

Electoral Selection, Strategic Challenger Entry, and the Incumbency Advantage

Scott Ashworth Princeton University
Ethan Bueno de Mesquita University of Chicago

We study the comparative statics of the incumbency advantage in a model of electoral selection and strategic challenger entry. The incumbency advantage arises in the model because, on average, incumbents have greater ability than challengers. This is true for two reasons: high-ability candidates are more likely to win election (electoral selection) and high-quality incumbents deter challengers (strategic challenger entry). We show that this quality-based incumbency advantage is expected to be greater for high visibility offices, in politics with relatively small partisan tides, in unpolarized electoral environments, and in electorates that are relatively balanced in their partisan preferences.

Electoral politicians are extraordinarily successful when they seek reelection. One prominent explanation of this incumbency advantage is that incumbents have, on average, higher quality than challengers. This is true for at least two reasons. First, there is *electoral selection* (Ashworth 2005; Banks and Sundaram 1998; Gowrisankaran, Mitchell, and Moro 2008; Londregan and Romer 1993; Samuelson 1984, 1987; Zaller 1998). Voters elect candidates whom they like on some dimension (e.g., ability or trustworthiness), so, compared to the average challenger, candidates who have won in the past are relatively more attractive on that dimension.¹ Second, there is *strategic challenger entry* (Cox and Katz 1996; Gordon, Huber, and Landa 2007; Jacobson 1980; Jacobson and Kernell 1983; Stone, Maisel, and Maestas 2004).² Challengers may be deterred from running against incumbents who are perceived to be of high ability.³ We refer to the combination of electoral selection and strategic challenger entry as the *quality difference* model of the incumbency advantage.

To evaluate the validity of the quality difference model, it is important to know what implications it has

beyond predicting the existence of an incumbency advantage—all theories of the incumbency advantage make that prediction. The literature to date has focused on predictions about the dynamics of the incumbency advantage over a politician's career: Samuelson (1984) proves that an incumbent's expected margin of victory will increase over time in a pure selection model. Samuelson (1987) provides a supportive test of this prediction, and Zaller (1998) argues that actual electoral careers resemble those from simulations of a pure selection model. This claim is confirmed by the structural estimates of Gowrisankaran, Mitchell, and Moro (2008).

These studies provide support for the claim that electoral selection is an important part of the incumbency advantage. However, they leave open the question of whether the quality difference model can account for the broad range of variation in incumbent success that is observed over time and across different offices and electorates. For instance, Ansolabehere, Snyder, and Stewart (2000) demonstrate that congressman from more partisan districts have smaller incumbency advantages, while Ansolabehere and

¹Diermeier, Keane, and Merlo (2005) study a related model in which candidate quality affects reelection probabilities and strategic decisions to run for higher office or to enter the private sector.

²Recently, Cox and Katz (2002) suggested that changes in redistricting since 1964 have contributed to an increased mismatch between challengers and incumbents, enhancing the incumbency advantage.

³A closely related idea is that challengers are less effective than open-seat candidates because they have less access to campaign funds (Gerber 1998).

Snyder (2002) find that the incumbency advantage is larger for more visible offices and has been increasing over time for all state-wide offices.

In light of these empirical findings, we study the comparative statics of the quality difference model in order to understand how variation in the office or in the electoral environment affects the magnitude of the advantage. We find that the quality difference model is consistent with the empirical observations discussed above. Moreover, our comparative statics yield a variety of additional predictions that constitute novel testable implications of the quality difference model. In particular, we show that the quality-difference-based incumbency advantage should be greater for high-visibility offices, when partisan tides are small, and in electorates that are either relatively unpolarized, non-partisan, or both.⁴ In the conclusion, we consider the relationship between the quality difference model and accounts of the incumbency advantage based on constituency service, pork barreling, and other factors that are specific to legislative politics.

The Model

At each of two dates, $t = 1$ and $t = 2$, a voter must elect a politician. Each candidate c has ability $\theta_c \in \mathbb{R}$, and these abilities are independent draws from a normal distribution with mean 0 and variance σ_θ^2 (denoted $\theta_c \sim \mathcal{N}(0, \sigma_\theta^2)$). No one has any private information about these abilities—the only information about them are the signals discussed below. In addition to her ability, each candidate has a policy position. In particular, there are left (L) and right (R) parties, each with a fixed policy platform in the one-dimensional policy space. We denote these platforms by $\mu_p \in \mathbb{R}$, where $p \in \{L, R\}$, with $\mu_L = -\mu_R$. In the first election (the *open-seat election*), each party runs one candidate at its location. In the second election, the incumbent is the candidate for her party, while recruiting a challenger costs the out party $\kappa \in [0, 1/2)$. A party makes its recruitment decision to maximize the probability of victory minus the recruitment cost, without knowing the realized ability of its candidate.

Before the first election, the voter gets a signal about each candidate's ability. This signal for candidate c is $s_{1c} = \theta_c + \varepsilon_{1c}$, where ε_{1c} is distributed $\mathcal{N}(0, \sigma_{\varepsilon_1}^2)$ and is independent of θ_c . The winner of this first

election (the *incumbent*) serves for the first period and then might face another candidate (the *challenger*) in a reelection contest. The voter also observes signals of the candidates' abilities before this second election. The second signal is $s_{2c} = \theta_c + \varepsilon_{2c}$, where $\varepsilon_{2c} \sim \mathcal{N}(0, \sigma_{\varepsilon_2}^2)$.

The voter has preferences represented by the expectation of $\theta - (x^* - x)^2$, where x is a policy and x^* is the voter's ideal point. That is, the voter cares both about candidate ability and about the policies enacted.⁵ The voter's ideal point in each election is an independent draw from a $\mathcal{N}(\gamma, \sigma_{x^*}^2)$ distribution. A negative value of γ indicates an electorate that leans left on average, while a positive value indicates an electorate that leans right on average. Thus γ represents persistent partisan leanings, while the residual $x^* - \gamma$ represents short-term partisan tides—the voter's ideal point can change from election to election due to idiosyncratic factors. This will lead to the partisan tides observed in empirical studies of elections.

Table 1 summarizes the notation introduced in this section. In several places in the text and conclusion we also explore the implications of various extensions of the model. In particular, we discuss implications of the voter receiving an extra signal about the incumbent's quality by observing the incumbent while in office, party primaries as an additional source of electoral selection, and endogenous, reputation-building effort by the incumbent while in office.

Preliminaries

Our main goal is to derive the comparative statics of a quality-difference-based incumbency advantage. In this section, we attend to two preliminaries. First, we derive the voter's optimal voting rule. Second, we define the incumbency advantage in terms of the model.

The Voter's Decision Rule

In each election, a candidate is characterized by the voter's beliefs about her ability. The updating rules follow from standard results on Bayes' rule with a

⁴The result on partisan tides formalizes a finding from Zaller's (1998) simulations. The other results have not previously been derived in any manner.

⁵There are several possible interpretations of the ability term, θ , including the ability of a candidate to secure pork-barrel spending for her district, charisma, or demagogic ability. It is important to note that one possible interpretation of ability, namely "ability to advance the party platform," is not consistent with the model's set-up, since ability and policy concerns are separable in the voter's utility function.

TABLE 1 Symbols introduced in Section 1

Symbol	Interpretation	Symbol	Interpretation
θ	Ability of a candidate	κ	Recruitment cost
σ_θ^2	Prior variance of ability	μ_p	Policy platform of party p
ε_i	Observation noise in election $i \in \{1, 2\}$	x^*	Voter's ideal point
$\sigma_{\varepsilon_i}^2$	Variance of ε_i	$\sigma_{x^*}^2$	Variance of x^*
		γ	Voter's partisan leaning

normal prior and normal signal (DeGroot 1970) and are formally described in the appendix (section A.1).

Intuitively, the voter's posterior beliefs about a candidate's ability are normally distributed with a mean that is a weighted average of the voter's prior mean belief and the signal the voter observes. Thus, the better the signal of candidate ability, the higher quality the voter believes the candidate to be. Moreover, the noisier the signal, the less weight the voter puts on it. The mean of the voter's belief about a candidate c 's ability is denoted m_c .

The voter chooses a candidate by comparing the expected utility of each choice. His decision rule is characterized by the following result, the proof of which follows from a direct comparison of the voter's expected utility for each candidate.

Lemma 1 *The voter votes for candidate L if and only if $m_L - m_R \geq 2(\mu_R - \mu_L)x^*$.*

The voter's voting rule takes into consideration two factors: his beliefs about the candidates' abilities (m_L, m_R) and his policy preferences. The term $2(\mu_R - \mu_L)x^*$ represents the policy component of the voter's choice. Since $\mu_R - \mu_L$ is always positive (i.e., the right-wing party is to the right of the left-wing party), when the voter is right-wing ($x^* > 0$), this term is positive—right-leaning voters are more likely to elect right-wing candidates. Similarly, when the voter is left-wing ($x^* < 0$), this term is negative—left-leaning voters are more likely to elect left-wing candidates. Intuitively, if the voter were exactly indifferent between the two parties on policy ($x^* = 0$), then the voter would vote for the left-wing candidate if and only if his expectation of the left-wing candidate's ability (m_L) was higher than that of the right-wing candidate's ability (m_R). Similarly, if the voter believed the two candidates were of identical ability ($m_L = m_R$) he would vote for the right-wing candidate if he was a right-leaning voter ($x^* > 0$) and vote for the

left-leaning candidate if he was a left-leaning voter ($x^* < 0$).⁶ Since it will play a major role in the analysis, we introduce the following notational simplification: $\eta(x^*) \equiv 2(\mu_R - \mu_L)x^*$. When no confusion can result, we simply write η , which we refer to as the “policy factor” in the voter's decision.

The Incumbency Advantage

The incumbency advantage is a measure of the improvement in a party's expected electoral performance when its candidate is an incumbent rather than a candidate in an open-seat election.

Definition 1 *The incumbency advantage for party p is the expected difference between the probability party p wins the election with an incumbent and the probability party p wins an open-seat election.*

Definition 2 *The incumbency advantage is a weighted average of the incumbency advantages for left- and right-wing incumbents, where the weight on a party p incumbent is given by the probability that party p wins the open-seat election.*

To see the intuition for these definitions, consider a simple example. Suppose that the probability of the left-wing party winning an open-seat election is 48% and the probability of the right-wing party winning an open-seat election is 52%. Further, suppose that the probability of a left-wing incumbent winning is 54% and the probability of a right-wing incumbent winning is 56%. The incumbency advantage for the left-wing party is $54\% - 48\% = 6\%$, while the incumbency advantage for the right-wing party is $56\% - 52\% = 4\%$. However, the right-wing party will hold office more frequently than the left-wing party, in expectation. Thus, the overall incumbency advantage is $.48(.54 - .48) + .52(.56 - .52) = 4.96\%$.

