When Do Elections Promote Policy Moderation by Privately-Informed Officials?[1]

Brandice Canes-Wrone[2] and Kenneth W. Shotts[3]

September 2005

[1]Prepared for delivery at the Wallis Institute's Annual Conference on Political Economy, University of Rochester, September 30 – October 1, 2005. We are grateful for helpful comments from Larry Bartels, Craig Volden, and participants in seminars at Harvard, Princeton, Stanford, Washington University, and Yale.

[2]Associate Professor of Politics and Public Affairs, Woodrow Wilson School, Princeton University. 214 Robertson Hall, Princeton NJ 08544. Phone: (609) 258-9047. Email: bcwrone@princeton.edu.

[3]Associate Professor, Stanford Graduate School of Business. 518 Memorial Way, Stanford CA 94305-5015. Phone (650) 725-4510. Email: kshotts@stanford.edu.

Abstract

The conventional wisdom is that elections induce ideological moderation in winner-take-all, single-member districts. Yet with rare exception, research does not examine how politicians' private information may affect their incentives for moderation. This omission is unfortunate given that politicians can have incentives to hide their ideological preferences, and that representative democracy is predicated upon the notion that elected officials will learn more about particular policies than will the typical voter. The following paper develops a theory based on private information about policy and preferences, and shows that under many circumstances elections will not promote ideological moderation. The key reason this occurs is that a politician's incentive to signal that he shares voters' ideological leanings can outweigh any electoral motivations to choose the policy that will most likely produce an outcome voters desire.

Since Downs's (1957) influential model of electoral competition, the conventional wisdom has been that elections induce ideological moderation in winner-take-all, single-member districts. As Ansolabehere, Snyder, and Stewart (2001, 137) surmise, this perspective forms the basis "for most contemporary theorizing about representative politics in the United States." Morton (forthcoming, 86) in her text *Analyzing Elections* supports this view, maintaining that "the majority rule nature of U.S. elections rewards candidates who moderate." Likewise, Besley and Coate (2003) highlight that this perspective guides research on elected regulators.

Of course (and as these scholars recognize), there exists research that suggests candidates will not converge to identical platforms that represent the median voter. Theory suggests that divergence can result from the pressures of parties (e.g., Snyder and Ting 2002), interest groups (e.g., Baron 1994), candidates' policy motivations (e.g., Calvert 1985), and valence issues (e.g., Ansolabehere and Snyder 2000; Aragones and Palfrey 2002; Groseclose 2001), among other factors.[1] This divergence does not, however, imply that candidates lack electoral incentives for moderation. For instance, in the work on personal motivations, a liberal candidate takes positions that are more liberal than those of the median voter, but elections still induce him to behave more like a centrist than he otherwise would.[2] More generally, the work on divergence indicates that even though candidates may not adopt identical policy positions, their need to appeal to voters will have a moderating impact.

The pervasiveness of this notion notwithstanding, scant attention has been paid to the possibility that the incentives for moderation may differ when a politician has private information,

[1] Grofman (2004) provides a review of the literature that predicts nonconvergence. He highlights that this work indicates electoral competition will produce "centripetal" pressures even when divergence occurs.

[2] Even in Callander (2004), where candidates have an electoral incentive not to converge fully to the median voter, elections cause them to take more centrist positions than they would otherwise.

or information the public lacks.[3] The following paper addresses the omission by analyzing the incentives for moderation by a re-election-seeking official who has private information about the expected consequences of policies and his ideological leanings. Voters have prior beliefs about which policy decision would most likely serve their interests, but know that the official has better, albeit not perfect information. Moreover, voters recognize that the official may have ideological leanings that encourage him to make policy decisions they might not desire if fully informed. We analyze two different scenarios concerning the public's knowledge at the time of the election. In the first, the electorate observes which policy the official has chosen but has no information about the policy consequences. This type of assumption has been adopted in models of economic cycles (e.g., Rogoff and Sibert 1988), and is particularly applicable to policy decisions directed towards long-term consequences. In the second scenario, the electorate knows not only which policy was chosen, but also can learn consequences of it.

The major result is that in a wide range of circumstances elections do not promote ideological moderation. When the electorate will only observe the policy choice but not its consequences, there are no incentives for moderation; in this circumstance, there cannot be pressures for both liberal and conservative officials to behave in a more centrist manner than they would in the absence of elections. The reason for this is that voters will reward either liberal behavior or conservative behavior and thereby push all politicians in one direction.

When the outcome of the policy choice can be observed, moderation is still not an electoral incentive under many conditions. In fact, the equilibrium results are surprisingly similar to the case where the outcome cannot be observed. For instance, when a district is somewhat conservative but would like the official to be willing to take a liberal position if his policy information recommends doing so, the official's desire to please voters can induce him to behave in a more conservative

[3] The major exception is Duggan (2000), which we discuss in detail below.

manner than he or the public, if it were fully informed, would like. It is because the public does not have the incumbent's policy information or know his preferences with certainty that these incentives arise.[4]

This equilibrium outcome—where the official takes an action that seems desirable to voters, but that they would not want if fully informed—generates a second result worth highlighting. In particular, the theory shows that the inability of a politician to convey his ideology to voters can create incentives for *pandering*, where the official chooses a popular policy that voters would not prefer if they had his policy information. As such, the theory contributes to a growing literature on the circumstances under which politicians have this incentive (Canes-Wrone, Herron, and Shotts 2001; Maskin and Tirole 2004; Fox 2005; Prat 2005).[5] In Canes-Wrone, Herron, and Shotts and in Prat, an incumbent will pander to signal that he is highly competent at obtaining policy-relevant information. In Maskin and Tirole and in Fox, like the theory subsequently developed, a politician will pander in an attempt to signal that his preferences are like those of the voter. A distinctive feature of our results is their relationship to moderation. Here we show how pandering can make an incumbent less ideologically moderate; the other theories do not address this issue.

The paper proceeds as follows. Section one reviews the relevant theoretical literature. Section two discusses key components of the theory. Sections three through five develop the formal model. They describe, respectively, the set-up, the case where voters do not learn about policy outcomes before the election, and the case where outcomes can be observed. Section six concludes

[4] Later we detail how the concept of moderation applies to situations where politicians are assumed to have private policy information.

[5] Maskin and Tirole (2004) define pandering more broadly to include any time an incumbent's electoral incentives affect his policy choices. We use the narrower definition, which is also used in Canes-Wrone, Herron, and Shotts (2001), because the broader definition includes incumbent behavior that voters would support if fully informed, and we do not want to label this behavior with the pejorative term "pandering." Pandering occurs in Maskin and Tirole (2004) even using the definition employed here.

by discussing implications of the results for theories of representation.

Related Literature

The most similar theory is Duggan's (2000) model of repeated elections in which an incumbent has private information about his policy preferences. Duggan shows that even with this assumption, electoral pressures from pivotal centrist voters induce moderation, where both liberal and conservative politicians have incentives to behave more like centrists.[6] However, unlike the theory developed here, Duggan does not incorporate the possibility that officials may have policy expertise, i.e., better information than voters have about expected policy consequences.

Various theories allow politicians to have policy expertise and private information about their preferences, but even the ones most similar to ours cannot examine the electoral incentives for moderation (Downs and Rocke 1994; Maskin and Tirole 2004; Fox 2005).[7] Downs and Rocke are interested in specifying a voting rule such that an incumbent will have an incentive to behave in a centrist fashion; they do not assess whether such a voting rule is actually optimal for the electorate or whether it would arise in equilibrium. In other words, Downs and Rocke simply assume that elections induce moderation when officials are privately informed, rather than analyzing whether this is indeed the case.

The other two papers cannot tackle this issue for a different reason. Politicians in these theories

[6] Enelow and Hinich (1981) develop a model that does not incorporate private information (of either policy consequences or politicians' preferences), but does allow for voter uncertainty concerning what politicians will do if elected to office. In that model there is convergence, with candidates locating at the position preferred by the median.

[7] Other papers that incorporate policy expertise by an official with private information about his preferences include Austen-Smith (1992) and Coate and Morris (1995). Austen-Smith focuses on how the informational asymmetries affect legislators' incentives to vote sophisticatedly. Coate and Morris examine how the asymmetries influence the efficiency of government transfers to interest groups.

can have only two types of preferences, and thus there is no way to distinguish centrism from movement in a particular direction. In Maskin and Tirole, the voters' objective is to avoid re-electing a bad type of politician whose preferences are opposed to theirs, e.g., a politician who wants to legalize a pesticide whenever voters want it banned and who wants to ban it whenever voters want to legalize it. In Fox, the voters try to avoid re-electing a dogmatic type of politician who always prefers the same action regardless of his private information, e.g., one who always wants to go to war.[8] In neither theory is there ideological tension that could result in different extremes moving towards a moderating center. Thus no existing theory analyzes how politicians' policy expertise may affect their incentives for ideological moderation.[9]

Modeling Issues

Preference Divergence

The key tension in the model is that an official's preferences may diverge from those of voters. As in other models that assume officials have private information about their preferences and expected policy consequences, we assume that there are two states of the world and two policy choices, where the correct choice is the one that matches the state of the world (Downs and Rocke 1994; Maskin and Tirole 2004; Fox 2005). For instance, an anti-terrorist policy that curbs civil liberties may or may not be necessary to prevent large-scale terrorist attacks, and the choice is whether to approve

[8] There are other electoral models that incorporate moral harzard and adverse selection but can not analyze the question at hand (e.g., Canes-Wrone, Herron and Shotts 2001; Fearon 1999). For instance, the Canes-Wrone, Herron, and Shotts model assumes an official has policy expertise, but because there is no ideological tension the theory cannot examine the issue of moderation.

[9] Harrington (1993) assumes that voters and politicians each have private information about their priors on which policy will produce a better outcome; policy expertise is not assumed. Harrington shows that in this circumstance elections (weakly) induce moderation to the median voter.

the policy. Preference divergence arises from the fact that the actors may place different weights on the two types of errors–a failure to approve a necessary policy versus mistakenly approving an unhelpful one. If the true state of the world were known, all actors would prefer the same policy choice–everyone would like to prevent large-scale terrorist attacks, and no one wants to curb civil liberties. However, because some players are more concerned about minimizing the risks of terrorism and others are more concerned about the negative side effects, they may weigh the two types of errors differently.

Information

Actors in the model do not know the true state of the world–e.g., no one knows for certain whether a given anti-terrorist policy is necessary or whether a new drug should be approved. Voters know the prior probability that a proposed policy is good, and an elected official has an additional signal. Although this signal is not perfect, it gives the official more information than the public has. As a result, it is possible that the official's information could indicate that a popular policy is not the one fully informed voters would want, and voters know this.

Yet since an official and voters may weigh possible policy errors differently, an official may wish to choose policies that run counter to what fully informed voters would want. This is why preference divergence creates the key tension in the model; voters cannot simply trust an official to use his policy expertise to advance their interests. Continuing with the example of anti-terrorist policy, there are two types of officials that a voter may worry about: a type who is too willing to curb civil liberties, and a type who is too reluctant to do so. The model allows for both types.

