1481–1498.
- Gordon, Sanford C, Gregory A Huber and Dimitri Landa. 2007. "Challenger entry and voter learning." *American Political Science Review* 101(02):303–320.
- 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.
- Horiuchi, Yusaku and Andrew Leigh. 2009. "Estimating Incumbency Advantage: Evidence from Multiple Natural Experiments." *Unpublished paper* .
- 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. 2015a. "Corruption and the Incumbency Disadvantage: Theory and Evidence."
- Klašnja, Marko. 2015b. "Increasing Rents and Incumbency Disadvantage." *Journal of Theoretical Politics* .
- Klašnja, Marko and Rocío Titunik. 2013. "Incumbency Disadvantage In Weak Party Systems: Evidence from Brazil."
- Kramer, Gerald H. 1977. "A dynamic model of political equilibrium." *Journal of Economic Theory* 16(2):310–334.
- 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.
- Thistlethwaite, Donald L and Donald T Campbell. 1960. "Regression-discontinuity analysis: An alternative to the ex post facto experiment." *Journal of Educational psychology* 51(6):309.
- Uppal, Yogesh. 2009. "The disadvantaged incumbents: estimating incumbency effects in Indian state legislatures." *Public Choice* 138(1-2):9–27.