The extra probability of winning associated with being an incumbent is not the only way to define the incumbency advantage. Indeed, empirical work often estimates the advantage as the difference in the incumbent's share of the two-party vote and the incumbent's party's share of the two-party vote in open-seat elections (the normal vote). Our results also apply to this definition, at least in a simple one-dimensional model of voter heterogeneity. Assume

⁶Notice that the voting rule implies that, for any given difference in the candidate's expected abilities, if the shock to the voter's ideal point is large enough (and in the correct direction), he will vote for the candidate with lower expected ability. This is a special feature of quadratic utility. Most of our results do not rely on it. Instead, our results rely on the existence of a tradeoff between policy and ability, which will exist quite generally.

that there is a continuum of voters, with policy preferences y_i distributed $\mathcal{N}(x^*, \sigma_y^2)$. We refer to this variant as a model with a *heterogeneous electorate*. As before, the median voter's ideal point, x^* , is itself a random variable, with distribution $\mathcal{N}(0, \sigma_{x^*}^2)$. Each voter's payoff is $-(\mu_p - y_i)^2 + m_p$. In this model, a voter will vote for the left-wing candidate if

$$m_L - m_R > 2(\mu_R - \mu_L)y \Leftrightarrow y < \frac{m_L - m_R}{2(\mu_R - \mu_L)} \equiv \bar{y}.$$

The left-wing vote share is the percent of the population to the left of \bar{y} . Note that this is the same rule as in Lemma 1, with x^* replaced by y so that we can consider the percent of the vote rather than the probability of winning. This definition will not lead to numerically identical results, since typically $\text{var}(y) = \sigma_y^2 + \sigma_{x^*}^2 > \sigma_{x^*}^2$. Nonetheless, the common structure of the two models means that the comparative statics will be the same whether we measure the incumbency advantage with respect to vote share or probability of winning.

A Baseline Model

This section studies the special case of the model which sets both the partisan bias (γ) and the recruitment cost (κ) to zero. We call this specification the *baseline model*. We start with this case for two reasons. First, the symmetry of the baseline model allows for a particularly clear development of the intuition behind our characterization of the incumbency advantage. Second, in this case we are able to derive analytically a rich collection of comparative static results. The next section relaxes the assumptions of the baseline model and uses a mix of analytical results and simulations to study the incumbency advantage and its comparative statics, along the way showing that the results of the current section are robust.

The Incumbency Advantage Exists

Since the candidates have identical *ex ante* expected abilities and the baseline model's electorate is evenly balanced between the parties on average, each candidate in the open-seat election wins with probability one half. Moreover, because the recruitment cost is zero, there will always be a challenger in the second election. The incumbency advantage, then, is the extent to which an incumbent's probability of defeating a challenger is greater than one half.

Consider an election with a left-wing incumbent whose posterior mean ability is m . By Lemma 1, she wins reelection exactly when

$$m_L - m_R - \eta \geq 0.$$

The left-hand side of this inequality is a mean m normal random variable; let σ be its variance (defined formally in Appendix A.1). Thus the L incumbent wins with probability $1 - \Phi\left(\frac{-m}{\sigma}\right) = \Phi\left(\frac{m}{\sigma}\right)$, where Φ is the cumulative distribution function of the standard normal.

To calculate the incumbency advantage, we average this probability conditional on m over the distribution of incumbent abilities. Since the innovation to the voter's belief about the incumbent as a result of the second signal is independent of m ,⁷ we can use the standard convolution formula for the sum of independent random variables to write the probability that an incumbent wins as

$$\int \Phi\left(\frac{m}{\sigma}\right) f(m) dm, \quad (1)$$

where f is the distribution of the posterior means.

This formula lets us formalize the idea that incumbents do well in their reelection contests because, on average, they are of higher ability than challengers. Specifically, f is better than the prior distribution in the sense of first-order stochastic dominance. Thus it is more likely that incumbents win than that a candidate of the same party wins an open-seat election.

Proposition 1 *In the baseline model, the incumbency advantage exists—an incumbent wins with probability greater than one half.*

This is a special case of Proposition 6, so we do not give an independent proof.

Comparative Statics

Having established the existence of the incumbency advantage within the baseline model, we can now turn to our main task: exploring how changes in the electoral environment change its magnitude.

All of the comparative statics discussed in this section result from the interaction of two factors, reflecting the two factors in the integral (1). The first is the degree of *electoral selection*. Changes in the environment that enhance the degree of selection for ability increase the incumbency advantage, other things equal. The second factor involves changes to

⁷This is a standard result on the normal distribution.

the probability that an incumbent wins, conditional on her reputation following selection. We call this the incumbent's *insulation*, since it is the probability that an incumbent with fixed reputation retains office—how well she is insulated from high-quality challengers and partisan tides. With low costs of entry (and consequently incumbents running, on average, as favorites), the more insulated the incumbent, the larger the incumbency advantage.

Our first two comparative statics are driven by the following intuitions. First, voters who are close to indifferent between the two parties are more likely to vote based on the quality dimension, rather than the policy dimension. When this is the case, voting in the open-seat election responds more to quality, leading to a larger electoral selection effect and, thus, a greater expected difference in quality between incumbents and challengers in the second election. Second, when voters are expected to be close to indifferent between the parties, it is less likely that an incumbent with above-average quality will be replaced as a result of partisan tides. That is, the incumbents are more insulated. Hence, any factor that increases the likelihood that the voter is close to indifferent between the two parties on the policy dimension tends to increase the incumbency advantage through both the selection and insulation effects. (Omitted proofs for this section are in Appendix B)

Lemma 2 *In the baseline model, the greater is the variance of the policy term (σ_η^2), the smaller is the incumbency advantage.*

Partisan Tides Recall that policy entered into the voter's decision through the term $\eta = 2(\mu_R - \mu_L)x^*$. This policy factor has two components: the polarization of the parties ($\mu_R - \mu_L$) and the voter's ideal point (x^*). This ideal point is distributed normally with mean 0 and variance $\sigma_{x^*}^2$. We can think of the variance of this distribution as a measure of the average size of shocks to partisanship. When the partisan shock is large, voters are strong partisans and, thus, are particularly likely to base their votes on policy. This implies that, on average, voter decisions are less likely to be influenced by quality considerations when partisan tides are expected to be large than when they are expected to be small, diminishing the quality difference based incumbency advantage. Thus, we have the following result.

Proposition 2 *In the baseline model, the larger is the average partisan shock ($\sigma_{x^*}^2$), the smaller is the incumbency advantage.*

Proof The result follows from Lemma 2 and the fact that $\sigma_\eta^2 = 4(\mu_R - \mu_L)^2 \sigma_{x^*}^2$.

This result implies that larger partisan tides decrease the incumbency advantage.⁸ As far as we know, this relationship has not been studied empirically. But the mechanisms behind the result seem to work in actual elections.⁹ Consider, for example, the congressional elections of 1994. The substantial partisan tide in favor of the Republicans swept out an unusual number of incumbent Democrats. And the small measured incumbency advantage through the rest of the 1990s reflects, on this view, the fact that the tide swept in many below-average-quality Republicans. Proposition 2 is just the statement that, when the typical partisan shock is large, these two mechanisms keep the incumbency advantage small.

From this perspective, the fortunes of the congressional class of 1974 seem anomalous—the Democrats who rode into office on the Watergate-inspired partisan tide fared extremely well in their reelection efforts. But the quality difference model has potential resources to account for these candidates as well. The partisan tide of 1974 was well anticipated, and, as a consequence, the Democratic challengers and open-seat candidates were above average in quality (Jacobson and Kernell 1983). Comparing the elections of 1974 and 1994 illustrates the importance of our assumption that potential candidates in the open-seat election have no information about the realization of the partisan shock.

Polarization of the Parties The other component of the policy factor in vote choice is the polarization of the parties. When the parties are highly polarized, voters are more likely to have strong preferences for one party or the other. This diminishes the quality advantage of incumbents, on average, because voters are willing to accept candidates with subpar quality if they are of the favored party and are willing to replace high quality candidates if they are of the wrong party. Thus, a highly polarized party system weakens the incumbency advantage in the quality-difference-based model.

Proposition 3 *In the baseline model, the more polarized the two parties are the larger is ($\mu_R - \mu_L$), the smaller the incumbency advantage.*

Proof The result follows from Lemma 2 and the fact that $\sigma_\eta^2 = 4(\mu_R - \mu_L)^2 \sigma_{x^*}^2$.

This result is not driven by quadratic utility over policy—any strictly concave utility function would

⁸This result is similar to Zaller's (1998) comparative static on electoral "luck."

⁹We thank an anonymous referee for prompting this and the following paragraphs.

lead to the same result. Instead, the key assumption behind this result is that the uncertainty driving the probabilistic voting is over the voter's ideal point, rather than over an additive popularity shock for one of the parties. As discussed in more detail by Ashworth and Bueno de Mesquita (2007), the random ideal point model does, and the random popularity shock model does not, make the policy/quality trade-off a function of platform polarization.

Visibility of the Office. Elections for some offices, such as governor or U.S. Senator, attract more media and voter attention than do elections for other offices, like state auditor. The natural way to formalize this in our model is to assume that voters receive more informative signals about high-visibility office seekers. Formally, this means that the signals of candidate ability have lower variance for more visible offices.

Consider first the impact of more informative signals in the open-seat election. When these signals are more informative, the voter is more likely to correctly identify the higher quality candidate. Thus, the selection effect is heightened.

There is another, reinforcing, effect. When the voter gets better information, he becomes confident about his beliefs more quickly, leading him to be less responsive to subsequent signals. Thus, once he has chosen a candidate, future signals are less likely to convince him to change his mind, which helps the incumbent (who runs as the favorite) by increasing the insulation effect.

While the voter is less responsive to information in the second election, if this information is, for some reason, significantly more precise than the information in the first election, then the voter will take note. This can only hurt the incumbent, as it decreases insulation. Thus, the incumbency advantage is increasing in the informativeness of the signals in the first election and decreasing in the informativeness of the signals in the second election. It may make intuitive sense to think that, on average, the first and second elections are equally informative (holding the office in question constant). That is, races for, say, governor in a given state attract about the same level of media attention from election to election.¹⁰ (This does not mean that the voter has the same amount of information about both candidates, since he has observed two signals on the incumbent versus only one for the challenger.) If the two elections

are equally informative, then the original argument above implies that the incumbency advantage is increasing in the overall informativeness of the signals. These results are stated in the following proposition.¹¹

Proposition 4 *In the baseline model, assume that there is no uncertainty about the voters' ideal points ($\sigma_\eta^2 = 0$). The incumbency advantage is increasing in the informativeness of the first election signal ($\frac{1}{\sigma_{\epsilon_1}^2}$) and is decreasing in the informativeness of the second election signals ($\frac{1}{\sigma_{\epsilon_2}^2}$). If both elections are equally informative, then the incumbency advantage is increasing in the informativeness of the election signals.*

Proposition 4 predicts that high-visibility offices will have larger incumbency advantages than low-visibility offices.¹² Ansolabehere and Snyder (2002) observe exactly this pattern in their empirical study of the incumbency advantage for all statewide elected offices. It is also worth noting that this result allows us to consider what would happen to the incumbency advantage and our comparative statics if the voter observed an extra signal about the incumbent that he did not observe about the challenger (since the voter observes the incumbent during her term in office). Formally, this would be equivalent to decreasing the variance of the second signal for the incumbent. Such a change would decrease the size of the incumbency advantage. This is for the same reason that the incumbency advantage is decreasing in the informativeness of the second period signal—on average, the incumbent runs as a favorite in the second election, so it is not in her interest for the voter to learn more information, which might change the status quo. However, this extra signal would have no effect on the structure of the underlying causes of the quality difference based incumbency advantage. Thus, none of our comparative statics would change.