These informational assumptions relate to the concept of moderation in the theory; importantly, this definition must account for the fact that officials have different policy information than the public does. We say that *moderation* occurs when at least some officials who prefer a left-leaning policy regardless of their signal will instead follow it if it recommends a right-leaning policy, and at least some officials who prefer a right-leaning policy regardless of their signal will instead follow the signal if it recommends a left-leaning policy. Consider the case where officials must decide whether to allow a pesticide to be used for farming, and the officials have a signal regarding the safety of the pesticide. Moderation occurs if elections influence the behavior of both left-leaning and right-leaning officials, e.g., officials who prefer to ban the pesticide regardless of the signal are electorally induced to legalize it when the signal indicates that it is safe, whereas officials who prefer to legalize regardless of the signal are electorally induced to ban the pesticide when the signal indicates that it is unsafe.

Notably, this definition of moderation differs from the concept of responsiveness to voters' interests. An official is not characterized as moderate if he caters to voters who want him to ignore his policy information. Thus, for example, we do not characterize behavior as moderate if a politician represents a district position that a pesticide should always be legalized regardless of the safety data. Importantly, however, this definitional distinction does not matter for the main results; even if we defined moderation as responsiveness to voters' interests, the theory suggests that elections often cause politicians to be less moderate than they otherwise would be.

Electoral Competition

When deciding to re-elect or remove an incumbent, voters know the policy choice and also may know about the success or failure of this policy. Solving the model both with and without the possibility of such policy feedback follows Maskin and Tirole (2004) and Fox (2005). By allowing for both possibilities, we learn that the observability of outcomes is a necessary condition for elections to induce policy moderation.

Like numerous theories of electoral competition (e.g., Ferejohn 1986; Harrington 1993; Fearon

1999; Duggan 2000) we assume that the challenger and incumbent come from the same pool of potential candidates.[10] Thus if the electorate decides to replace the incumbent, the challenger could be either to the "right" or to the "left" of the incumbent. While this assumption is common, we note that it may make the theory more applicable to certain types of policy contexts than others. The most direct application is a policy decision for which an incumbent could plausibly be replaced from the right or the left.[11] For instance, a chief executive with a war prerogative could be replaced by a challenger who is either more pro- or anti-war. Likewise, an elected prosecutor could be replaced by someone who is more or less punitive. Consistent with these examples, models of war (Downs and Rocke 1994) and elected prosecutors (Gordon and Huber 2002) presume that all office-holders are drawn from the same pool of candidates.

The Model

The model has two time periods. In each there is one policy decision and the state of the world is $\omega \in \{A, B\}$. The periods are independent and the prior probability of A in each is $\pi \in (0, 1)$. At the beginning of a given period the elected official receives a private signal $s \in \{A, B\}$ about the state of the world, where the quality of the signal is $q = \Pr(s = \omega) \in \left(\frac{1}{2}, 1\right)$. Given this information, the official must choose a policy $x \in \{A, B\}$ where A represents taking an action and B represents inaction.

There are three actors in the model: a voter, incumbent official, and challenger, all of whom are policy motivated and prefer $x = \omega$ over $x \neq w$. The actors can differ, however, in their degree of concern over each type of error (choosing $x = B$ when $\omega = A$ versus choosing $x = A$ when

[10] One of the few models that does not make this assumption is Meirowitz (2004).

[11] Clearly, the theory is also applicable to the decisions of officials selected through nonpartisan elections, which are prevalent in the U.S. in state judicial and city council contests as well as in other countries.

$\omega = B$). The relative importance that each actor places on choosing $x = A$ when $\omega = A$ is denoted by $\alpha \in [0,1]$, subscripted α_V for the voter, α_I for the incumbent, and α_C for the challenger. If the correct outcome occurs, i.e., $x = \omega$, then the actor receives zero utility. If $x = B$ when $\omega = A$ then the actor receives utility $-\alpha$. On the other hand if $x = A$ when $\omega = B$ the actor receives utility $-(1-\alpha)$. An actor's total utility from the game is the sum of the utility in the two periods:

$$\begin{aligned} U &= -\alpha \{\text{Total number of mistaken } B \text{ policy choices}\} \\ &\quad - (1-\alpha)\{\text{Total number of mistaken } A \text{ policy choices}\}. \end{aligned}$$

The voter and official may have divergent preferences, and an official's α is her private information. Specifically $\alpha_V \in [0,1]$ is common knowledge, but α_I is a random draw from a uniform distribution $F : [0,1] \to [0,1]$. The challenger's α_C is also independently drawn from F. We specify an official's strategy as a function of her preference parameter and private signal about the state of the world. For the first period let the incumbent's strategy, which determines whether she chooses $x = A$ or $x = B$, be $\sigma^1(\alpha_I, s) : [0,1] \times \{A,B\} \to \{A,B\}$. Similarly, let $\sigma^2(\alpha, s) : [0,1] \times \{A,B\} \to \{A,B\}$ be the strategy for the second period official.

After the first period the voter must decide either to retain the incumbent or replace her. Let the voter's strategy be ν, a vector that specifies the probability that he re-elects the incumbent in each information set. We consider two variants of the model. In the first variant the voter only observes the policy choice $x \in \{A,B\}$ and decides whether to re-elect simply based on this policy choice. Thus the voter's strategy is $\nu = (\nu_A, \nu_B)$ where $\nu_A \in [0,1]$ denotes the probability that he re-elects the official when $x = A$ and $\nu_B \in [0,1]$ denotes the probability that he re-elects the official when $x = B$.

In the second variant of the model, the voter observes the policy choice $x \in \{A,B\}$ and if the official chooses to take action ($x = A$) the voter also observes whether this action succeeds or fails, i.e. the voter learns the value of ω. If the official chooses not to take action ($x = B$) then the voter only observes x and does not learn ω. Thus the voter's strategy $\nu = (\nu_{AA}, \nu_{AB}, \nu_B)$ must specify behavior in three information sets: ν_{AA} when $x = A$ and $\omega = A$, ν_{AB} when $x = A$ and $\omega = B$, and ν_B when $x = B$. It is worth noting that this assumption that the voter fails to learn ω if the official fails to take action is not critical to our central findings regarding the lack of incentives for moderation. In particular, even if voters learned ω when $x = B$ there would be conditions under which an incumbent would have incentives to make more ideologically extreme policy choices than she otherwise would.

The equilibrium concept employed is Perfect Bayesian. We specify voter beliefs in more detail below, but for now it is worth noting that we do not need to specify voter beliefs off the equilibrium path since each possible outcome occurs with strictly positive probability.

Policy Choice

In both variants of the model, the official's action in the second period depends on her preference parameter α and her beliefs about the probability that A is the state of the world. Let these beliefs, which are straightforward to calculate via Bayes's Rule, be denoted as $\theta^A(A) = \Pr(\omega = A|s = A) = \frac{\pi q}{\pi q + (1-\pi)(1-q)}$ and $\theta^A(B) = \Pr(\omega = A|s = B) = \frac{\pi(1-q)}{\pi(1-q)+(1-\pi)q}$. Given these beliefs, the second period private signal s, and the second period official's preferences α, we can characterize optimal second period behavior.

Proposition 1 (Second Period Policy Choice) *There exist cutpoints $\underline{\alpha}^2$ and $\overline{\alpha}^2$ such that:*

1. $0 < \underline{\alpha}^2 < \overline{\alpha}^2 < 1$

2. *If $\alpha < \underline{\alpha}^2$ then the official chooses $x = B$ regardless of s*

3. *If $\alpha \in \left(\underline{\alpha}^2, \overline{\alpha}^2\right)$ then the official chooses $x = s$*

4. *If $\alpha > \overline{\alpha}^2$ then the official chooses $x = A$ regardless of s.*

As shown in Figure 1, there are three types of officials, categorized according to their second period behavior: a *type* B official always chooses $x = B$ regardless of the signal; a *type* R, or signal-responsive, official always chooses $x = s$; and a *type* A official always chooses $x = A$.

[Figure 1 about here]

It is worth emphasizing that an official's type depends not only on her preferences α but also the quality of her signal q and the ex-ante likelihood π that A is the correct policy. Thus an official who is type B when the quality of her signal is low may be signal-responsive if q is high. At the same time, for a given q and π, type R officials are the only ones who are open to choosing either A or B in the second period on the basis of the policy information. Accordingly, within this set-up, type R officials are more moderate than type A or B ones.

Given these official-types, we also can categorize a voter as being type B, type R, or type A depending on whether $\alpha_V < \underline{\alpha}^2$, $\alpha_V \in \left(\underline{\alpha}^2, \overline{\alpha}^2\right)$, or $\alpha_V > \overline{\alpha}^2$. A type B voter, for instance, wants to elect a type B official for the second period. Since all officials within a given category behave identically in the second period, the voter's expected utility from an official taking office depends upon his beliefs about the probability that the official falls into each of these three categories. Let $\phi_B = F(\underline{\alpha}^2)$ be the probability that a randomly drawn official is type B, let $\phi_R = F(\overline{\alpha}^2) - F(\underline{\alpha}^2)$ be the probability that a randomly drawn official is type R, and let $\phi_A = 1 - F(\overline{\alpha}^2)$ be the probability that the randomly drawn official is type A. These ϕ's are also relevant for the first period official's behavior, since she is policy motivated and cares about the action her replacement will take if she fails to win re-election.

The incumbent's probability of re-election depends on the voter's re-election strategy ν. For an official who observes signal s let $r_A(s)$ denote the probability of re-election if she chooses policy A and let $r_B(s)$ denote her probability of winning if she chooses policy B. In the model without uncertainty resolution these re-election probabilities are determined directly by the voter's strategy after observing policy choice A or B, i.e., for $s \in \{A, B\}$ the re-election probabilities are $r_A(s) = \nu_A$ and $r_B(s) = \nu_B$. With uncertainty resolution, in contrast, an official's probability of re-election after taking action $x = A$ may depend on whether the policy choice succeeded or failed:

$$r_A(s) = \nu_{AA}\theta^A(s) + \nu_{AB}\left(1 - \theta^A(s)\right). \tag{1}$$

Thus the probability of re-election after choosing $x = A$ depends on how likely A is to be the true state of the world given signal s. Since uncertainty is not resolved when she chooses inaction the incumbent's probability of re-election when $x = B$ is $r_B(s) = \nu_B$.

We can now partially characterize first period policy behavior. The incumbent's decision depends not only on first period policy considerations but also second period ones, which are based on the probability that she will be re-elected as well as the behavior of her replacement if she is not. Specifically, for any voter behavior, the incumbent's policy choice can be characterized by cutpoints like the ones for second period behavior.