Polarization of the Electorate. Earlier, we discussed an interpretation of our model that measures the incumbency advantage as the increase in share of the vote won by an incumbent, rather than as the increase in the probability of winning. In this version,

¹¹The proof of these results relies on the special assumption that $\sigma_\eta^2 = 0$. However, we demonstrate the robustness of the results to relaxing this assumption in our computational simulations in the next section.

¹²The results about variation in the first-period signal variances would be enhanced if the voter was risk-averse over ability—then she would prefer the candidate with lower posterior variance, all else equal. In fact, we could get those results in a model driven entirely by risk aversion, without electoral selection. But risk aversion alone is insufficient to derive the other comparative statics in the paper, since the other changes we consider do not change the posterior variance.

¹⁰On the other hand, many offices, e.g., legislative races, probably attract more attention in open-seat races. We can capture this by letting the open-seat signals have variance σ_ϵ^2 and the second election signals have variance $\beta\sigma_\epsilon^2$, with $\beta < 1$. The same logic that applies to the case of equal variances would then imply that the incumbency advantage is increasing in σ_ϵ^2 .

we can also ask what happens to the incumbency advantage as the electorate becomes more polarized. A more polarized electorate corresponds, in the model, to a greater dispersion of individual voter policy preferences (σ_y^2). This increase in polarization does not affect the degree of quality difference because the winner of the open-seat election is determined by the median voter. However, more polarization does affect the share of the vote received by an incumbent of any given ability. When the electorate becomes more polarized, policy considerations loom large for a greater number of voters, diminishing quality-based voting. Thus, as polarization increases, the incumbency advantage—measured as additional vote share—decreases.

Proposition 5 *In the baseline model with a heterogeneous electorate, the more polarized the electorate (σ_y^2), the smaller is the incumbency advantage in terms of vote share.*

Two Measures of the Incumbency Advantage. As we noted earlier, there are at least two reasonable measures of the incumbency advantage: increased probability of winning and increased share of the vote. Further, as we discussed above, while the two measures share comparative statics, they may not generate identical point predictions. Because of this, our model has implications for how estimates of the incumbency advantage will differ depending on which measure is employed. The intuition is simple—the probability of winning definition corresponds to the vote share definition for an electorate that is identical in every respect except that polarization of the electorate is set to zero. Thus, Proposition 5 implies the following result.

Corollary 1 *In the baseline model, the incumbency advantage measured by probability of winning will be greater than the incumbency advantage measured by share of vote.*

And, indeed, this difference is observed across empirical analyses of the incumbency advantage. For example, Gowrisankaran, Mitchell, and Moro (2008) find that incumbency increases the probability of winning a Senate election by approximately 15 percentage points, while Ansolabehere and Snyder (2002) report that incumbency increases vote shares in Senate elections by approximately 9 percentage points.

The Full Model

Now we relax the baseline model’s assumptions that the district is perfectly balanced between the two

parties and that there is no recruitment cost. Because the model is symmetric, we can focus attention on L incumbents.

The Incumbency Advantage Exists

As before, an L incumbent facing a challenge wins reelection exactly when

$$m_L - m_R - \eta \geq 0.$$

There are two differences between reelection probabilities in the general case and the baseline case. First, the left-hand side of the inequality now has mean

$$m - 2(x_R - x_L)\gamma \equiv m - \bar{\eta},$$

yielding a reelection probability conditional on a challenge of $\Phi\left(\frac{m - \bar{\eta}}{\sigma}\right)$.

Second, not all incumbents will face challengers. Instead, the R party will choose to field a candidate if and only if the probability she will win exceeds the recruitment cost. Thus the R party will mount a challenge if and only if

$$1 - \Phi\left(\frac{m - \bar{\eta}}{\sigma}\right) \geq \kappa.$$

The left-hand side is strictly decreasing in the incumbent’s expected ability (m), so there is a cutoff m^* such that there is a challenge if and only if the incumbent’s reputation is less than or equal to m^* . This cutoff can be written as a function of the costs of mounting a challenge (κ), the partisan leaning ($\bar{\eta}$), and the variance (σ):

$$m^*(\kappa, \bar{\eta}, \sigma) = \sigma\Phi^{-1}(1 - \kappa) + \bar{\eta}.$$

Since $\kappa < 1/2$, we have $m^* > \bar{\eta}$. Thus the marginal challenger runs as an “underdog” in the sense of having lower expected ability (net of expected partisanship) than the incumbent. In addition, this function is differentiable in κ and σ , with derivatives

$$\frac{\partial m^*}{\partial \kappa} < 0 \text{ and } \frac{\partial m^*}{\partial \sigma} > 0.$$

These comparative statics are intuitive: higher costs directly make running less attractive, while a higher variance increases the likelihood of an upset, making running more attractive for the marginal (underdog) challenger.

Let $V(m, m^*, \bar{\eta})$ be the probability an incumbent with expected ability m is reelected, given a cutoff for mounting a challenge of m^* . We have

$$V(m, m^*, \bar{\eta}) = \begin{cases} \Phi\left(\frac{m-\bar{\eta}}{\sigma}\right) & \text{if } m \leq m^* \\ 1 & \text{if } m > m^* \end{cases}.$$

Because $\Phi(m/\sigma) < 1$ for all m , the function V is decreasing in m^* in the following sense:

$$\begin{aligned} m^* > m^{**} &\Rightarrow \\ V(m, m^*, \bar{\eta}) &\leq V(m, m^{**}, \bar{\eta}) \\ &\text{with strict inequality if } m \in (m^{**}, m^*). \end{aligned}$$

That is, the higher m^* (so the more likely there is to be a challenger), the less likely the incumbent is to be elected. This is because, if no challenge is mounted, the incumbent wins for sure, whereas, if there is a challenge, there is some chance that the incumbent loses.

As before, we define the incumbency advantage as the weighted average of the incumbency advantages for each party, with weights given by the probability that each party wins an open-seat election. Recall that the incumbency advantage for party p is the difference between the probability that an incumbent from party p wins reelection and the probability that a candidate from party p wins the open-seat election. In the baseline model, the probability of a party winning the open-seat election was $1/2$, the same for each party. In the current framework, the probability that a party wins the open-seat election depends on the voter's partisan leanings. Of course, these partisan leanings also affect the probability that an incumbent is reelected.

As in the previous section, we can use the standard convolution formula for the sum of independent random variables to write the probability that an incumbent wins as

$$\int V(m, m^*, \bar{\eta}) f^\gamma(m) dm, \quad (2)$$

where f^γ is the distribution of the posterior means (derived formally in Appendix A.1). Using this formula, we have the following (see Appendix A.2 for the proof).

Proposition 6 1. *The incumbency advantage exists—the probability an incumbent wins is greater than the probability the incumbent's party wins the open-seat election.*

2. *The incumbency advantage is continuous in the entry cost (κ) and the partisan bias (γ).*
3. *The incumbency advantage is increasing in the entry cost, κ .*

As before, the existence of the incumbency advantage is driven by two effects: *selection* and *insulation*. The insulation effect is now augmented (relative to the baseline model) because sufficiently strong incumbents ($m > m^*$) deter a challenger from entering and so are fully insulated. The incumbency advantage is increasing in the cost of challenger entry because this entry deterrence component of the insulation effect grows smoothly with this cost.

Do Costs of Challenger Entry Change the Comparative Statics?

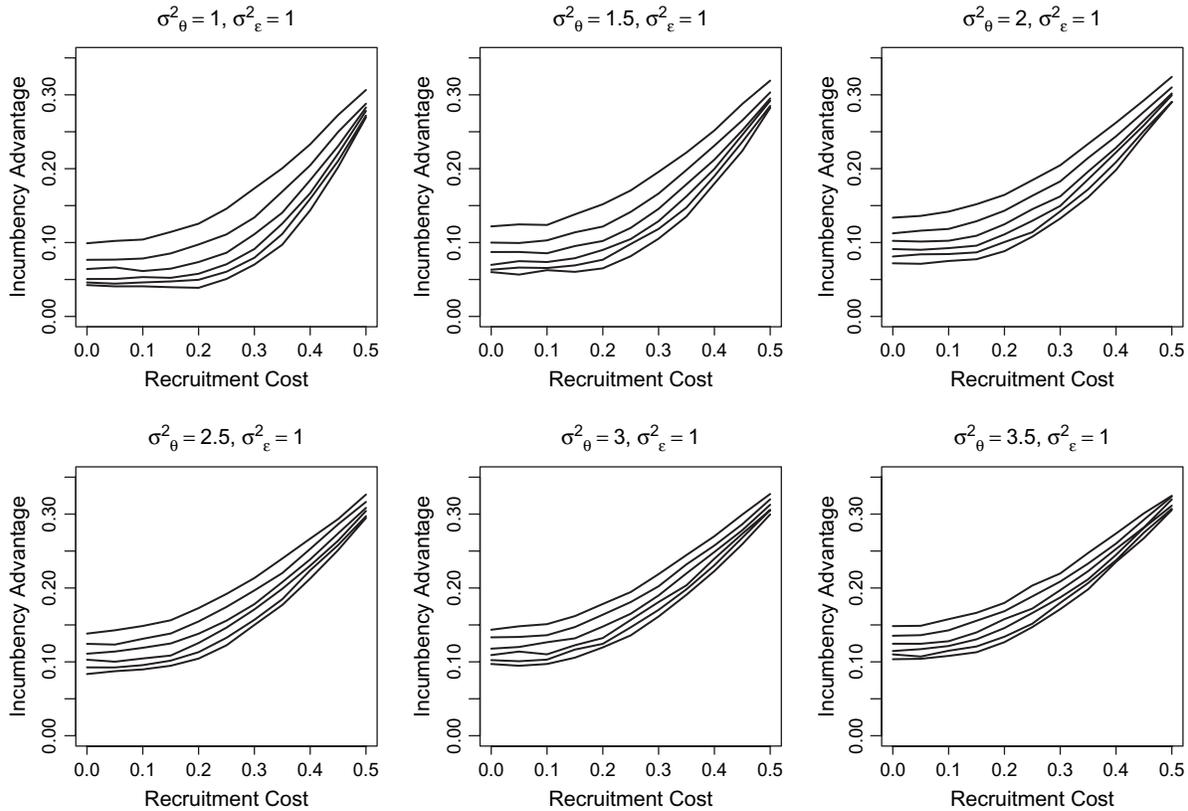
In the baseline model we derived our comparative statics under the restrictive assumption that there were no costs of challenger entry. Proposition 6 demonstrates that the incumbency advantage is continuous in the costs of entry. Hence, the comparative statics clearly hold for sufficiently small entry costs. The question remains as to what happens when the costs of entry are large.

To see what is at issue, recall that all of the comparative statics in the baseline model followed the same basic logic—some exogenous change in a parameter increases the incumbency advantage because it enhances both electoral selection and insulation. But when recruitment costs are great enough, the changes in these effects start to work in opposing directions. The reason is that, for sufficiently high costs, most incumbents *who face a challenge* will run as underdogs. To see this most clearly, consider the case of maximal costs of entry ($\kappa = 1/2$). In this case, there will be a challenge if and only if the incumbent's expected ability net of partisan bias is less than the *prior* mean of zero (i.e., $m - \bar{\eta} < 0$). Thus, for maximal costs, all challenged incumbents will run as underdogs because high costs deter potential candidates from challenging those incumbents who would run as favorites. For slightly lower costs, most challenged incumbents will run as underdogs.

Unlike an incumbent running as a favorite, an incumbent who runs as an underdog benefits from increased variability in the outcome of the election. If the status quo favors the challenger, anything that decreases the likelihood that the outcome will reflect the status quo will help the incumbent.¹³ Thus, when recruitment costs are sufficiently high, changes in

¹³It is important to note that we are not arguing that the incumbency advantage can decrease as the costs of entry increase. (Indeed, the incumbency advantage increases because more incumbents run uncontested.) Rather, we are arguing that, as the costs of entry increases, the other comparative statics might change sign.

FIGURE 1 Simulated comparative statics on κ and σ_η^2 , with $\gamma = 0$. Each frame represents a vector of parameter values $(\sigma_\theta^2, \sigma_\varepsilon^2)$. The x -axis is the recruitment cost (κ) and the y -axis is the incumbency advantage. Thus, moving along the x -axis is equivalent to increasing the cost of mounting a challenge. Each separate curve in a given cell represents a value for σ_η^2 (the variability of the policy component of the voter’s decision). The value of σ_η^2 is decreasing as the curves increase on the vertical axis. A reversal of the comparative statics on σ_η^2 would involve these curves crossing, which they never do for the parameter values explored.



parameter values that used to increase both selection and insulation, now increase selection but decrease insulation.

Increasing electoral selection tends to increase the incumbency advantage. But decreasing insulation tends to decrease the incumbency advantage. Thus, these effects are in tension with one another. If the magnitude of the decrease in insulation is larger than the magnitude of the increase in electoral selection, our overall comparative statics will be reversed.

Unfortunately, we cannot give a simple analytic characterization of when the comparative statics might reverse. However, we explore this question computationally, through a series of simulations.

As Figure 1 makes clear, our main comparative static (on σ_η^2) is robust to the inclusion of a large recruitment cost. Each cell in the figure represents a different vector of parameters. As one moves up the y -axis within a cell, each curve represents a lower value of σ_η^2 —that is, the incumbency advantage is decreas-

ing in σ_η^2 , just as in Lemma 2. Moreover, for the parameter values explored here, the comparative statics never reverse, even for very high values of κ (i.e., the curves within a cell never cross). Thus, at least for these parameter values, the increase in insulation associated with an increase in σ_η^2 (for high values of κ) is never large enough to dominate the decrease in electoral selection associated with an increase in σ_η^2 .

It is, of course, possible to find cases where the decrease in insulation does become strong enough to reverse a comparative static. The easiest way to do this is to allow the variance of the two election signals to differ. Recall that the variance of the second signal ($\sigma_{\varepsilon_2}^2$) only affects insulation (since selection happens in the first election). As a result, for low values of κ a move from a small to a large $\sigma_{\varepsilon_2}^2$ increases the incumbency advantage by increasing insulation. However, for high values of κ , this comparative static reverses—making the second election less informative actually decreases the incumbency advantage by

decreasing insulation.¹⁴ This reversal of comparative statics is illustrated in Figure 2.

Districts with Partisan Leanings

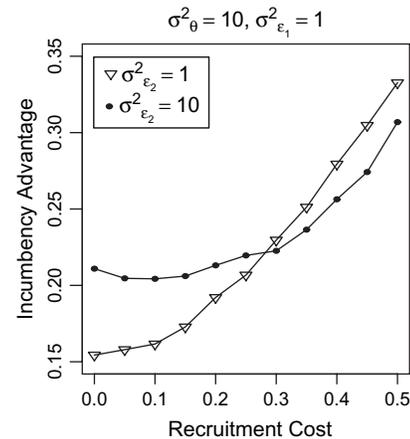
In this section, we explore one more comparative static of the quality difference model: changes in partisan balance. What happens to the magnitude of the incumbency advantage when the electorate becomes more evenly divided between the two parties? Intuitively, as the voter becomes more likely to be indifferent between the two parties on the policy dimension, the probability that he will make his decision based on quality increases, leading to a stronger quality difference based incumbency advantage. This is similar to the intuition of Proposition 2, but we will see that persistent partisanship (as opposed to short-run partisan tides) introduces new complications. To keep the model tractable, we return to the assumption that $\kappa = 0$.

Although the formula for the incumbency advantage in Appendix A.1 is complicated, we can build some intuition about how the incumbency advantage changes with shifts in the partisan balance of the electorate. In particular, we will ask what happens to the incumbency advantage for a left-wing incumbent when the electorate becomes more right-wing (γ increases). Since the problem is symmetric for left- and right-wing candidates, this will also tell us about right-wing incumbents.

A rightward shift in the voter’s partisan bias has three effects on the incumbency advantage for a left-wing candidate, formalized in the three following lemmata. The first effect is that a left-wing candidate has a harder time winning the open-seat election the more right-wing is the voter. This means that left-wing candidates only win open-seat elections in right-wing districts if they can convince the voter they are of particularly high ability. Thus, the more right-wing the district, the more electoral selection left-wing candidates face, and the higher the voters believe a left-wing incumbent’s ability to be. This increase in selection tends to increase the left-wing incumbency advantage.

¹⁴This does not necessarily mean that the substantive comparative static—the incumbency advantage should be greater in more visible offices—reverses. Increasing the visibility of the office means decreasing both variances, a case where our simulations do not reveal a reversal. In addition, it is probably easier for parties to recruit acceptable candidates for less visible offices, meaning κ should increase with visibility. This has the direct effect of increasing the incumbency advantage, along with the indirect effect of potentially reversing the comparative statics on the signals. We thank an anonymous referee for the last point.

FIGURE 2 For low values of κ the incumbency advantage is increasing in the second signal’s variance. For high values of κ the incumbency advantage is decreasing in the second signal’s variance.



Lemma 3 *The distribution of left-wing incumbent abilities is better (in the sense of first-order stochastic dominance) the more right leaning is the electorate (higher γ).*

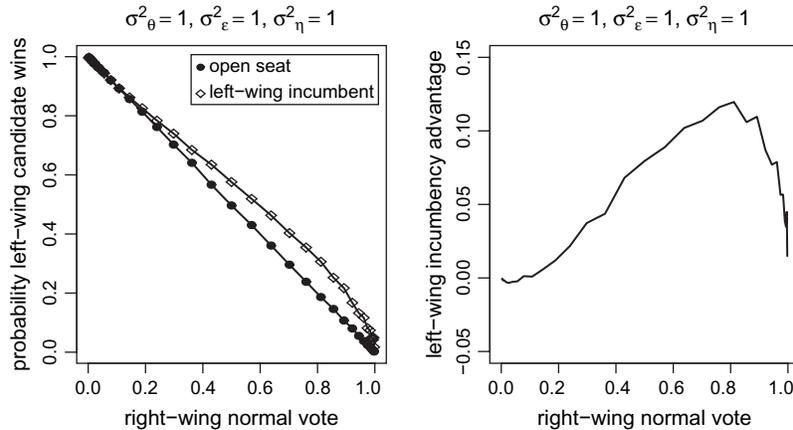
Lemma 3 identifies the effect of partisan leanings on electoral selection. All else equal, an increase in electoral selection, like that identified in Lemma 3, implies an increase in the incumbency advantage. However, in the current model, increased partisanship has an effect not only on selection, but on the two probabilities in the definition of the incumbency advantage. These probabilities are the probability a left-wing candidate wins an open-seat election and the probability of a given left-wing incumbent being reelected (i.e., insulation). The next two lemmata describe the effects on these probabilities.

As the electorate becomes more right-wing, the probability the left-wing party wins the open-seat election decreases. The left-wing incumbency advantage is the probability a left-wing incumbent wins minus the probability the left-wing party wins the open-seat election. Hence, this second effect also increases the incumbency advantage.

Lemma 4 *The more right-wing the voter, the lower the probability the left-wing party wins the open-seat election.*

The third effect is that the left-wing incumbent in a more right-wing district faces a more right-wing voter on average. Thus, as the district becomes more right-wing, the probability a left-wing incumbent

FIGURE 3 The left-hand panel shows how changes in partisan balance affect the simulated probability a left-wing candidate wins as an incumbent and in an open seat election. The right-hand panel shows how partisan balance affects the left-wing incumbency advantage—the difference between the two lines from the left-hand panel. The horizontal axis measures partisan balance as the probability that the right-wing candidate wins the open seat election. By a slight abuse of the standard terminology, we call this the “normal vote”—since, in our model, the probability of winning and the expected vote share have the same comparative statics, this should not cause any confusion.



wins reelection decreases, holding her expected ability constant. This decrease in insulation tends to reduce the incumbency advantage.

Lemma 5 Consider a fixed distribution of left-wing incumbent expected abilities. The more right-wing is the voter, the less likely is a left-wing incumbent to win reelection.

These results show that changing the partisanship of the voter by making him more right-wing has competing effects on the left-wing incumbency advantage. On the one hand, it increases the expected quality of left-wing incumbents by increasing electoral selection (Lemma 3) and decreases their chance of winning the open seat (Lemma 4)—both of which work to increase the left-wing incumbency advantage. On the other hand, it diminishes insulation (Lemma 5)—working to decrease the left-wing incumbency advantage. Unfortunately, we cannot give a complete analytic account of how these competing effects balance out to affect the incumbency advantage. However, we can use simulations to build on the analysis above.

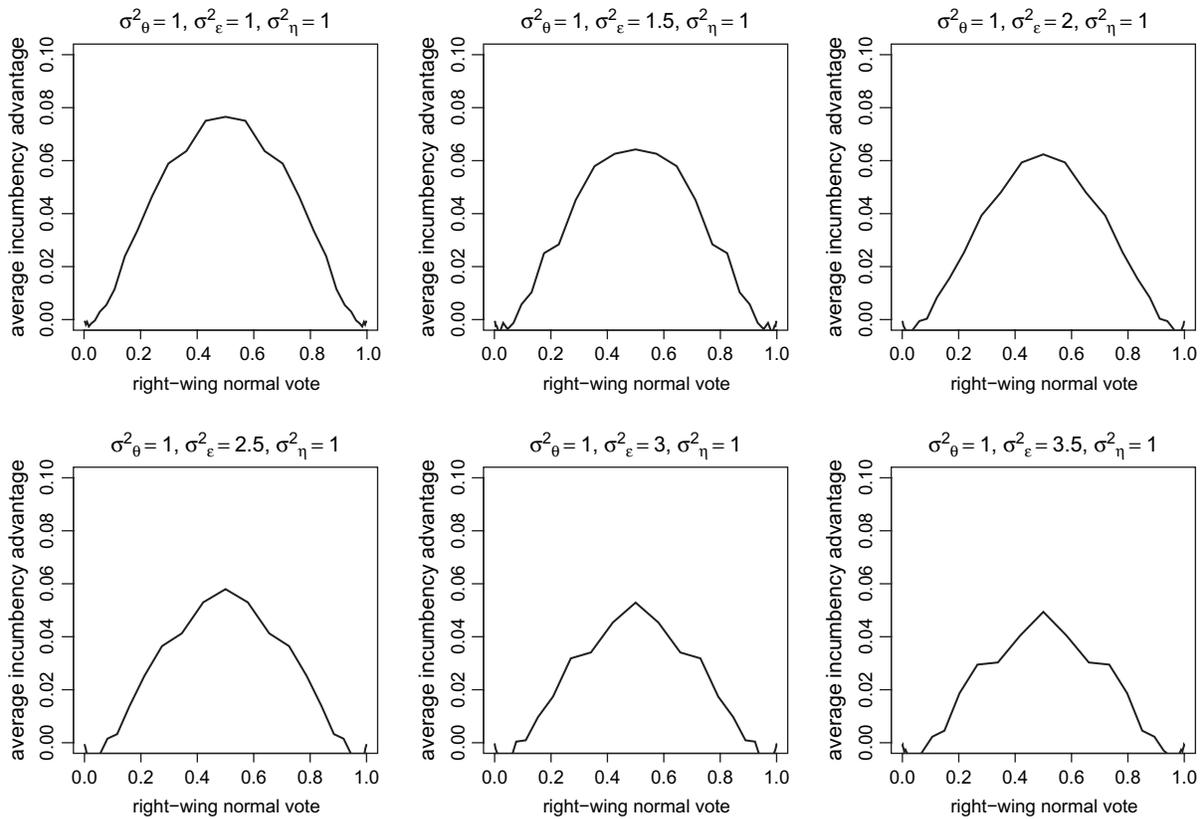
The three effects are illustrated in simulations reported in Figure 3.¹⁵ Consistent with Lemma 4, the simulated probability that the left-wing candidate wins the open-seat election decreases as the voter becomes more right-wing (the line marked with solid