Proposition 2 (Incumbent's Best Response to Voter Strategy) *For any voter strategy ν there exist cutpoints $\underline{\alpha}^1$ and $\overline{\alpha}^1$ such that in the first period:*

1. $0 \leq \underline{\alpha}^1 < \overline{\alpha}^1 \leq 1$, *and either* $0 < \underline{\alpha}^1$ *or* $\overline{\alpha}^1 < 1$
2. *If $\alpha_I < \underline{\alpha}^1$ then the official chooses $x = B$ regardless of s*
3. *If $\alpha_I \in \left(\underline{\alpha}^1, \overline{\alpha}^1\right)$ then the official chooses $x = s$*
4. *If $\alpha_I > \overline{\alpha}^1$ then the official chooses $x = A$ regardless of s.*

Propositions 1 and 2 enable us to categorize possible effects that voter behavior may have on first period policy choices. One possibility, shown in Figure 2a, is that for both signals $s \in \{A, B\}$ the official has an electoral incentive to choose policy B, i.e., $r_A(s) < r_B(s)$.

[Figures 2a, 2b, and 2c about here]

In this case, $r_A(A) < r_B(A)$ implies that $\underline{\alpha}^1 > \underline{\alpha}^2$, i.e., some signal-responsive or type R officials, who prefer $x = A$ when $s = A$, instead choose $x = B$. Likewise, $r_A(B) < r_B(B)$ implies that $\overline{\alpha}^1 > \overline{\alpha}^2$, i.e., some type A officials, who prefer $x = A$ even when $s = B$, are electorally induced to choose $x = B$. On the flip side, as shown in Figure 2b, if $r_A(s) > r_B(s)$ for both signals $s \in \{A, B\}$ then a first period official always has an incentive to choose $x = A$ so that $\underline{\alpha}^1 < \underline{\alpha}^2$ and $\overline{\alpha}^1 < \overline{\alpha}^2$. Thus in Figures 2a and 2b, some types of officials have an electoral incentive to take more extreme actions than they otherwise would.

A third possibility, shown in Figure 2c, is that $r_A(B) < r_B(B)$ but $r_A(A) > r_B(A)$ so that $\underline{\alpha}^1 < \underline{\alpha}^2$ but $\overline{\alpha}^1 > \overline{\alpha}^2$.[12] In this case officials have electoral incentives to follow their signals when choosing first period policy. As a result, an official never has an incentive to be more extreme and can have an incentive to be moderate than she would be if not running for re-election. In particular, type B officials in the region $\left(\underline{\alpha}^1, \underline{\alpha}^2\right)$ are electorally induced to moderate their behavior by choosing $x = A$ when $s = A$. Likewise, type A officials in the region $\left(\overline{\alpha}^2, \overline{\alpha}^1\right)$ moderate their behavior by choosing $x = B$ when $s = B$.

It is worth noting that moderation here is not complete – an incumbent with $\alpha_I < \underline{\alpha}^1$ or $\alpha_I > \overline{\alpha}^1$ does not behave like a centrist but rather supports a particular policy regardless of her signal. The moderation that occurs here is analogous to moderation in Duggan (2000), in which some types of liberal or conservative officials are induced to alter their behavior and act like centrists.[13] It is also worth noting that Figures 2a-c show how we distinguish the concept of moderation from that of responsiveness. In particular, when the voter is type A, so that he always wants the official

[12] The fourth possibility, $r_A(A) < r_B(A)$ and $r_A(B) > r_B(B)$, which would give the official incentives to go against her signal, cannot occur in equilibrium.

[13] The behaviors are not completely analogous, of course, since Duggan's official does not have policy information, and our official has a discrete rather than continuous set of policies from which to choose.

to choose policy A regardless of her policy information, we do not characterize policy behavior as moderate if the official behaves in this way. At the same time, and as will soon become clear, the lack of incentives for moderation are not limited to type A and type B electorates.

Equilibria Without Uncertainty Resolution

We first analyze the model without uncertainty resolution. The equilibrium depends on the voter's type. Not surprisingly if the voter has a low α_V, and thus is relatively pro-B, he re-elects an official who chooses policy B, whereas a voter who is relatively pro-A re-elects an official who chooses policy A. To see how this works, we define voter beliefs in various information sets. If the incumbent chooses policy $x = A$ in the first period let the voter's beliefs about the probability that the incumbent is type A, R, or B, respectively, be $\mu_A(A), \mu_R(A)$, and $\mu_B(A)$, where $\mu_A(A) + \mu_R(A) + \mu_B(A) = 1$. Similarly if $x = B$ let the voter's beliefs be $\mu_A(B), \mu_R(B)$, and $\mu_B(B)$.

It is relatively straightforward to characterize equilibria in circumstances where the voter has extreme preferences, i.e., either a type B voter with $\alpha_V < \underline{\alpha}^2$ or a type A voter with $\alpha_V > \overline{\alpha}^2$. Suppose that the voter observes first period policy choice $x = B$. What does he infer about the incumbent's type? Proposition 2 shows that a pro-B incumbent (with $\alpha_I < \underline{\alpha}^1$) always chooses $x = B$, a centrist incumbent (with $\alpha_I \in (\underline{\alpha}^1, \overline{\alpha}^1)$) chooses $x = B$ if and only if her signal s indicates that B is the correct policy choice, and a pro-A incumbent (with $\alpha_I > \overline{\alpha}^1$) never chooses $x = B$ in the first period. Thus, the voter's belief about the probability that the incumbent is type B goes up if she chooses $x = B$, i.e., $\mu_B(B) > \phi_B$.[14] Similarly, if the incumbent chooses $x = B$ then the voter's belief about the probability that she is type A goes down, i.e., $\mu_A(B) < \phi_A$.

What does this imply for a voter's re-election decision? Consider first a type B voter. From

[14] This is non-trivial to show since the first period cutpoints $\underline{\alpha}^1$ and $\overline{\alpha}^1$ are not the same as the second period cutpoints $\underline{\alpha}^2$ and $\overline{\alpha}^2$. For details see Lemmas 3 and 4 in the appendix.

this voter's perspective, a type B official is better than a type R official, who is in turn better than a type A official. Thus if this voter observes $x = B$ as the first period policy choice he strictly prefers to re-elect the incumbent, and this is true regardless of the first period cutpoints $\underline{\alpha}^1$ and $\overline{\alpha}^1$ for the incumbent's behavior. A type A voter, in contrast, has exactly the opposite preferences over the three types of officials, and thus for any cutpoints $\underline{\alpha}^1$ and $\overline{\alpha}^1$ he strictly prefers to remove an incumbent who chooses $x = B$.

By similar reasoning, if the official chooses $x = A$ in the first period, a type A voter strictly prefers to re-elect her whereas a type B voter strictly prefers to remove her.

What if the voter is type R, i.e., with $\alpha_V \in \left(\underline{\alpha}^2, \overline{\alpha}^2\right)$? This voter faces a tradeoff, since he prefers a signal-responsive official over either a type B or a type A official. For example, if the incumbent chooses $x = B$ the voter learns that $\mu_A(B) < \phi_A$, i.e., the incumbent is less likely to be type A than a randomly drawn new official; this is good news from the voter's perspective. However, $\mu_B(B) > \phi_B$, i.e., the incumbent is more likely to be type B than her replacement; this is bad news from a type R voter's perspective. How these tradeoffs balance out in equilibrium is characterized in the following proposition.

Proposition 3 (Equilibria Without Uncertainty Resolution) *For any voter type α_V there exists an equilibrium. The equilibrium depends on the value of α_V relative to cutpoints $\underline{\alpha}_V$ and $\overline{\alpha}_V$:*

1. $\underline{\alpha}^2 < \underline{\alpha}_V < \overline{\alpha}_V < \overline{\alpha}^2$

2. *If $\alpha_V < \underline{\alpha}_V$ there exists a unique equilibrium in which the voter re-elects the official if and only if she chooses $x = B$, i.e., $\nu_B = 1$ and $\nu_A = 0$*

3. *If $\alpha_V \in (\underline{\alpha}_V, \overline{\alpha}_V)$ there exists a continuum of equilibria in which the voter uses mixed strategies*

4. *If $\alpha_V > \overline{\alpha}_V$ there exists a unique equilibrium in which the voter re-elects the official if and only if she chooses $x = A$, i.e., $\nu_B = 0$ and $\nu_A = 1$.*

Example 1 To illustrate Proposition 3, suppose that the prior probability that $\omega = A$ is $\pi = \frac{1}{2}$ and the quality of the official's signal is $q = \frac{3}{4}$. In this example, the cutpoints for second period official behavior are $\underline{\alpha}^2 = 0.25$ and $\overline{\alpha}^2 = 0.75$. Figure 3 shows how different types of voters would ideally want the incumbent to act: a voter to the left of 0.25 always prefers $x = B$, a voter to the right of 0.75 always prefers $x = A$, and voter in the region $[0.25, 0.75]$ prefers that the incumbent follow her signal, choosing $x = A$ when $s = A$ and $x = B$ when $s = B$. Figure 3 also shows the type of electoral incentives that arise in equilibrium within the example.

[Figure 3 about here]

A few features of this equilibrium (and Proposition 3 more generally) are worth noting. If the voter is type B or type A ($\alpha_V < 0.25$ or $\alpha_V > 0.75$) the incumbent always has an electoral incentive to choose the voter's preferred policy. Thus not surprisingly, when voters have extreme preferences officials are responsive to these extreme preferences. What is more interesting is that for many R voter-types, who would like the official to choose A or B on the basis of the policy information she receives, equilibrium behavior is exactly the same as for a type A or type B voter. In Example 1, the equilibrium for a type R voter with $\alpha_V \in [0.25, 0.48]$ is the same as the equilibrium for a type B voter; the voter re-elects the official if and only if she chooses $x = B$ (Region I of Figure 3). Likewise, for $\alpha_V \in [0.52, 0.75]$ the equilibrium is the same as for a type A voter (Region III of Figure 3); the voter re-elects the official when $x = A$ and removes her when $x = B$. Notably, neither of these types of voter behavior promotes moderation, since the official is always rewarded for choosing one of the two policy options regardless of her signal.

A very centrist type R voter, with $\alpha_V \in [0.48, 0.52]$, uses a mixed strategy after observing at least one of the policy choices $x = A$ or $x = B$.[15] However, moderation still does not occur, because the incumbent's electoral incentive is always to choose whatever policy the voter initially desires. If

[15] The difference between ν_A and ν_B induces incumbent behavior such that the voter is indifferent between re-

$\nu_B > \nu_A$, as in Region II(i) of Figure 3, the incumbent always has an electoral incentive to choose B rather than A, regardless of her signal. If $\nu_A > \nu_B$, as in Region II(ii), then the incumbent always has an electoral incentive to choose A rather than B, regardless of her signal.[16] Consequently, the game plays out as in Figures 2a or 2b, but never 2c.

This example illustrates why in general there are no incentives for moderation when uncertainty does not resolve. The incumbent, lacking any expectation that the voter will gain updated information about the optimal policy choice, focuses instead on trying to signal that her ideological preferences are ones the voter would like. She does this even when the policy that advances her electoral prospects is one that she does not desire, and is more extreme than a fully informed voter would want.

Equilibria With Uncertainty Resolution

In the model with uncertainty resolution, moderation can occur. Yet at the same time, behavior is more similar to the case with no uncertainty resolution than one might expect.