circles in the left-hand panel is decreasing). Consistent with Lemma 5, the simulated probability that a left-wing incumbent wins reelection (i.e., the level of insulation) also decreases as the voter becomes more right-wing (the curve marked with open diamonds in the left-hand panel is also decreasing). However, the curve marked by open diamonds in the left-hand panel of the figure, which represents the probability of an incumbent winning, also shows the selection effect from Lemma 3—the more right-wing the electorate, the higher the expected ability of a left-wing incumbent. Hence, in the figure, expected ability is not constant as partisanship changes, rather ability is increasing with partisanship. Thus, the curve marked with open diamonds has a shallower slope than the line marked with solid circles.

The vertical distance between the two curves represents the left-wing incumbency advantage for a given level of partisanship—the probability a left-wing incumbent wins minus the probability a left-wing candidate wins an open-seat election. This difference is plotted in the right-hand panel of the figure. Figure 3 shows that the left-wing incumbency advantage is maximized with a moderately right-wing voter. This maximum is achieved where the decrease in insulation associated with increased right-wing partisanship begins to more than compensate for the positive effects of stronger electoral selection and a smaller normal vote. The right-wing incumbency advantage, of course, is the mirror image. Thus the

¹⁵The R code used for all simulations is available at <http://home.uchicago.edu/~bdm>.

FIGURE 4 Simulated comparative statics on σ_ε^2 and γ . Each frame represents a vector of parameter values $(\sigma_\theta^2, \sigma_\varepsilon^2, \sigma_\eta^2)$. The x -axis is the normal vote for a right-wing candidate. Thus, moving along the x -axis is equivalent to changing the level of partisan balance (γ), where the electorate's partisan preferences are exactly balanced between the two parties at the midpoint.



right-wing incumbency advantage is maximized in a moderately left-wing district.

The overall incumbency advantage is a weighted average of the left- and right-wing incumbents' advantages, weighted by the probability of each party winning an open-seat election. As Lemma 4 shows, left-wing candidates will be more likely to win an open-seat election in left-leaning districts and right-wing candidates will be more likely to win an open-seat election in right-leaning districts. Thus, in calculating the overall incumbency advantage, the majority of weight in any given district will be put on the candidate whose incumbency advantage is increasing in partisan balance. This suggests that the overall incumbency advantage will be increasing in partisan balance and will be maximized at a normal vote of one half for each party—i.e., in a perfectly balanced electorate.

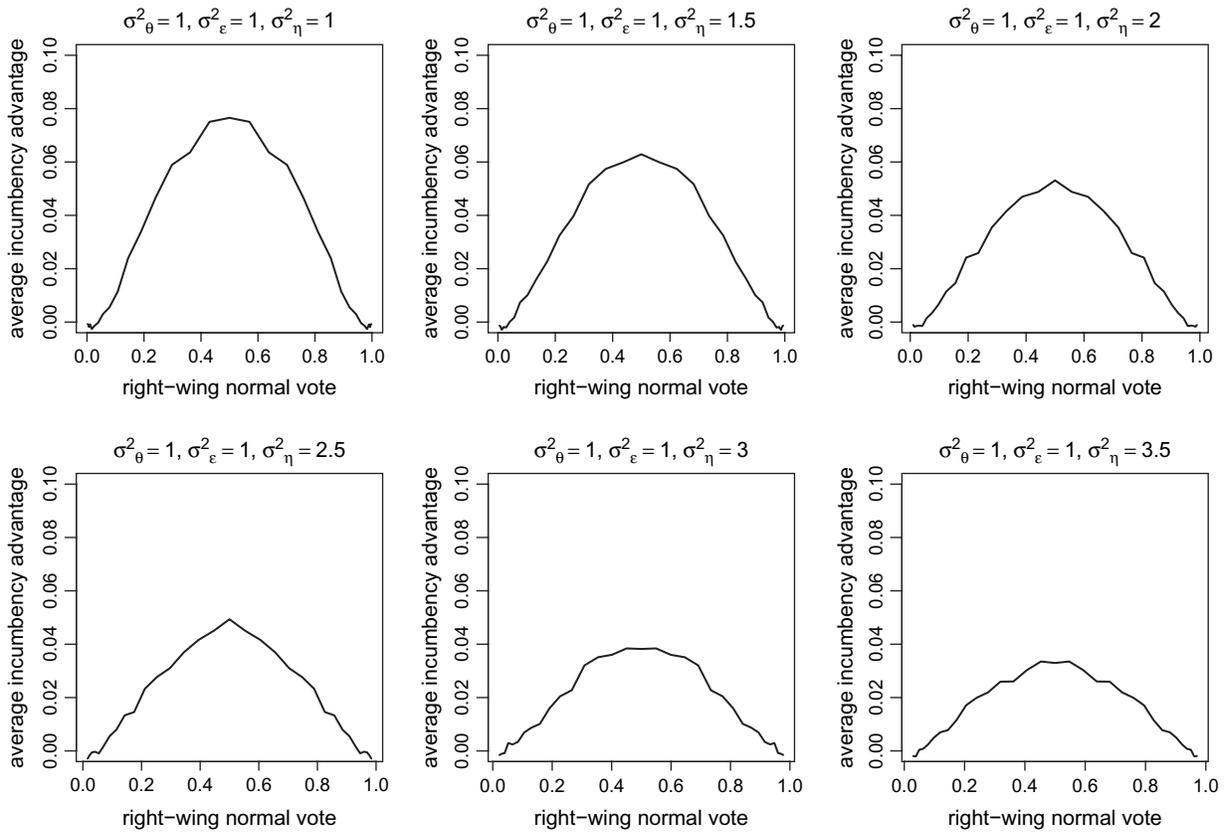
This intuition is confirmed in the simulations reported in Figures 4 and 5. The incumbency advantage is increasing as the electorate's normal vote moves closer to one half, i.e., the partisan balance increases. This can be seen in each frame of both Figures 4 and 5. As expected, for all vectors of parameter values that

we have investigated, the incumbency advantage (the y -axis) is increasing as partisan balance increases.

The simulations also serve as a check on the robustness of our earlier comparative statics on the visibility of the office (σ_ε^2), expected size of partisan tides ($\sigma_{x^*}^2$), and polarization ($\mu_R - \mu_L$). The comparative statics for the visibility of office can be seen by moving across the panels of Figure 4, and the comparative statics for both partisan tides and polarization can be seen by moving across the panels of Figure 5 (since $\eta = 2(\mu_R - \mu_L)x^*$). In each case the simulated comparative statics in the more general model are the same as the analytically derived results from the baseline model. The incumbency advantage is decreasing as the visibility of office decreases (Figure 4) and as party polarization and the size of partisan tides increase (Figure 5).

Given that we had to use computational methods to explore the interaction of all three effects, one might wonder whether the effect of partisan balance on the incumbency advantage is actually caused by the electoral selection effect identified in Lemma 3 or whether it is driven by the nonselection based effects

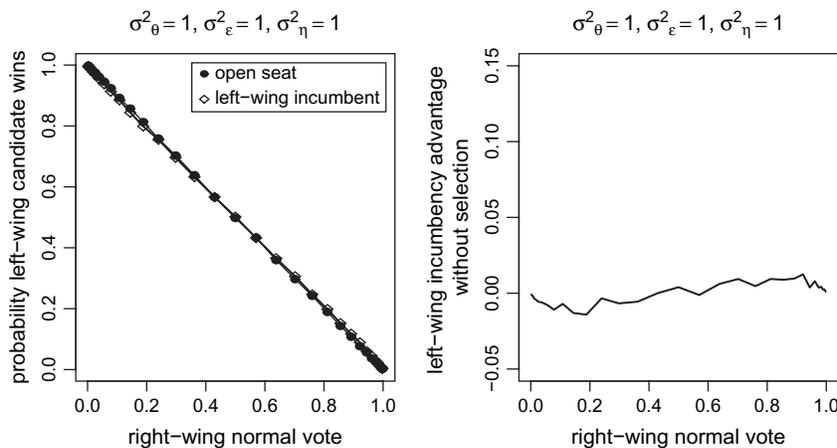
FIGURE 5 Simulated comparative statics on σ_η^2 and γ .



identified in Lemmata 4 and 5. In order to investigate this possibility, we simulate a version of the model that includes the nonselection effects, but excludes electoral selection. This is done by assuming that,

after the signals of candidate ability are revealed in the open-seat election, the winner is chosen at random. Thus, there is no selection for ability. These simulations are shown in Figure 6.

FIGURE 6 The left-hand panel shows the effect of partisan balance on the probability a left-wing incumbent wins, if there is no selection on quality, and the probability a left-wing candidate wins an open seat. The difference between these (shown in the right-hand panel) shows how partisan balance affects the left-wing incumbency advantage given the effects identified in Lemmata 4 and 5, but without electoral selection. Notice that the right-hand panel has the same vertical scale as the right-hand panel of Figure 3.



The line marked by dark circles represents the simulated probability that a left-wing candidate wins the open-seat election, as a function of partisan leanings. This probability is unchanged in this altered model (thus, the curve looks the same as in Figure 3). The curve marked by open diamonds represents how partisan balance affects the probability that the incumbent wins reelection, in the absence of electoral selection. The downward slope of the curve reflects the decreased insulation effect described in Lemma 5—for any distribution of ability, it is more difficult for a left-wing incumbent to win reelection the more right-wing the electorate. The fact that this curve has a steeper slope than the analogous curve in Figure 3 reflects the absence of electoral selection.

The right-hand side of Figure 6 shows how partisan balance impacts the left-wing incumbency advantage given these two effects, but without electoral selection. As is evident from the figure, removing the selection effect essentially eliminates any systematic impact of partisan balance on the incumbency advantage. Thus, although we are unable to prove it analytically, the computational results indicate fairly decisively that decreased partisan balance increases the incumbency advantage because of electoral selection.