Recall from Figure 2c that the key conditions for moderation are $r_A(B) < r_B(B)$ and $r_A(A) > r_B(A)$, which imply that the incumbent has an electoral incentive to follow her signal. Recall also that Equation 1 establishes that $r_A(s) = \nu_{AA}\theta^A(s) + \nu_{AB}\left(1 - \theta^A(s)\right)$ for $s \in \{A, B\}$. Since uncertainty does not resolve if the official chooses inaction, $r_B(B) = r_B(A) = \nu_B$ and the conditions for moderation simplify to $r_A(B) < \nu_B < r_A(A)$. There exists a ν_B that satisfies this inequality whenever $\nu_{AA}\theta^A(B) + \nu_{AB}\left(1 - \theta^A(B)\right) < \nu_{AA}\theta^A(A) + \nu_{AB}\left(1 - \theta^A(A)\right)$, which reduces to $\nu_{AA} > \nu_{AB}$, i.e., that the incumbent is strictly more likely to win re-election after a policy success than a

electing and removing the incument after either $x \in \{A, B\}$. Since there are only two information sets, any time the voter is indifferent in one information set he is also indifferent in the other.

[16] The final possiblility, $\nu_A = \nu_B$, only occurs for one specific value of α_V and in this case there is no electoral incentive to choose either A or B so moderation does not occur.

policy failure.

This condition, which seems at first to be very intuitive, and indeed essentially trivial, is actually rather difficult to satisfy, due to the voter's concern over re-electing the wrong type of incumbent. The condition is never satisfied for an extreme voter, i.e. one who is either type A or type B and thus always prefers the same policy action regardless of the state of the world. More notably, even a type R voter, who wants an extreme official to moderate her behavior by responding to the signal, will often set $\nu_{AB} = \nu_{AA}$ in equilibrium. In such cases the official is rewarded or punished simply for her policy choice, and *not* for the success or failure of this policy. And just as in the previous section, the lack of outcome-based rewards and punishments eliminates the possibility of elections inducing moderation. Furthermore, the condition $\nu_{AA} > \nu_{AB}$ (that an incumbent is more likely to win re-election after policy success than after policy failure) is only a necessary condition to have electoral incentives for moderation. It is not a sufficient condition, since the official's re-election incentives also depend on ν_B, her probability of winning re-election after choosing $x = B$.[17]

To understand the role of ν_B, consider an extension of Example 1 and suppose that $\nu_{AB} = 0$ and $\nu_{AA} = 1$, i.e., the incumbent is maximally punished for policy failure and rewarded for policy success. From Equation 1, the official's probability of winning re-election if she chooses $x = A$ when $s = A$ is $r_A(A) = \theta^A(A) = 0.75$, and her probability of winning re-election if she chooses $x = A$ when $s = B$ is $r_A(B) = \theta^A(B) = 0.25$. For moderation, the official must have an incentive to choose $x = A$ when $s = A$, i.e., $r_A(A) > r_B(A) = \nu_B$. Also the official must have an incentive to choose $x = B$ when $s = B$, i.e., $r_A(B) < r_B(B) = \nu_B$. Combining these two requirements, we see that if $\nu_{AB} = 0$ and $\nu_{AA} = 1$ moderation can only occur if $\nu_B \in (0.25, 0.75)$.[18]

We now turn to the general model and determine when in equilibrium moderation will actually

[17] If one assumed the voter learned ω after the official chose B (inaction), moderation would still not occur for type A voters, type B ones, and some type R voters.

[18] If $\nu_{AB} > 0$ or $\nu_{AA} < 1$ then the range of ν_B values that produce moderation is even more narrow.

occur. The following two propositions, one for type A and B voters, and one for type R voters, summarize the results.[19]

Proposition 4 (Equilibrium for Extreme Voters) *For any type A or type B voter there exists a unique equilibrium:*

1. *If $\alpha_V < \underline{\alpha}^2$ then the voter re-elects the official if and only if she chooses $x = B$, i.e., $\nu_B = 1$ and $\nu_{AA} = \nu_{AB} = 0$*

2. *If $\alpha_V > \overline{\alpha}^2$ then the voter re-elects the official if and only if she chooses $x = A$, i.e., $\nu_B = 0$ and $\nu_{AA} = \nu_{AB} = 1$.*

Proposition 4 shows that equilibrium behavior for a type A or type B voter in the model with uncertainty resolution is exactly the same as in the model without uncertainty resolution. Moderation cannot occur for these voter types, and there is always an electoral incentive to adopt the voter's favored policy.

Proposition 5 (Equilibria for Moderate Voters) *For any type R voter there exists an equilibrium, and all equilibria are one of the following five types, each of which occurs for some values of $\alpha_V \in \left(\underline{\alpha}^2, \overline{\alpha}^2\right)$:*

[19] If we allow for a non-uniform distribution of official types then equilibria like the ones in Propositions 4 and 5 always still exist. Under certain parameter values there also exists an equilibrium in which the incumbent always follows her signal regardless of α_I, and any type of voter is always indifferent between re-electing or removing the incumbent. This additional equilibrium can only occur if the official's signal is extremely accurate, i.e., for q close to 1. Also, the distribution of official types must be highly polarized; there must be a substantial probability that the challenger has $\alpha_C < \underline{\alpha}^2$ and also a substantial probability that $\alpha_C > \overline{\alpha}^2$, i.e., despite the fact that the signal is highly accurate, it must be likely that the challenger will simply ignore the signal if elected. If either of these conditions fails to hold then the equilibrium we characterize here is unique, even allowing for a non-uniform F.

1. $\nu_B = 1, \nu_{AA} = \nu_{AB} = 0$

2. $\nu_B = 1, \nu_{AA} \in (0,1], \nu_{AB} = 0$

3. $\nu_B \in (0,1), \nu_{AA} = 1, \nu_{AB} = 0$

4. $\nu_B = 0, \nu_{AA} = 1, \nu_{AB} \in [0,1)$

5. $\nu_B = 0, \nu_{AA} = 1, \nu_{AB} = 1$.

As in the model without uncertainty resolution, for some type R voters the equilibrium is exactly the same as the equilibrium for a type B voter (part 1 of the proposition) or for a type A voter (part 5 of the proposition). However, there are also several additional equilibria (parts 2-4). In which of these equilibria does moderation occur?

Consider part 2 of the proposition. Here $\nu_B = 1$ so if the incumbent chooses $x = B$ in the first period she is re-elected with certainty. If she chooses $x = A$, in contrast, there is some probability that the policy will fail and in that case she will lose re-election as $\nu_{AB} = 0$. Regardless of her signal the official is more likely to win re-election when she chooses B than when she chooses A. Similarly, in part 4 of the proposition, the official is always more likely to win re-election when she chooses A than when she chooses B since $\nu_B = 0$. Thus moderation cannot occur in these equilibria.

The only equilibrium for which moderation can occur is type 3, in which $\nu_B \in (0,1), \nu_{AA} = 1$, and $\nu_{AB} = 0$. However, moderation does not occur for all equilibria of this type; it occurs only for α_V values such that the equilibrium ν_B is sufficiently far away from zero and one to give the incumbent an incentive to follow her signal regardless of whether she believes A or B is the true state of the world. If ν_B is too close to one, the incumbent has an electoral incentive to choose B even if her signal suggests A is the optimal policy. Conversely, if ν_B is too close to zero, then the incumbent has an incentive to choose A even if her signal suggests B is the optimal policy.

It is worth noting that in any equilibrium for a type R voter where moderation does not occur, there will be pandering. Although any type R voter wants the official to follow her signal, his concern over selecting the official for the second period can cause him to adopt an electoral strategy that induces perverse behavior by an incumbent official with $\alpha_I \in \left(\underline{\alpha}^2, \overline{\alpha}^2\right)$, i.e., an official who prefers to follow her signal. Specifically, for one of the two signals $s \in \{A, B\}$ some types of signal-responsive officials are electorally induced to pander, going against the private signal in an attempt to prove to the voter that they are likely to choose the type of policies that he wants in the future.

Example 2. To demonstrate exactly when moderation occurs within a specific context, we extend our previous Example 1 ($\pi = \frac{1}{2}$, $q = \frac{3}{4}$). Figure 4 shows the equilibria of this example.

[Figure 4 about here]

Just as in the model without uncertainty resolution a type B voter ($\alpha_V < 0.25$), who prefers that the incumbent always choose $x = B$, bases his electoral decision solely on the incumbent's policy choice. He re-elects her if $x = B$ and removes her if $x = A$, regardless of whether policy A turns out to be the correct policy. Thus the incumbent has an electoral incentive to choose $x = B$, which is the voter's preferred policy. Similarly, a type A voter ($\alpha_V > 0.75$), who prefers that the incumbent always choose $x = A$, re-elects the incumbent if and only if $x = A$, regardless of whether the true state of the world is revealed to be $\omega = A$ or $\omega = B$.

A moderate voter, $\alpha_V \in (0.25, 0.75)$, prefers that the incumbent follow her signal, choosing $x = s$. Yet only in certain circumstances will the equilibrium electoral incentives actually induce moderation in the incumbent's policy choice. For a relatively pro-B moderate voter ($\alpha_V \in (0.25, 0.43)$), the only equilibrium is one in which the voter behaves exactly like a type B voter, rewarding the incumbent for choosing policy B and punishing her for choosing policy A, regardless of whether A turned out to be the correct policy choice. The reason the voter removes an official who chooses policy A is that this policy choice reveals information about the incumbent's type, i.e., that she is likely to be a type A official. Similarly, a relatively pro-A moderate voter $\alpha_V \in (0.57, 0.75)$ will remove an incumbent who chooses $x = B$ and re-elect an incumbent who chooses $x = A$, even if A turns out to be the wrong policy.

For the most centrist type R voters, i.e., $\alpha_V \in (0.43, 0.57)$, which corresponds to Region II of Figure 4, the equilibrium electoral behavior when $x = A$ depends on whether $\omega = A$ or $\omega = B$, and the voter is more likely to re-elect the incumbent after a successful policy A than after an unsuccessful one. Thus there is the potential for electoral incentives to induce moderation. However the only α_V values for which moderation occurs are in Region II(ii). Here, when $\alpha_V \in (0.49, 0.51)$, the voter is almost completely indifferent between the two types of policy mistakes – choosing A when B is the correct policy and choosing B when A is the correct policy; for other parts of Region II, the voter's greater concern over one type of policy mistake gives the incumbent the incentive to show that she is not type A or type B. For instance, when $\alpha_V \in (0.43, 0.49)$, the incumbent has an electoral incentive to choose B to show that she is not a dogmatic type A official. Thus, for only a small portion of the signal-responsive voter range does moderation actually occur in equilibrium.

This example illustrates the general finding that elections will often not induce moderation. Even when the electorate might learn the policy consequences of the official's policy decision, moderation will not occur unless the voter wants the official to follow her private information. The voter must also base his electoral choice on whether the decision produced a successful outcome. Moreover, moderation will not occur unless electoral behavior is balanced so that it provides an incentive for the official to follow her signal regardless of its content. These conditions seem innocuous at first glance, but as the example has highlighted, their combination can be rare once one allows for information- and preference-based differences between voters and politicians.