The result in Figure 3 is consistent with cross-sectional, district-level data for the House of Representatives. In particular, Ansolabehere, Snyder, and Stewart report that “the more highly partisan a district, the smaller is the incumbency advantage enjoyed by an incumbent of the favored party” (2000, 27).¹⁶ Erikson and Wright (2001) find further evidence for this claim, particularly in their Figure 6.

Ansolabehere and Snyder (2002) also identify significant changes over time in the incumbency advantage in the United States. In particular, for all statewide elected offices, the average incumbency advantage was small to nonexistent at midcentury, grew dramatically in the 1960s and 1970s, and has since stabilized at a relatively high level (8–10%).

Traditional explanations for the increase in the incumbency advantage for the House of Representatives over this period stress changes in the nature of Congressional politics. Specific examples include constituency service (Cain, Ferejohn, and Fiorina 1987) and redistricting (Cox and Katz 2002). While these are

important parts of the story for the House, Ansolabehere and Snyder (2002) show that the incumbency advantage increased for all statewide offices, including those, such as governor or state auditor, for which arguments based on legislative politics do not apply. Given this, it is worth speculating on whether or not the quality difference model can also contribute to understanding the overtime trends.

The possibility that the quality difference model can explain the increase in the incumbency advantage for statewide offices is based on the following intuition. At midcentury, most states were relatively uncompetitive. In terms of the model, the “solid south” was a group of states with very low γ (they leaned heavily Democratic), while most of the north and west had very high γ (they leaned heavily Republican). Since the 1960s, the partisan balance of these states has increased (Erikson, Wright, and McIver 1993; Ansolabehere and Snyder 2002). Voters in the south became more willing to consider electing Republicans to statewide offices and voters in the north and west became more willing to consider electing Democrats. Our model suggests that future empirical work should explore the possibility that this decrease in pure partisan voting might have increased selection on quality, thereby strengthening the incumbency advantage.¹⁷

It is also worth noting that one potential complication with this stories is that, prior to the demise of the solid south, competitive primaries in Democratic election may have played a similar role to general elections in the more balanced, later era. Thus, selection and strategic entry may still have been occurring, at an earlier stage in the electoral process (Key 1949). In a recent paper, Ansolabehere, et al. (2005) examine the evidence for this mechanism on the universe of statewide elections between 1908 and 2004. While they do find that the fraction of primaries that were contested declined as the states became more balanced, they also find that there was never a period without heterogeneity across elections in the extent to which primaries were contested. For example, in the south, even at the high point in the 1930s, only just over 60% of races with an incumbent featured contested primaries. And the numbers for the south

¹⁶There are two possible interpretations of this finding. First, “old voters” may have a particular proclivity for their incumbents, having nothing to do with learning about quality, that “new voters” do not have. Second, and more consistent with our model, “old voters” may have information about their incumbents that “new voters” do not. Thus, “old voters” may simply believe that the incumbent is higher quality than do “new voters.”

¹⁷An important subtlety is that the district-level heterogeneity that we model (γ) need not necessarily reflect ideology. It can simply be interpreted as the level of commitment to the party due to some valence term. Thus, when we describe the solid south as having γ s to the left of zero, this need not imply that the south was “liberal” relative to other parts of the country. Rather, it implies that the south was more reliably committed to the Democratic party. All of our results are consistent with this more nuanced interpretation.

declined to about one third by the 1950s. Thus, while the primary mechanism will certainly moderate the comparative statics about partisan balance, the data suggest that contested primaries have rarely been frequent enough to completely suppress the effect.

Conclusion

We derived the comparative statics of a model of the incumbency advantage based on electoral selection and strategic challenger entry. This contributes to the literature on the incumbency advantage in two ways. First, we have derived a variety of new testable implications of the quality-difference-based model. Second, we discussed how the quality based model might account for over time trends in the incumbency advantage. It is important to have an account, such as ours, that addresses changes in the incumbency advantage without reference to legislative politics because, as Ansolabehere and Snyder (2002) point out, trends in the incumbency advantage in the United States hold for all statewide offices.

The model also suggests some natural extensions. As mentioned in the previous section, including the selection and entry dynamics that occur in primaries could add an additional, important layer to the quality difference story, as could giving potential candidates informative signals about the realization of the partisan shock prior to entry. Another interesting extension would add incumbents' strategic behavior while in office, thereby relating our account to the literature on the personal vote (Cain, Ferejohn, and Fiorina 1987). Ashworth and Bueno de Mesquita (2006) model voter learning about candidate ability and the provision of constituency service or pork by incumbents attempting to manipulate that learning. That model has only one election, so there is no possibility of an incumbency advantage. However, they show that the changes in electoral environment that lead to an increased incumbency advantage in this paper also lead to increased constituency service or pork (and thus, an increased personal vote) in that model.

The question remains what would happen to the comparative statics in a model that simultaneously included constituency service or pork and enough elections to give rise to a quality-difference-based incumbency advantage. We have some reason to believe that our comparative statics would persist. In particular, Ashworth(2005) studies a model with electoral selection, candidate effort, and committee assignments. That model does not include strategic entry, and it does not explore the comparative statics we study here.

However, it does find that when effort and ability are additively separable, the dynamics of the incumbency advantage are the same as they would be in a model, like the one studied here, without effort.

Appendix A. Proofs and Derivations

We will focus on L incumbents and use the symmetry of the problem to deduce R incumbency advantages.

A.1. Preliminary Results

Updating. The updating rule for voter beliefs follow from standard results on Bayes' rule with a normal prior and normal signal (DeGroot 1970).

In the first election, candidate c has the ability distribution $\mathcal{N}(m_c, \nu_1)$, where $\lambda_1 = \frac{\sigma_\theta^2}{\sigma_\theta^2 + \sigma_{\varepsilon_1}^2}$, $m_c = \lambda_1 s_c$ and $\nu_1 = \lambda_1 \sigma_{\varepsilon_1}^2$. In the second election, we need to separately account for the updating on the incumbent and on the challenger. The voter believes that the incumbent's ability is distributed $\mathcal{N}(m_{2,inc}, \nu_{inc})$, where $\lambda_{inc} = \frac{\nu_1}{\nu_1 + \sigma_{\varepsilon_2}^2}$, $m_{2,inc} = \lambda_{inc} s_{inc} + (1 - \lambda_{inc}) m_{1,inc}$ and $\nu_{inc} = \lambda_{inc} \sigma_{\varepsilon_2}^2$. Similarly, the voter believes that the challenger's ability is distributed $\mathcal{N}(m_{chall}, \nu_{chall})$, where $\lambda_{chall} = \frac{\sigma_\theta^2}{\sigma_\theta^2 + \sigma_{\varepsilon_2}^2}$, $m_{chall} = \lambda_{chall} s_{chall}$ and $\nu_{chall} = \lambda_{chall} \sigma_{\varepsilon_2}^2$. We will drop subscripts when no confusion will result.

The Normal Vote and the Incumbency Advantage.

Next we provide formal versions of Definitions 1 and 2. Let τ be the variance of $m_L - m_R - \eta$ in the open seat election and τ_- be the variance of $m_R + \eta$ in the open seat election. Then an L candidate wins the open-seat election with probability $1 - \Phi(\bar{\eta}/\tau)$, and an R candidate wins with complementary probability.

Recalling that $\bar{\eta}$ is a function of γ , write $\mathcal{I}\mathcal{A}_p(\gamma)$ for the incumbency advantage of an incumbent from party p when the voter has expected ideal point γ :

$$\mathcal{I}\mathcal{A}_L(\gamma) = \Pr(L \text{ wins} \mid \gamma, L \text{ an incumbent}) \\ - (1 - \Phi(\bar{\eta}/\tau))$$

and

$$\mathcal{I}\mathcal{A}_R(\gamma) = \mathcal{I}\mathcal{A}_L(-\gamma).$$

Then we can write the overall incumbency advantage as

$$\mathcal{I}\mathcal{A}(\gamma) = (1 - \Phi(\bar{\eta}/\tau))\mathcal{I}\mathcal{A}_L(\gamma) + \Phi(\bar{\eta}/\tau)\mathcal{I}\mathcal{A}_R(\gamma).$$

All that's left is to calculate $\Pr(L \text{ wins} \mid \gamma, L \text{ an incumbent})$.

Let f^γ denote the density of posterior means for L candidates who have won reelection once. By Bayes's rule, this is given by

$$\begin{aligned} f^\gamma(m_L) &= \frac{\Pr(L \text{ won} \mid m_L) \Pr(m_L)}{\Pr(L \text{ won})} \\ &= \frac{\Phi\left(\frac{m_L - \bar{\eta}}{\tau_-}\right)}{1 - \Phi\left(\frac{\bar{\eta}}{\tau}\right)} \frac{1}{\sigma_m} \phi\left(\frac{m_L}{\sigma_m}\right). \end{aligned}$$

To derive these expressions, recall that the L candidate wins if and only if $m_L \geq \eta + m_R$. Since $\eta + m_R \sim \mathcal{N}(\bar{\eta}, \tau_-)$, we see that

$$\Pr(L \text{ won} \mid m_L) = \Phi\left(\frac{m_L - \bar{\eta}}{\tau_-}\right).$$

Similarly, the unconditional probability is

$$\Pr(L \text{ won}) = \Phi\left(\frac{-\bar{\eta}}{\tau}\right).$$

Now consider the second election, and consider an incumbent who enters the reelection campaign with expected ability m . She wins reelection unless a challenger runs and the innovation to the belief about her ability pushes her below $\eta + m_{chall} \sim \mathcal{N}(\bar{\eta}, \sigma_\eta^2 + \lambda_{chall}\sigma_\theta^2)$. Since the new posterior for the incumbent is distributed $\mathcal{N}(m, \lambda_{inc}\nu_{inc})$, she wins with probability $1 - \Phi\left(\frac{\bar{\eta} - m}{\sigma}\right)$, where

$$\sigma = \sqrt{\lambda_{inc}\nu_{inc} + \sigma_\eta^2 + \lambda_{chall}\sigma_\theta^2}.$$

An unchallenged incumbent wins with probability 1, so the probability that an incumbent with posterior m wins reelection is $V(m, m^*, \bar{\eta})$ (defined in the text). To calculate the overall probability an incumbent wins, average the probability conditional on m over the distribution of incumbent abilities, f^γ . Since the innovation to the voter's belief about the incumbent is independent of m ,¹⁸ we can use the standard convolution formula for the sum of independent random variables to write the probability that an incumbent wins as

$$\int V(m, m^*, \gamma) f^\gamma(m) dm. \tag{3}$$

¹⁸This is a standard result on the normal distribution.

A.2. Proof of Proposition 6

We proceed in two steps. First, we prove the result for the case where $\kappa = 0$. Second, we show that the incumbent's reelection probability is increasing in κ , which establishes the result in general. The continuity results follow trivially from the characterization.

Say that density g dominates density h in the likelihood ratio (MLR) order if, for all $x > z$,

$$\frac{g(x)}{g(z)} > \frac{h(x)}{h(z)}.$$

This implies that the distributions conditional on any interval are ordered by first-order stochastic dominance.¹⁹ Fixing $m > m'$, the likelihood ratio for the incumbent's ability is

$$\frac{\Phi((m - \bar{\eta})/\tau_-)\phi(m/\sigma_m)}{\Phi((m' - \bar{\eta})/\tau_-)\phi(m'/\sigma_m)} > \frac{\phi(m/\sigma_m)}{\phi(m'/\sigma_m)},$$

where the inequality follows from monotonicity of the normal cdf. Thus the distribution of posteriors conditional on election likelihood ratio dominates the distribution of posteriors not conditioned on reelection.

Now, consider the reelection probability at the second election under the (counterfactual) assumption that the incumbent ability distribution is just the posterior distribution for a single candidate, unconditioned on victory. In this case, the distribution of posteriors is $\Phi(m/\sigma_m)$, and the reelection probability at the second election is

$$\begin{aligned} Prob &= \int \frac{1}{\sigma_m} \phi\left(\frac{m}{\sigma_m}\right) \left(1 - \Phi\left(\frac{\bar{\eta} - m}{\sigma}\right)\right) dm \\ &= \int \int \frac{1}{\sigma_m} \phi(m/\sigma_m) \frac{1}{\sigma} \phi((y - m)/\sigma) dm dy \\ &= \int \frac{1}{\tau} \phi(y/\tau) dy \\ &= 1 - \Phi\left(\frac{\bar{\eta}}{\tau}\right), \end{aligned}$$

where the third equality follows from the facts that $\sigma_m^2 + \sigma^2 = \tau^2$ and that the integrand in the second line is the convolution giving the density of the sum of two normals.

¹⁹See the appendix of Krishna (2002) for more discussion of the likelihood ratio order.

Above, we showed that the actual distribution of posteriors MLR-dominates the no selection distribution we considered above. Since the MLR order is a subset of the FOSD order and $1 - \Phi(-m/\sigma)$ is increasing in m , the incumbent’s probability of winning is greater than her party’s probability of winning an open-seat election.

Now we show that the advantage is increasing in κ . The advantage in the general case is

$$\int_{-\infty}^{m^*(\kappa)} \Phi\left(\frac{m - \bar{\eta}}{\sigma}\right) f^\gamma(m) dm + \int_{m^*(\kappa)}^{\infty} f^\gamma(m) dm - \left(1 - \Phi\left(\frac{\bar{\eta}}{\tau}\right)\right).$$

Differentiate with respect to κ to get

$$\left(\Phi\left(\frac{m^*(\kappa) - \bar{\eta}}{\sigma}\right) - 1\right) f^\gamma(m^*(\kappa)) \frac{dm^*(\kappa)}{d\kappa} > 0,$$

where the inequality follows from the fact that m^* is decreasing in κ .

Appendix B. Comparative Statics in the Baseline Model

To handle the baseline model, we make extensive use of the following:

Lemma 6 *Assume that $\kappa = 0$ and fix a distribution f so that $f(x) > f(-x)$ for all $x > 0$. Then the integral in equation (3) is decreasing in σ .*

Proof Differentiate the reelection probability to get

$$\begin{aligned} \frac{d}{d\sigma} \text{Pr} &= \int -\phi\left(\frac{-m}{\sigma}\right) \cdot \frac{m}{\sigma^2} f(m) dm \\ &= \frac{-1}{\sigma^2} \left(\int_{-\infty}^0 m\phi(-m/\sigma) f(m) dm \right. \\ &\quad \left. + \int_0^{\infty} m\phi(-m/\sigma) f(m) dm \right). \end{aligned}$$

In the last line, the first integrand is negative, while the second is positive. Thus we will be done as soon as we show that the absolute value of the second integral is greater than the absolute value of the first.

We have

$$\begin{aligned} &\left| \int_{-\infty}^0 m\phi(-m/\sigma) f(m) dm \right| \\ &< \int_{-\infty}^0 |m\phi(-m/\sigma) f(m)| dm \\ &= \int_{-\infty}^0 -m\phi(-m/\sigma) f(m) dm \\ &= \int_0^{\infty} m\phi(m/\sigma) f(-m) dm \\ &= \int_0^{\infty} m\phi(-m/\sigma) f(-m) dm \\ &< \int_0^{\infty} m\phi(-m/\sigma) f(m) dm, \end{aligned}$$

where the first equality is the definition of the absolute value, the second equality is a change of variable, the third equality is from the symmetry of the normal density, and the last inequality is from the hypothesis on f . □

B.1. Proof of Lemma 2

Let $f_{\sigma_\eta^2}$ be the density of the date 1 posterior means (given reelection at date 1). Observe that the (pre-election) density of date 0 posteriors is $(1/\sigma_m)\phi(\cdot/\sigma_m)$. The postelection density at x is thus proportional to $(1/\sigma_m)\phi(x/\sigma_m)(1 - \Phi(-x/\sigma_\eta))$. Since the normal density is symmetric about 0, we can rewrite this as $(1/\sigma_m)\phi(x/\sigma_m)\Phi(x/\sigma_\eta)$.

Lemma 7 *The density $f_{\sigma_\eta^2}$ is indexed by $-\sigma_\eta$ in the MLRP sense.*

Proof Fix $x > x'$. The likelihood ratio is

$$\mathcal{LR} = \frac{\phi(x/\sigma_m)\Phi(x/\sigma_\eta)}{\phi(x'/\sigma_m)\Phi(x'/\sigma_\eta)}.$$

We need to show that this is decreasing in σ_η . Differentiate to see that

$$\begin{aligned} \text{sgn}\left(\frac{\partial \mathcal{LR}}{\partial \sigma_\eta}\right) &= (-x/\sigma_\eta^2)\phi(x/\sigma_\eta)\Phi(x'/\sigma_\eta) \\ &\quad - (-x'/\sigma_\eta^2)\phi(x'/\sigma_\eta)\Phi(x/\sigma_\eta), \end{aligned}$$

so $\partial \mathcal{LR} / \partial \sigma_\eta < 0$ if and only if

$$-x \frac{\phi(x/\sigma_\eta)}{\Phi(x/\sigma_\eta)} < -x' \frac{\phi(x'/\sigma_\eta)}{\Phi(x'/\sigma_\eta)}.$$

This inequality follows from $x > x'$ and the logconcavity of ϕ . \square

We will use the following result from Athey (2002): If $f(x, y)$ is logsupermodular in x and y and $g(y, z)$ is logsupermodular in y and z , then $h(x, z) = \int f(x, y)g(y, z) dy$ is logsupermodular in x and z .

Lemma 8 For any $\sigma_\eta^2, f_{\sigma_\eta^2}(x) > f_{\sigma_\eta^2}(-x)$.

Proof If $x > x'$, we can use the monotonicity of the normal cdf to get

$$\frac{\phi(x/\sigma_m)(1 - \Phi(-x/\sigma_\eta))}{\phi(x'/\sigma_m)(1 - \Phi(-x'/\sigma_\eta))} > \frac{\phi(x/\sigma_m)}{\phi(x'/\sigma_m)},$$

so the postelection date 0 posterior MLR dominates the preelection posterior distribution. Thus the Lemma from Athey implies that the convolution of the postelection density with the posterior innovation density MLR dominates the convolution of the preelection density with the innovation density. But this second convolution yields a mean-zero normal density. Since this density is symmetric, the result follows. \square

Because the MLR order is a subset of the FOSD order and $1 - \Phi(-m/\sigma)$ is increasing in m , Lemma 6 implies that the reelection probability is decreasing in σ_η^2 , proving the proposition.

B.2. Proof of Proposition 4

We will proceed in three steps. First, we will derive the comparative statics with respect to $\sigma_{e_2}^2$. Second, we will prove a stochastic order result for order statistics and use it to derive the comparative statics with respect to $\sigma_{e_1}^2$. Third, we will consider the case when $\sigma_{e_1}^2 = \sigma_{e_2}^2$.

Since this proposition refers to the case where $\sigma_\eta^2 = 0$, the standard deviation σ is a monotone transformation of $\lambda_{inc}\nu_{inc} + \lambda_{chall}\sigma_\theta^2$. Notice that both λ s are decreasing in $\sigma_{e_2}^2$. Thus σ is decreasing in $\sigma_{e_2}^2$, and Lemma 6 will imply that the incumbency advantage is increasing in $\sigma_{e_2}^2$ as soon as we verify the hypothesis on f . Since $\sigma_\eta^2 = 0$, the winner of the first election is just the candidate with greater expected ability, so the density of incumbent abilities is the density of the first-order statistic:

$$f(m) = \frac{2}{\sigma_m} \Phi\left(\frac{m}{\sigma_m}\right) \phi\left(\frac{m}{\sigma_m}\right).$$

The symmetry of the normal density implies $\phi(m/\sigma_m) = \phi(-m/\sigma_m)$, and the monotonicity of the normal cdf implies that $\Phi(m/\sigma) > \Phi(-m/\sigma)$ when $m > 0$. Combining these facts gives us $f(m) > f(-m)$ when $m > 0$.

Now we turn to the comparative statics for $\sigma_{e_1}^2$. Before turning to the details, we will outline the proof for this case. If we let $f(m; \sigma_{e_1}^2) = \frac{1}{\sigma_m} \Phi(m/\sigma_m) \phi(m/\sigma_m)$ and $h(m; \sigma_{e_1}^2) = \Phi(m/\sigma)$, then the incumbent's probability of winning is

$$\mathcal{P}(\sigma_{e_1}^2) = \int h(m; \sigma_{e_1}^2) f(m; \sigma_{e_1}^2) dm.$$

Fix $\overline{\sigma_{e_1}^2} > \sigma_{e_1}^2$. We will prove the claim in the proposition by establishing the following chain of inequalities:

$$\begin{aligned} \mathcal{P}(\overline{\sigma_{e_1}^2}) &= \int h(m; \overline{\sigma_{e_1}^2}) f(m; \overline{\sigma_{e_1}^2}) dm \\ &> \int h(m; \overline{\sigma_{e_1}^2}) f(m; \sigma_{e_1}^2) dm \\ &> \int h(m; \overline{\sigma_{e_1}^2}) f(m; \sigma_{e_1}^2) dm = \mathcal{P}(\overline{\sigma_{e_1}^2}). \end{aligned}$$

The first inequality will follow from Lemma 6 and the fact that σ is increasing in $\sigma_{e_1}^2$. The second inequality will follow from the fact that σ_m is decreasing in $\sigma_{e_1}^2$ and the following result.

Lemma 9 Let h have a first derivative that is positive and symmetric about 0, in the sense that $h'(m) = h'(-m)$. Then the integral $\frac{2}{\sigma_m} \int h(m) \Phi(m/\sigma_m) \phi(m/\sigma_m) dm$ is increasing in σ_m .