More generally, the theory has shown that elections in single-member, winner-take-all districts do not inherently promote moderation. Indeed, when an incumbent has private information about her preferences and expected policy consequences, there are many conditions under which her incentives are to take an action that is at least as extreme, if not more extreme, than she would if not running for re-election. Among other circumstances, this occurs when there is no possibility for uncertainty resolution or when the election does not depend on the success of the enacted policy. It can also occur when voters are so concerned about preventing a particular type of policy error that even though they want an incumbent to utilize her private policy information, they will only re-elect her if she endorses a particular option.

Conclusion

This paper has examined the incentives for moderation by an elected official who has private information about his preferences and the expected effects of policies. Contrary to conventional wisdom, the theory shows that there are not pervasive electoral incentives for moderation in single-member, winner-take-all districts. In fact, even in moderate districts, which want politicians to be open to choosing different policies on the basis on information about their expected consequences, officials can have electoral incentives to disregard any such information. This occurs because the official wants to signal that his ideological preferences are similar to those of the district.

Paradoxically, there are situations in which both the incumbent and voters, if they were fully informed, would want the incumbent to take a more moderate position. Thus as in several other recent theories, politicians can have the incentive to pander to voters by taking positions that the politicians do not want and that voters would not want if they had full information (Canes-Wrone, Herron, and Shotts 2001; Maskin and Tirole 2004; Fox 2005; Prat 2005). What is new about this theory is that the pandering is coincident with incentives for ideological extremism. We show that extreme, dogmatic actions can be popular simply because voters are worried about electing an incumbent with a different type of policy preferences; if voters knew the incumbent shared their preferences, then these extreme actions would not be popular.

The results highlight that even minor changes in district opinion, such as those that can result from gerrymandering or more natural changes in district composition, can have a large impact on elected officials' incentives. For instance, if a district goes from being quite moderate to slightly conservative, an official's electoral incentive can be to move considerably to the right, even more than the district would want if fully informed about his preferences. Likewise, if a district shifts from being moderate or slightly conservative to slightly liberal, an official can have an incentive to move more to the left than the electorate would want if fully informed. Our theory thus provides a new rationale for elites to be more polarized than the electorate (Fiorina, Abrams, and Pope 2004); most explanations depend on the existence of strong parties or interest groups that influence elite behavior. Here we have shown that even in the absence of such actors, officials can have incentives to be more extreme than their electorates.

These results pose a natural question: what might induce moderation? The theory highlights the important role of information. If voters knew the politicians' preferences with certainty, or knew their policy information, moderation would be a pervasive electoral incentive. The paper thus highlights the important role that information can play in altering elected officials' policy incentives.

Appendix

Proof of Proposition 1

We use a superscript t to denote expected utility from the period t policy choice. The official's utility is $U^2(B|s) = -\alpha\theta^A(s)$ from $x = B$ and $U^2(A|s) = -(1-\alpha)\left(1-\theta^A(s)\right)$ from $x = A$. The difference is $U^2(A|s) - U^2(B|s) = \alpha - 1 + \theta^A(s)$. The derivative with respect to α is 1 so we set $U^2(A|s) = U^2(B|s)$ for the cutpoints: $\underline{\alpha}^2 = 1 - \theta^A(A)$ and $\overline{\alpha}^2 = 1 - \theta^A(B)$. ■

To prove Proposition 2 we first prove a lemma.

Lemma 1 *For any voter strategy ν and probabilities ϕ_B, ϕ_R, and ϕ_A that the challenger is type B, R, and A, the incumbent's expected utility difference between $x = A$ and $x = B$ in the first period is a continuous, piecewise linear, strictly increasing function of α_I.*

Proof. Let $U(x|s)$ denote the official's total expected utility from choosing policy $x \in \{A, B\}$ in the first period given signal s. We first consider a type B official, and find $U(A|s) - U(B|s)$. The first component of the difference between these utilities is the first period utility difference from the two actions, which by the reasoning in the proof of Proposition 1 is $\alpha - 1 + \theta^A(s)$.

The second component is the second period effect of her first period action. The difference in her probability of being re-elected as a result of her policy choice is $r_A(s) - r_B(s)$. If the type B incumbent loses office, there is an increased chance of an incorrect policy A in the second period, which results in $-(1 - \alpha_I)$ utility. Specifically, with probability $\phi_R(1 - \pi)(1 - q)$ a type R challenger observes $s = A$ when $\omega = B$, and chooses $x = A$. Or, with probability $\phi_A(1 - \pi)$, the challenger is type A and chooses $x = A$ when $\omega = B$. If a type B official loses office, there is also a decreased chance of an incorrect policy B in the second period, which results in $-\alpha$ utility. With probability $\phi_R q\pi$ the challenger is type R and $s = A$ when $\omega = A$. Or, with probability $\phi_A\pi$, the challenger is type A and $\omega = A$. Combining these terms, for a type B incumbent, $\alpha_I < \underline{\alpha}^2$:

$$\begin{aligned}
U(A|s) - U(B|s) &= \alpha_I - 1 + \theta^A(s) + [r_A(s) - r_B(s)](1 - \alpha_I)(1 - \pi)(\phi_R(1 - q) + \phi_A) \\
&\quad - [r_A(s) - r_B(s)]\alpha_I\pi(\phi_R q + \phi_A) \\
&= \theta^A(s) - 1 + [r_A(s) - r_B(s)](1 - \pi)(\phi_R(1 - q) + \phi_A) \\
&\quad + \alpha_I\{1 - [r_A(s) - r_B(s)][\pi(\phi_R q + \phi_A) + (1 - \pi)(\phi_R(1 - q) + \phi_A)]\}
\end{aligned}$$

This expression is linear in α_I, with slope strictly greater than zero since $[r_A(s) - r_B(s)] \in [-1, 1]$, $\pi \in (0, 1)$, $(\phi_R q + \phi_A) \in (0, 1)$, and $(\phi_R(1 - q) + \phi_A) \in (0, 1)$.

A type R incumbent's reasoning is similar, except that a type R challenger will act like the incumbent. For $\alpha_I \in \left(\underline{\alpha}^2, \overline{\alpha}^2\right)$

$$\begin{aligned} U(A|s) - U(B|s) &= \alpha_I - 1 + \theta^A(s) + [r_A(s) - r_B(s)](1-\alpha_I)(1-\pi)(\phi_A q - \phi_B(1-q)) \\ &\quad + [r_A(s) - r_B(s)]\alpha_I \pi (\phi_B q - \phi_A(1-q)) \\ &= \theta^A(s) - 1 + [r_A(s) - r_B(s)](1-\pi)(\phi_A q - \phi_B(1-q)) \\ &\quad + \alpha_I \{1 + [r_A(s) - r_B(s)][\pi(\phi_B q - \phi_A(1-q)) + (1-\pi)(\phi_B(1-q) - \phi_A q)]\}. \end{aligned}$$

This expression is linear in α_I, with slope strictly greater than zero since $(\phi_B q - \phi_A(1-q)) \in (-1, 1)$ and $(\phi_B(1-q) - \phi_A q) \in (-1, 1)$.

For a type A official, $\alpha_I > \underline{\alpha}^2$,

$$\begin{aligned} U(A|s) - U(B|s) &= \alpha_I - 1 + \theta^A(s) - [r_A(s) - r_B(s)](1-\alpha_I)(1-\pi)(\phi_B + \phi_R q) \\ &\quad + [r_A(s) - r_B(s)]\alpha_I \pi (\phi_B + \phi_R(1-q)) \\ &= \theta^A(s) - 1 - [r_A(s) - r_B(s)](1-\pi)(\phi_B + \phi_R q) \\ &\quad + \alpha_I \{1 + [r_A(s) - r_B(s)][\pi(\phi_B + \phi_R(1-q)) + (1-\pi)(\phi_B + \phi_R q)]\}. \end{aligned}$$

This expression is linear in α_I, with slope strictly greater than zero since $(\phi_B + \phi_R(1-q)) \in (0, 1)$ and $(\phi_B + \phi_R q) \in (0, 1)$. Finally, from the expressions for the three types of officials it is straightforward to confirm that $U(A|s) - U(B|s)$ is continuous at $\underline{\alpha}^2$ and $\overline{\alpha}^2$.■

Lemma 2 *There exists a cutpoint $\tilde{\alpha}$ such that for any voter strategy ν it is strictly optimal for an incumbent with $\alpha_I \leq \tilde{\alpha}$ to choose $x = B$ when $s = B$ the first period and it is strictly optimal for an incumbent with $\alpha_I \geq \tilde{\alpha}$ to choose $x = A$ when $s = A$ the first period.*

Proof. In period t, an official can observe $s = B$ or $s = A$. A type B official cares more about current-period effects of choosing $x = B$ when $s = B$ than when $s = A$, i.e., for any $\alpha_I < \underline{\alpha}^2$,

$$\begin{aligned} U^t(B|B) - U^t(A|B) &> U^t(B|A) - U^t(A|A) \\ 1 - \alpha - \theta^A(B) &> 1 - \alpha - \theta^A(A) \\ \theta^A(A) &> \theta^A(B). \end{aligned}$$

Similarly, for a type A official, i.e., $\alpha_I > \overline{\alpha}^2$, $U^t(A|A) - U^t(B|A) > U^t(A|B) - U^t(B|B)$.

For a type R official we find $\tilde{\alpha}$ such that the official cares equally about the two decisions, $U^t(A|A) - U^t(B|A) = U^t(B|B) - U^t(A|B)$. This reduces to $\tilde{\alpha} = 1 - \frac{\theta^A(A) + \theta^A(B)}{2}$. It is straightforward to confirm that $\underline{\alpha}^2 < \tilde{\alpha} < \overline{\alpha}^2$.

We still must establish that for $\alpha_I \leq \tilde{\alpha}$, $x = B$ is optimal in the first period when $s = B$. For $\alpha_I < \tilde{\alpha}$, the biggest possible difference between an official's expected utility from her own policy choice when re-elected versus a challenger's policy choice occurs when $s = B$ in the second period. She can lose up to $U^2(B|B) - U^2(A|B)$ if the challenger is type A. However, this occurs with probability strictly less than 1, since the second period signal may be $s = A$ or the replacement official may be type B. Thus a strict upper bound on the official's expected second period utility loss from choosing $x = A$ when $s = B$ in the first period is $U^2(B|B) - U^2(A|B)$. Since the official's first period utility difference when $s = B$ is $U^1(B|B) - U^1(A|B) = U^2(B|B) - U^2(A|B)$ it is strictly optimal for her to choose $x = B$. For $\alpha_I \geq \tilde{\alpha}$ a symmetric argument shows that $x = A$ is strictly optimal when $s = A$ in the first period.■

Proof of Proposition 2

We first find the cutpoint $\overline{\alpha}^1 \in (\tilde{\alpha}, 1]$ for first period official behavior when $s = B$. Lemma 2 establishes that choosing $x = B$ is strictly optimal for all $\alpha_I < \tilde{\alpha}$, and that $U(A|B) - U(B|B) < 0$ for $\alpha_I = \tilde{\alpha}$. Lemma 1 establishes that $U(A|B) - U(B|B)$ is continuous and strictly increasing in $[\tilde{\alpha}, 1]$. Thus $\overline{\alpha}^1 = \min\{\alpha_I : U(A|B) - U(B|B) \geq 0\}$ or if there is no $\alpha_I \leq 1$ for which

$U(A|B) - U(B|B) \geq 0$ then $\overline{\alpha}^1 = 1$. A similar argument yields a cutpoint $\underline{\alpha}^1 \in [0, \tilde{\alpha})$, where $\underline{\alpha}^1 = \max\{\alpha_I : U(A|A) - U(B|A) \leq 0\}$, for first period official behavior when $s = A$.