Proof Integrate by parts to get

$$\begin{aligned} &\frac{2}{\sigma_m} \int h(m) \Phi(m/\sigma_m) \phi(m/\sigma_m) dm \\ &= h(m) \Phi(m/\sigma_m)^2 \Big|_{-\infty}^{\infty} - \int h'(m) \Phi(m/\sigma_m)^2 dm \\ &= 1 - \int_0^{\infty} h'(m) (\Phi(m/\sigma_m) + \Phi(-m/\sigma_m)) dm, \end{aligned}$$

where the second equality uses the symmetry of h' . Differentiate to get

$$\begin{aligned} & \frac{d}{d\sigma_m} (\Phi(m/\sigma_m) + \Phi(-m/\sigma_m)) \\ &= \cdot - 2\Phi(m/\sigma_m)\phi(m/\sigma_m) \frac{m}{\sigma_m^2} \\ & \quad + 2\Phi(m/\sigma_m)\phi(m/\sigma_m) \frac{m}{\sigma_m^2} \\ &= \frac{2m}{\sigma_m^2} \phi(m/\sigma_m) (\Phi(-m/\sigma_m) - \Phi(m/\sigma_m)) < 0, \end{aligned}$$

where the inequality follows from the monotonicity of Φ . Since h' is positive, this implies that the overall integral is increasing in σ_m . \square

To complete this argument, we calculate the derivatives of the relevant variances. First, $\sigma_m = \frac{\sigma_\theta^4}{\sigma_\theta^2 + \sigma_{\varepsilon_1}^2}$ is clearly decreasing in $\sigma_{\varepsilon_1}^2$. Next, we have

$$\begin{aligned} \frac{d}{d\sigma_{\varepsilon_1}^2} v_{inc} &= \frac{d}{d\sigma_{\varepsilon_1}^2} \frac{\sigma_\theta^2 \sigma_{\varepsilon_1}^2}{\sigma_\theta^2 + \sigma_{\varepsilon_1}^2} \\ &= \frac{\sigma_\theta^2 (\sigma_\theta^2 + \sigma_{\varepsilon_1}^2) - \sigma_\theta^2 \sigma_{\varepsilon_1}^2}{(\sigma_\theta^2 + \sigma_{\varepsilon_1}^2)^2} \\ &= \frac{\sigma_\theta^4}{(\sigma_\theta^2 + \sigma_{\varepsilon_1}^2)^2} > 0, \end{aligned}$$

so v_{inc} is increasing in $\sigma_{\varepsilon_1}^2$. Finally,

$$\begin{aligned} \frac{d}{dv_{inc}} \lambda_{inc} v_{inc} &= \frac{d}{dv_{inc}} \frac{v_{inc}^2}{v_{inc} + \sigma_{\varepsilon_2}^2} \\ &= \frac{2v_{inc} (v_{inc} + \sigma_{\varepsilon_2}^2) - v_{inc}^2}{(v_{inc} + \sigma_{\varepsilon_2}^2)^2} > 0. \end{aligned}$$

$$\frac{d}{d\bar{\eta}} \frac{\Phi((m - \bar{\eta})/\tau_-)}{\Phi((m' - \bar{\eta})/\tau_-)} = \frac{-(1/\tau_-)\phi((m - \bar{\eta})/\tau_-)\Phi((m' - \bar{\eta})/\tau_-) + (1/\tau_-)\phi((m' - \bar{\eta})/\tau_-)\Phi((m - \bar{\eta})/\tau_-)}{(\Phi((m' - \bar{\eta})/\tau_-))^2}.$$

Combined with the previous derivative, this implies that $\lambda_{inc} v_{inc}$, and thus σ , is increasing in $\sigma_{\varepsilon_1}^2$.

Finally, we consider the case of $\sigma_{\varepsilon_1}^2 = \sigma_{\varepsilon_2}^2$. From the previous steps, it suffices to show that σ decreases with the common value of σ_ε^2 . After imposing the equality and differentiating, substitution and rearranging (available on request) yield

$$\begin{aligned} \frac{\partial}{\partial \sigma_\varepsilon^2} \sigma^2 &= \frac{\sigma_\theta^4}{(\sigma_\theta^2 + \sigma_\varepsilon^2)^2 (\frac{\sigma_\theta^2 \sigma_\varepsilon^2}{\sigma_\theta^2 + \sigma_\varepsilon^2} + \sigma_\varepsilon^2)^2} \\ &\cdot (-2\sigma_\varepsilon^4) < 0. \end{aligned}$$

B.3. Proof of Proposition 5

Since the median voter is decisive in the open-seat election, changes in σ_y^2 do not affect f . Given this, the proposition follows immediately from Lemma 6.

Appendix C. Comparative Statics for Partisan Bias

Recall that γ affects the incumbency advantage only through $\bar{\eta} = 2(x_R - x_L)\gamma$, so the comparative statics with respect to γ have the same sign as those with respect to $\bar{\eta}$.

C.1. Proof of Lemma 3

The likelihood ratio is

$$\frac{f^\gamma(m)}{f^\gamma(m')} = \frac{\Phi((m - \bar{\eta})/\tau_-)\phi(m/\sigma_m)}{\Phi((m' - \bar{\eta})/\tau_-)\phi(m'/\sigma_m)}.$$

The second factor on the RHS does not depend on $\bar{\eta}$, so we can find out what happens as $\bar{\eta}$ changes by differentiating the first factor on its own. We have

This is positive if and only if

$$\begin{aligned} & \phi((m' - \bar{\eta})/\tau_-)\Phi((m - \bar{\eta})/\tau_-) \\ & > \phi((m - \bar{\eta})/\tau_-)\Phi((m' - \bar{\eta})/\tau_-) \end{aligned}$$

if an only if

$$\frac{\phi((m' - \bar{\eta})/\tau_-)}{\Phi((m' - \bar{\eta})/\tau_-)} > \frac{\phi((m - \bar{\eta})/\tau_-)}{\Phi((m - \bar{\eta})/\tau_-)}$$

which is true since Φ is logconcave. Thus increasing $\bar{\eta}$ improves the distribution of posteriors in the MLR sense.

C.2. Proof of Lemma 4

The normal vote for party L is $1 - \Phi(\bar{\eta}/\tau)$, which is decreasing in $\bar{\eta}$.

C.3. Proof of Lemma 5

The probability that an L incumbent wins is $\int f(m)(1 - \Phi((\bar{\eta} - m)/\sigma)) dm$, which is decreasing in $\bar{\eta}$.

Acknowledgment

We are indebted to Andrew Martin for teaching us how to do simulations. We also benefited from the comments of Amanda Friedenber, Roger Myerson, Tom Palfrey, Howard Rosenthal, Steve Smith, and seminar audiences at Kellogg, Yale, and the PIER Conference on Political Economics.

Manuscript submitted 30 June 2006

Manuscript accepted for publication 13 October 2007

References

- Ansolabehere, Stephen, and James M. Snyder, Jr. 2002. "The Incumbency Advantage in U.S. Elections: An Analysis of State and Federal Offices, 1942–2000." *Election Law Journal* 1 (3): 315–38.
- Ansolabehere, Stephen, James M. Snyder, Jr., and Charles-Stewart, III. 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, John Mark Hansen, Shigeo Hirano, and James M. Snyder, Jr. 2005. "The Decline of Competition in U.S. Primary Elections, 1908–2004." MIT. Unpublished paper.
- Ashworth, Scott. 2005. "Reputational Dynamics and Political Careers." *Journal of Law, Economics and Organization* 21 (2): 441–66.
- Ashworth, Scott, and Ethan Bueno de Mesquita. 2006. "Delivering the Goods: Legislative Particularism in Different Electoral and Institutional Settings." *Journal of Politics* 68 (1): 169–79.
- Ashworth, Scott, and Ethan Bueno de Mesquita. 2007. "Valence Competition and Platform Divergence." University of Chicago. Unpublished Paper.
- Athey, Susan. 2002. "Monotone Comparative Statics Under Uncertainty." *Quarterly Journal of Economics* CXVII (1): 187–223.
- Banks, Jeffrey S., and Rangarajan K. Sundaram. 1998. "Optimal Retention in Agency Problems." *Journal of Economic Theory* 82: 293–323.
- Cain, Bruce, John Ferejohn, and Morris Fiorina. 1987. *The Personal Vote: Constituency Service and Electoral Independence*. Cambridge, MA: Harvard University Press.
- Cox, Gary W., and Jonathan N. Katz. 1996. "Why Did the Incumbency Advantage in U.S. House Elections Grow?" *American Journal of Political Science* 40 (2): 478–97.
- Cox, Gary W., and Jonathan N. Katz. 2002. *Elbridge Gerry's Salamander: The Electoral Consequences of the Reapportionment Revolution*. New York: Cambridge University Press.
- DeGroot, Morris H. 1970. *Optimal Statistical Decisions*. New York: McGraw-Hill.
- Diermeier, Daniel, Michael Keane, and Antonio Merlo. 2005. "A Political Economy Model of Congressional Careers." *American Economic Review* 95 (1): 347–73.
- Erikson, Robert S., and Gerald C. Wright. 2001. Voters, Candidates, and Issues in Congressional Elections. In *Congress Reconsidered*, ed. Lawrence C. Dodd and Bruce I. Oppenheimer (7th ed.). Washington, DC: Congressional Quarterly Press, 67–95.
- Erikson, Robert S., Gerald C. Wright, and John P. McIver. 1993. *Statehouse Democracy: Public Opinion and Policy in the American States*. Cambridge: Cambridge University Press.
- Gerber, Alan. 1998. "Estimating the Effect of Campaign Spending on Senate Election Outcomes Using Instrumental Variables." *American Political Science Review* 92 (2): 401–11.
- Gordon, Sanford C., Gregory A. Huber, and Dimitri Landa. 2007. "Challenger Entry and Voter Learning." *American Political Science Review* 101 (2): 303–20.
- Gowrisankaran, Gautam, Matthew F. Mitchell, and Andrea Moro. 2008. "Electoral Design and Voter Welfare from the U.S. Senate: Evidence from a Dynamic Selection Model." *Review of Economic Dynamics* 11 (January): 1–17.
- Jacobson, Gary C. 1980. *Money in Congressional Elections*. New Haven, CT: Yale University Press.
- Jacobson, Gary C., and Samuel Kernell. 1983. *Strategy and Choice in Congressional Elections, 2nd Edition*. New Haven, CT: Yale University Press.
- Key, V.O., Jr. 1949. *Southern Politics in State and Nation*. New York: A. A. Knopf.
- Krishna, Vijay. 2002. *Auction Theory*. San Diego: Academic Press.
- Londregan, John and Thomas Romer. 1993. Polarization, Incumbency, and the Personal Vote. In *Political Economy: Institutions, Competition, and Representation*, ed. William A. Barnett, Melvin J. Hinich, and Norman J. Schofield. Cambridge: Cambridge University Press, 355–77.
- Samuelson, Larry. 1984. "Electoral Equilibria with Restricted Strategies." *Public Choice* 43 (3): 307–27.
- Samuelson, Larry. 1987. "A Test of the Revealed-Preference Phenomenon in Congressional Elections." *Public Choice* 54 (2): 141–69.
- Stone, Walter J., L. Sandy Maisel and Cherie D. Maestas. 2004. "Quality Counts: Extending the Strategic Politician Model of Incumbent Deterrence." *American Journal of Political Science* 48 (3): 479–95.
- Zaller, John. 1998. Politicians as Prize Fighters: Electoral Selection and the Incumbency Advantage. In *Politicians and Party Politics*, ed. John G. Geer. Baltimore, MD: Johns Hopkins University Press, 125–85.

Scott Ashworth is assistant professor of politics, Princeton University, Princeton, NJ 08544. Ethan Bueno de Mesquita is associate professor of public policy, University of Chicago, Chicago, IL 60637.