We now prove by contradiction that one of the inequalities in part 1 of the proposition must be strict. To have $\underline{\alpha}^1 = 0$ requires that the most extreme type B incumbent, with $\alpha_I = 0$, choose $x = A$ when $s = A$. From the utility difference functions for a type B incumbent in the proof of Lemma 1 this requires

$$\begin{aligned} 0 &< U(A|A) - U(B|A) \\ 0 &< \theta^A(A) - 1 + \left[\nu_{AA}\theta^A(A) + \nu_{AB}\left(1 - \theta^A(A)\right) - \nu_B\right](1-\pi)\left((\phi_R(1-q) + \phi_A)\right) \\ \nu_B &< \left[\nu_{AA}\theta^A(A) + \nu_{AB}\left(1 - \theta^A(A)\right)\right] - \frac{1 - \theta^A(A)}{(1-\pi)(\phi_R(1-q) + \phi_A)}. \end{aligned}$$

Similarly, to have $\overline{\alpha}^1 = 1$ requires that for the most extreme type A incumbent, with $\alpha_I = 1$

$$\begin{aligned} 0 &> U(A|B) - U(B|B) \\ 0 &> \theta^A(B) + \left[\nu_{AA}\theta^A(B) + \nu_{AB}\left(1 - \theta^A(B)\right) - \nu_B\right]\pi\left(\phi_B + \phi_R(1-q)\right) \\ \nu_B &> \left[\nu_{AA}\theta^A(B) + \nu_{AB}\left(1 - \theta^A(B)\right)\right] + \frac{\theta^A(B)}{\pi(\phi_B + \phi_R(1-q))}. \end{aligned}$$

For there to exist a ν_B that satisfies both of these inequalities would require that $0 < \left[\nu_{AA}\left(\theta^A(A) - \theta^A(B)\right) - \nu_{AB}\left(\theta^A(A) - \theta^A(B)\right)\right] - \frac{1-\theta^A(A)}{(1-\pi)(\phi_R(1-q)+\phi_A)} - \frac{\theta^A(B)}{\pi(\phi_B+\phi_R(1-q))}$. Since $\theta^A(A) - \theta^A(B) > 0$, if this inequality is satisfied for any voter strategies ν_{AA} and ν_{AB} it holds for $\nu_{AA} = 1$ and $\nu_{AB} = 0$, i.e., $0 < \theta^A(A) - \theta^A(B) - \frac{1-\theta^A(A)}{(1-\pi)(\phi_R(1-q)+\phi_A)} - \frac{\theta^A(B)}{\pi(\phi_B+\phi_R(1-q))}$. The distribution of official types F is uniform so $\phi_B = \underline{\alpha}^2 = 1 - \theta^A(A)$, $\phi_A = 1 - \overline{\alpha}^2 = \theta^A(B)$, and $\phi_R = \theta^A(A) - \theta^A(B)$. Rearranging, substituting, and cancelling terms in the previous inequality yields $\frac{1}{\Pr(s=A)\left(\theta^A(A) - \theta^A(B) + \frac{\pi}{\Pr(s=B)}\right)} + \frac{1}{\Pr(s=B)\left(\frac{1-\pi}{\Pr(s=A)} + \theta^A(A) - \theta^A(B)\right)} < \theta^A(A) - \theta^A(B)$. Since $\theta^A(A) - \theta^A(B) \in (0,1)$ and both terms on the left hand side are positive, if we replace $\theta^A(A) - \theta^A(B)$ with 1 the left hand side will be made smaller. And substituting in 1 for $\theta^A(A) - \theta^A(B)$ makes the right hand side larger, so if the previous inequality holds $\frac{1}{\Pr(s=A)\left(1 + \frac{\pi}{\Pr(s=B)}\right)} + \frac{1}{\Pr(s=B)\left(\frac{1-\pi}{\Pr(s=A)} + 1\right)} < 1$.

Since each left hand side term is strictly positive, the inequality requires that each of the two terms be strictly less than 1. We find a contradiction, depending on whether $\pi < 1/2$ or $\pi \geq 1/2$. For $\pi < 1/2$, note that the inequality requires $\frac{1}{\Pr(s=A)\left(1+\frac{\pi}{\Pr(s=B)}\right)} < 1$, i.e., $\frac{1}{\left(\frac{\Pr(s=A)}{\Pr(s=B)}\right)(1-\Pr(s=A)+\pi)} < 1$. When $\pi < 1/2$, it is straightforward to show that $\Pr(s = A) < \Pr(s = B)$ and $\Pr(s = A) > \pi$. Thus each of the two terms in the denominator is strictly less than 1 so the left hand side is strictly greater than 1. By similar reasoning it can be shown that $\frac{1}{\Pr(s=B)\left(\frac{1-\pi}{\Pr(s=A)}+1\right)} \geq 1$ when $\pi \geq 1/2$, which yields a contradiction.■

For Propositions 3 and 4, we characterize voter beliefs.

Lemma 3 *In the model without uncertainty resolution, for any first period official strategy as in Proposition 2,* $\mu_B(B) > \phi_B > \mu_B(A)$ *and* $\mu_A(B) < \phi_A < \mu_A(A)$.

Proof. Very similar to the proof of Lemma 4. Omitted and available upon request.

Lemma 4 *In the model with uncertainty resolution, for any first period official strategy as in Proposition 2,* $\mu_B(B) > \phi_B > \mu_B(AA) \geq \mu_B(AB)$ *and* $\mu_A(B) < \phi_A < \mu_A(AA) \leq \mu_A(AB)$.

There are four cases: **(i)** $\underline{\alpha}^1 < \underline{\alpha}^2$ and $\overline{\alpha}^1 \leq \overline{\alpha}^2$, **(ii)** $\underline{\alpha}^1 \geq \underline{\alpha}^2$ and $\overline{\alpha}^1 > \overline{\alpha}^2$, **(iii)** $\underline{\alpha}^1 < \underline{\alpha}^2$ and $\overline{\alpha}^1 > \overline{\alpha}^2$, **(iv)** $\underline{\alpha}^1 \geq \underline{\alpha}^2$ and $\overline{\alpha}^1 \leq \overline{\alpha}^2$. We show that in each case, voter beliefs about the probability that the incumbent official is type B after the first period policy outcome is revealed can be ordered as follows: $\mu_B(B) > \phi_B > \mu_B(AA) \geq \mu_B(AB)$. Since the first period outcome must be B, AA, or AB, ϕ_B is a weighted average of $\mu_B(B), \mu_B(AA)$, and $\mu_B(AB)$. Thus it is sufficient to prove the latter two inequalities. Similarly for beliefs about the probability that the official is type A, we show that $\phi_A < \mu_A(AA)$ and $\mu_A(AA) \leq \mu_A(AB)$ so that $\mu_A(B) < \phi_A < \mu_A(AA) \leq \mu_A(AB)$.

For voter beliefs in **case (i)**, $\mu_B(AB) = \frac{\left[F(\underline{\alpha}^2)-F\left(\underline{\alpha}^1\right)\right](1-\pi)(1-q)}{[1-F(\underline{\alpha}^1)](1-\pi)(1-q)+\left[1-F\left(\overline{\alpha}^1\right)\right](1-\pi)q} = \frac{\left[F(\underline{\alpha}^2)-F\left(\underline{\alpha}^1\right)\right]}{[1-F(\underline{\alpha}^1)]+\left[1-F\left(\overline{\alpha}^1\right)\right]\left(\frac{q}{1-q}\right)}$ and $\mu_B(AA) = \frac{\left[F(\underline{\alpha}^2)-F\left(\underline{\alpha}^1\right)\right]\pi q}{[1-F(\underline{\alpha}^1)]\pi q+\left[1-F\left(\overline{\alpha}^1\right)\right]\pi(1-q)} = \frac{\left[F(\underline{\alpha}^2)-F\left(\underline{\alpha}^1\right)\right]}{[1-F(\underline{\alpha}^1)]+\left[1-F\left(\overline{\alpha}^1\right)\right]\left(\frac{1-q}{q}\right)}$. The only difference is the last term of the denominator, and $q > 1/2$ so $\mu_B(AB) < \mu_B(AA)$. To show that $\mu_B(AA) < \phi_B = F(\underline{\alpha}^2)$, note that the second term in the denominator of $\mu_B(AA) = \frac{\left[F(\underline{\alpha}^2) - F\left(\underline{\alpha}^1\right)\right]}{[1-F(\underline{\alpha}^1)] + \left[1-F\left(\overline{\alpha}^1\right)\right]\left(\frac{1-q}{q}\right)}$ is strictly greater than zero. Thus it's sufficient to show that $\frac{\left[F(\underline{\alpha}^2) - F\left(\underline{\alpha}^1\right)\right]}{[1-F(\underline{\alpha}^1)]} \leq F(\underline{\alpha}^2)$, which reduces to $F(\underline{\alpha}^2) - F\left(\underline{\alpha}^1\right) \leq F(\underline{\alpha}^2) - F(\underline{\alpha}^2)F\left(\underline{\alpha}^1\right)$.

We now turn to beliefs about the probability that the official is type A in case (i). Here $\mu_A(AA) = \frac{\left[1-F(\overline{\alpha}^2)\right]\pi}{\left[1-F(\overline{\alpha}^1)\right]\pi + \left[F(\overline{\alpha}^1) - F(\underline{\alpha}^1)\right]\pi q}$ and $\mu_A(AB) = \frac{\left[1-F(\overline{\alpha}^2)\right](1-\pi)}{\left[1-F(\overline{\alpha}^1)\right](1-\pi) + \left[F(\overline{\alpha}^1) - F(\underline{\alpha}^1)\right](1-\pi)(1-q)}$. After cancelling out π and $(1-\pi)$ terms the only difference is the last term of the denominator, and $q > 1/2$ so $\mu_A(AA) < \mu_A(AB)$.

For $\phi_A = 1 - F(\overline{\alpha}^2) < \mu_A(AA)$, we re-write $\mu_A(AA)$ as $\frac{\phi_A}{\left[1-F(\overline{\alpha}^1)\right] + \left[F(\overline{\alpha}^1) - F(\underline{\alpha}^1)\right]q}$ and note that the denominator is strictly less than 1.

For **case (ii)**, since $\underline{\alpha}^1 \geq \underline{\alpha}^2$ no type B incumbent ever chooses $x = A$ so $\mu_B(AA) = \mu_B(AB) = 0$ and thus $\mu_B(B) > \phi_B > \mu_B(AA) \geq \mu_B(AB)$. Beliefs about the probability that the official is type A are $\mu_A(AA) = \frac{\left[1-F\left(\overline{\alpha}^1\right)\right] + \left[F\left(\overline{\alpha}^1\right) - F(\overline{\alpha}^2)\right]q}{\left[1-F\left(\overline{\alpha}^1\right)\right] + \left[F\left(\overline{\alpha}^1\right) - F\left(\underline{\alpha}^1\right)\right]q}$ and $\mu_A(AB) = \frac{\left[1-F\left(\overline{\alpha}^1\right)\right] + \left[F\left(\overline{\alpha}^1\right) - F(\overline{\alpha}^2)\right](1-q)}{\left[1-F\left(\overline{\alpha}^1\right)\right] + \left[F\left(\overline{\alpha}^1\right) - F\left(\underline{\alpha}^1\right)\right](1-q)}$. Straightforward algebra shows $\phi_A < \mu_A(AA) < \mu_A(AB)$.

For **case (iii)**, the proof that $\mu_B(B) > \phi_B > \mu_B(AA) \geq \mu_B(AB)$ follows case (i) and the proof that $\mu_A(B) < \phi_A < \mu_A(AA) \leq \mu_A(AB)$ follows case (ii). For **case (iv)**, the proof that $\mu_B(B) > \phi_B > \mu_B(AA) \geq \mu_B(AB)$ follows case (ii) and the proof that $\mu_A(B) < \phi_A < \mu_A(AA) \leq \mu_A(AB)$ follows case (i).■

We now state without proof a trivial lemma for preferences of type A and type B voters over the three types of officials. The equilibria for type A and type B voters in Propositions 3 and 4 follow directly from Lemmas 3-5.

Lemma 5 *A type B voter strictly prefers a type B official over a type R official and strictly prefers a type R official over a type A official. A type A voter has the opposite strict ordinal preferences.*

We prove Proposition 5 through a series of lemmas. The proof of part 3 of Proposition 3, which is very similar, is available upon request.

Lemma 6 *If first period official behavior is characterized by cutpoints* $\underline{\alpha}^1$ *and* $\overline{\alpha}^1$ *as in Proposition 2 then there exist cutpoints* α^{AA}, α^B*, and* α^{AB} *such that:*

1. $\underline{\alpha}^2 < \alpha^{AA} \leq \alpha^B \leq \alpha^{AB} < \overline{\alpha}^2$

2. *If* $x = A$ *and* $\omega = A$*, a voter* $\alpha_V < \alpha^{AA}$ *strictly prefers to remove the official, a voter* $\alpha_V > \alpha^{AA}$ *strictly prefers to re-elect the official, and a voter* $\alpha_V = \alpha^{AA}$ *is indifferent*

3. *If* $x = B$*, a voter* $\alpha_V < \alpha^B$ *strictly prefers to re-elect the official, a voter* $\alpha_V > \alpha^B$ *strictly prefers to remove the official, and a voter* $\alpha_V = \alpha^B$ *is indifferent*

4. *If* $x = A$ *and* $\omega = B$*, a voter* $\alpha_V < \alpha^{AB}$ *strictly prefers to remove the official, a voter* $\alpha_V > \alpha^{AB}$ *strictly prefers to re-elect the official, and a voter* $\alpha_V = \alpha^{AB}$ *is indifferent.*

Proof. We first establish existence of the cutpoints. For α^B, note that the difference in the voter's expected utility difference from re-electing versus removing the incumbent is a linear, and hence monotonic, function of α_V. We denote these utilities as $U(old|B)$ and $U(new)$.

$$\begin{aligned}
U(old|B) - U(new) &= -\mu_B(B)\alpha_V\pi - \mu_R(B)(1-q)\left[\alpha_V\pi + (1-\alpha_V)(1-\pi)\right] - \mu_A(B)(1-\alpha_V)(1-\pi) - \\
&\quad \{-\phi_B\alpha_V\pi - \phi_R(1-q)\left[\alpha_V\pi + (1-\alpha_V)(1-\pi)\right] - \phi_A(1-\alpha_V)(1-\pi)\} \\
&= \alpha_V\pi\left\{[\phi_B - \mu_B(B)] + [\phi_R - \mu_R(B)](1-q)\right\} + \\
&\quad (1-\alpha_V)(1-\pi)\left\{[\phi_R - \mu_R(B)](1-q) + [\phi_A - \mu_A(B)]\right\}
\end{aligned}$$

Also, from Lemmas 4 and 5 a voter at $\alpha_V = \underline{\alpha}^2$ strictly prefers to re-elect the official when $x = B$ and a voter at $\alpha_V = \overline{\alpha}^2$ strictly prefers to remove her. Thus there exists a cutpoint $\alpha^B \in \left(\underline{\alpha}^2, \overline{\alpha}^2\right)$ such that a voter at $\alpha_V < \alpha^B$ prefers to re-elect whereas a voter at $\alpha_V > \alpha^B$ strictly prefers to remove. A similar argument establishes the cutpoints α^{AA} and α^{AB} for optimal voter behavior when $x = A$.

We now order the cutpoints. If $\alpha^{AA} < \alpha^{B}$ and $\alpha^{AB} < \alpha^{B}$ then a voter $\alpha_V \in \left(\max\left\{\alpha^{AA}, \alpha^{AB}\right\}, \alpha^{B}\right)$ strictly prefers to re-elect the incumbent after all possible first period outcomes. This is a contradiction since the challenger is drawn from the same pool as the incumbent. A similar contradiction occurs if $\alpha^{AA} > \alpha^{B}$ and $\alpha^{AB} > \alpha^{B}$.

The final part of the argument is to show that $\alpha^{AA} \leq \alpha^{AB}$ so that $\alpha^{AA} \leq \alpha^{B} \leq \alpha^{AB}$. To prove that $\alpha^{AA} \leq \alpha^{AB}$ we show that if a voter's utility from re-electing after $x = A$ and $\omega = B$ is greater than or equal to his utility from re-electing after $x = A$ and $\omega = A$ then his utility from re-electing after $x = A$ and $\omega = A$ is greater than or equal to his utility from a new official. We denote these utilities as $U\left(old|AA\right)$ and $U\left(old|AB\right)$. We use $U\left(\alpha > g\right)$ to denote a voter's expected utility from an official randomly drawn from the portion of the distribution F that is greater than $g \in (0,1)$. Similarly $U\left(\alpha \in (g,h)\right)$ denotes expected utility from an official drawn from F restricted to the interval $(g,h) \subseteq (0,1)$.

First note that if $\underline{\alpha}^1 > \underline{\alpha}^2$ then for a type R voter $U\left(old|AA\right) > U\left(old|AB\right)$. This holds since a type B official never chooses $x = A$ when $\underline{\alpha}^1 > \underline{\alpha}^2$ and $\mu_A(AA) < \mu_A(AB)$ as shown in the proof of Lemma 4.

The argument is more complicated when $\underline{\alpha}^1 \leq \underline{\alpha}^2$ and we proceed in 4 steps. If $\overline{\alpha}^1 = 1$ then $\mu_A\left(AA\right) = \mu_A\left(AB\right)$ and $\mu_B\left(AA\right) = \mu_B\left(AB\right)$ so $U\left(old|AA\right) = U\left(old|AB\right)$ and if $U\left(old|AB\right) \geq U\left(new\right)$ then $U\left(old|AA\right) \geq U\left(new\right)$. We therefore restrict attention to $\overline{\alpha}^1 < 1$ and thus $1 - F\left(\overline{\alpha}^1\right) > 0$.

Step 1. We first show that $\Pr\left(\alpha_I > \overline{\alpha}^1 | AB\right) > \Pr\left(\alpha_I > \overline{\alpha}^1 | AA\right) > 1 - F\left(\overline{\alpha}^1\right)$:

$$\begin{aligned}
\Pr\left(\alpha_I > \overline{\alpha}^1 | AB\right) &> \Pr\left(\alpha_I > \overline{\alpha}^1 | AA\right) \\
\frac{\left[1 - F\left(\overline{\alpha}^1\right)\right](1-\pi)}{\left[1 - F\left(\overline{\alpha}^1\right)\right](1-\pi) + \left[F\left(\overline{\alpha}^1\right) - F\left(\underline{\alpha}^1\right)\right](1-\pi)(1-q)} &> \frac{\left[1 - F\left(\overline{\alpha}^1\right)\right]\pi}{\left[1 - F\left(\overline{\alpha}^1\right)\right]\pi + \left[F\left(\overline{\alpha}^1\right) - F\left(\underline{\alpha}^1\right)\right]\pi q} \\
q &> 1 - q.
\end{aligned}$$

For the second inequality,

$$\begin{aligned}
\Pr\left(\alpha_I > \overline{\alpha}^1 | AA\right) &> 1 - F\left(\overline{\alpha}^1\right) \\
\frac{\left[1 - F\left(\overline{\alpha}^1\right)\right]\pi}{\left[1 - F\left(\overline{\alpha}^1\right)\right]\pi + \left[F\left(\overline{\alpha}^1\right) - F\left(\underline{\alpha}^1\right)\right]\pi q} &> 1 - F\left(\overline{\alpha}^1\right) \\
1 &> \left[1 - F\left(\overline{\alpha}^1\right)\right] + \left[F\left(\overline{\alpha}^1\right) - F\left(\underline{\alpha}^1\right)\right] q.
\end{aligned}$$

Step 2. We show that if $U\left(old|AB\right) > U\left(old|AA\right)$ then $U\left(\alpha > \overline{\alpha}^1\right) > U\left(\alpha \in \left(\underline{\alpha}^1, \overline{\alpha}^1\right)\right)$.

$$\begin{aligned}
U\left(old|AB\right) &> U\left(old|AA\right) \\
\begin{array}{c}\Pr\left(\alpha_I > \overline{\alpha}^1 | AB\right) U\left(\alpha > \overline{\alpha}^1\right) + \\ \Pr\left(\alpha_I \in \left(\underline{\alpha}^1, \overline{\alpha}^1\right) | AB\right) U\left(\alpha \in \left(\underline{\alpha}^1, \overline{\alpha}^1\right)\right)\end{array} &> \begin{array}{c}\Pr\left(\alpha_I > \overline{\alpha}^1 | AA\right) U\left(\alpha > \overline{\alpha}^1\right) + \\ \Pr\left(\alpha_I \in \left(\underline{\alpha}^1, \overline{\alpha}^1\right) | AA\right) U\left(\alpha \in \left(\underline{\alpha}^1, \overline{\alpha}^1\right)\right)\end{array}
\end{aligned}$$

Since $\Pr\left(\alpha_I < \underline{\alpha}^1 | AB\right) = \Pr\left(\alpha_I < \underline{\alpha}^1 | AA\right) = 0$, we can substitute $1 - \Pr\left(\alpha_I > \overline{\alpha}^1 | AB\right) = \Pr\left(\alpha_I \in \left(\underline{\alpha}^1, \overline{\alpha}^1\right) | AB\right)$ and $1 - \Pr\left(\alpha_I > \overline{\alpha}^1 | AA\right) = \Pr\left(\alpha_I \in \left(\underline{\alpha}^1, \overline{\alpha}^1\right) | AA\right)$ to get

$$\begin{aligned}
\begin{array}{c}\Pr\left(\alpha_I > \overline{\alpha}^1 | AB\right)\left[U\left(\alpha > \overline{\alpha}^1\right) - U\left(\alpha \in \left(\underline{\alpha}^1, \overline{\alpha}^1\right)\right)\right] + \\ U\left(\alpha \in \left(\underline{\alpha}^1, \overline{\alpha}^1\right)\right)\end{array} &> \begin{array}{c}\Pr\left(\alpha_I > \overline{\alpha}^1 | AA\right)\left[U\left(\alpha > \overline{\alpha}^1\right) - U\left(\alpha \in \left(\underline{\alpha}^1, \overline{\alpha}^1\right)\right)\right] + \\ U\left(\alpha \in \left(\underline{\alpha}^1, \overline{\alpha}^1\right)\right)\end{array} \\
\begin{array}{c}\left[\Pr\left(\alpha_I > \overline{\alpha}^1 | AB\right) - \Pr\left(\alpha_I > \overline{\alpha}^1 | AA\right)\right] \cdot \\ \left[U\left(\alpha > \overline{\alpha}^1\right) - U\left(\alpha \in \left(\underline{\alpha}^1, \overline{\alpha}^1\right)\right)\right]\end{array} &> 0.
\end{aligned}$$

In Step 1, we proved that the first term is strictly greater than zero, so $U\left(\alpha > \overline{\alpha}^1\right) > U\left(\alpha \in \left(\underline{\alpha}^1, \overline{\alpha}^1\right)\right)$.

Step 3. We show that $U\left(\alpha \in \left(\underline{\alpha}^1, \overline{\alpha}^1\right)\right) > U\left(\alpha < \underline{\alpha}^1\right)$. There are two cases: $\overline{\alpha}^1 \leq \overline{\alpha}^2$ and $\overline{\alpha}^1 > \overline{\alpha}^2$. For the first case, if $\alpha \in \left(\underline{\alpha}^1, \overline{\alpha}^1\right)$ then the official is either type B or type R, and if $\alpha < \underline{\alpha}^1$ the official is type B with probability 1. The type R voter strictly prefers type R over type B officials so $U\left(\alpha \in \left(\underline{\alpha}^1, \overline{\alpha}^1\right)\right) > U\left(\alpha < \underline{\alpha}^1\right)$.

For the second case, $\overline{\alpha}^1 > \overline{\alpha}^2$ implies that if $\alpha > \overline{\alpha}^1$ then the official is surely type A. From Step 2, $U\left(\alpha > \overline{\alpha}^1\right) > U\left(\alpha \in \left(\underline{\alpha}^1, \overline{\alpha}^1\right)\right)$. Note that $\left(\underline{\alpha}^1, \overline{\alpha}^1\right)$ includes some officials of each type (R, B, and A). Also, a type R voter most prefers a type R official so if $U\left(\alpha > \overline{\alpha}^1\right) > U\left(\alpha \in \left(\underline{\alpha}^1, \overline{\alpha}^1\right)\right)$ then a type B official is the voter's least preferred type. Since $\alpha < \underline{\alpha}^1$ implies that the official is type B for sure, $U\left(\alpha \in \left(\underline{\alpha}^1, \overline{\alpha}^1\right)\right) > U\left(\alpha < \underline{\alpha}^1\right)$.

Step 4. We show that $U\left(old|AA\right) > U\left(new\right)$.

$$
\begin{array}{ccc}
U\left(old|AA\right) & > & U\left(new\right) \\
\begin{array}{c} \Pr\left(\alpha_I > \overline{\alpha}^1|AA\right) U\left(\alpha > \overline{\alpha}^1\right) + \\ \Pr\left(\alpha_I \in \left(\underline{\alpha}^1, \overline{\alpha}^1\right)|AA\right) U\left(\alpha \in \left(\underline{\alpha}^1, \overline{\alpha}^1\right)\right) \end{array} & > & \begin{array}{c} \left[1 - F\left(\overline{\alpha}^1\right)\right] U\left(\alpha > \overline{\alpha}^1\right) + \\ \left[F\left(\overline{\alpha}^1\right) - F\left(\underline{\alpha}^1\right)\right] U\left(\alpha \in \left(\underline{\alpha}^1, \overline{\alpha}^1\right)\right) + \\ F\left(\underline{\alpha}^1\right) U\left(\alpha < \underline{\alpha}^1\right) \end{array}
\end{array}
$$

From Step 3, $U\left(\alpha \in \left(\underline{\alpha}^1, \overline{\alpha}^1\right)\right) > U\left(\alpha < \underline{\alpha}^1\right)$ so the inequality will hold if $\Pr\left(\alpha_I > \overline{\alpha}^1|AA\right) U\left(\alpha > \overline{\alpha}^1\right) + \Pr\left(\alpha_I \in \left(\underline{\alpha}^1, \overline{\alpha}^1\right)|AA\right) \ U\left(\alpha \in \left(\underline{\alpha}^1, \overline{\alpha}^1\right)\right) \ > \ \left[1 - F\left(\overline{\alpha}^1\right)\right] \ U\left(\alpha > \overline{\alpha}^1\right) + \ F\left(\overline{\alpha}^1\right) U\left(\alpha \in \left(\underline{\alpha}^1, \overline{\alpha}^1\right)\right)$. Since $\underline{\alpha}^1 > \underline{\alpha}^2$ no type B official chooses $x = A$ so $\Pr\left(\alpha_I > \overline{\alpha}^1|AA\right) + \Pr\left(\alpha_I \in \left(\underline{\alpha}^1, \overline{\alpha}^1\right)|AA\right) = 1$. Substituting, the inequality simplifies to $\left[\Pr\left(\alpha_I > \overline{\alpha}^1|AA\right) - \left(1 - F\left(\overline{\alpha}^1\right)\right)\right] \left[U\left(\alpha > \overline{\alpha}^1\right) - U\left(\alpha \in \left(\underline{\alpha}^1, \overline{\alpha}^1\right)\right)\right]$ > 0. Steps 1 and 2 establish that each term is strictly greater than zero so the inequality holds.■

Proposition 5 is a restatement of the following lemma.

Lemma 7 *Each of the following is an equilibrium voter strategy for some* $\alpha_V \in \left(\underline{\alpha}^2, \overline{\alpha}^2\right)$ *: (1)* $\nu_B = 1, \nu_{AA} = \nu_{AB} = 0$, *(2)* $\nu_B = 1, \nu_{AA} \in (0,1), \nu_{AB} = 0$, *(3)* $\nu_B = 1, \nu_{AA} = 1, \nu_{AB} = 0$, *(4)* $\nu_B \in (0,1), \nu_{AA} = 1, \nu_{AB} = 0$, *(5)* $\nu_B = 0, \nu_{AA} = 1, \nu_{AB} = 0$, *(6)* $\nu_B = 0, \nu_{AA} = 1, \nu_{AB} \in (0,1)$, *(7)* $\nu_B = 0, \nu_{AA} = 1, \nu_{AB} = 1$.

Proof. For (1), set $\nu_B = 1, \nu_{AA} = \nu_{AB} = 0$, which, by Proposition 2, implies cutpoints $\underline{\alpha}^1$ and $\overline{\alpha}^1$ for first period official behavior. By Lemma 6, this official behavior implies voter cutpoints α^{AA}, α^B, and α^{AB}. The voter behavior in part (1) is optimal for any $\alpha_V \leq \alpha^{AA}$.

For (2), set $\nu_B = 1$, $\nu_{AB} = 0$, and take any $\nu_{AA} \in (0,1)$. Proposition 2 gives cutpoints $\underline{\alpha}^1$ and $\overline{\alpha}^1$, and Lemma 6 gives the resulting voter cutpoints α^{AA}, α^B, and α^{AB}. For $\alpha_V = \alpha^{AA}$ it is optimal to play $\nu_B = 1$ and $\nu_{AB} = 0$ and since this voter type is indifferent after observing $x = A$ and $\omega = A$, he can mix using the particular $\nu_{AA} \in (0,1)$ that was used to generate α^{AA}. The arguments for (3)-(7) are similar.■

Lemma 8 *Any equilibrium voter strategy must be one of the 7 types in Lemma 7.*

Proof. Consider any (possibly mixed) voter strategy ν. Given ν, Proposition 2 implies cutpoints $\underline{\alpha}^1$ and $\overline{\alpha}^1$ for first period official behavior. Given these cutpoints, Lemma 6 characterizes cutpoints for voter behavior. It is straightforward to check that the only voter strategies ν that are compatible with these cutpoints are the types listed in Lemma 7.■

Lemma 9 *For any $\alpha_V \in \left[\underline{\alpha}^2, \overline{\alpha}^2\right]$ there exists an equilibrium.*

We set up a correspondence $\lambda(z)$:$[0,3] \to \left[\underline{\alpha}^2, \overline{\alpha}^2\right]$. Each ν from Lemma 7 is specified by a unique value of z, and we use an intermediate value theorem for upper semi-continuous correspondences to show that for any $\alpha \in \left[\underline{\alpha}^2, \overline{\alpha}^2\right]$ there is some z such that $\alpha \in \lambda(z)$, i.e., there is an equilibrium. Define $\lambda(z)$ as follows

$$\lambda(z) = \begin{cases} \alpha \in \left[\underline{\alpha}^2, \overline{\alpha}^2\right] : \nu_B = 1, \nu_{AA} = \nu_{AB} = 0 \text{ is an equilibrium} & \text{for } z = 0 \\ \alpha \in \left[\underline{\alpha}^2, \overline{\alpha}^2\right] : \nu_B = 1, \nu_{AA} = z, \nu_{AB} = 0 \text{ is an equilibrium} & \text{for } z \in (0,1) \\ \alpha \in \left[\underline{\alpha}^2, \overline{\alpha}^2\right] : \nu_B = \nu_{AA} = 1, \nu_{AB} = 0 \text{ is an equilibrium} & \text{for } z = 1 \\ \alpha \in \left[\underline{\alpha}^2, \overline{\alpha}^2\right] : \nu_B = 2 - z, \nu_{AA} = 1, \nu_{AB} = 0 \text{ is an equilibrium} & \text{for } z \in (1,2) \\ \alpha \in \left[\underline{\alpha}^2, \overline{\alpha}^2\right] : \nu_B = 0, \nu_{AA} = 1, \nu_{AB} = 0 \text{ is an equilibrium} & \text{for } z = 2 \\ \alpha \in \left[\underline{\alpha}^2, \overline{\alpha}^2\right] : \nu_B = 0, \nu_{AA} = 1, \nu_{AB} = -2 + z \text{ is an equilibrium} & \text{for } z \in (2,3) \\ \alpha \in \left[\underline{\alpha}^2, \overline{\alpha}^2\right] : \nu_B = 0, \nu_{AA} = 1, \nu_{AB} = 1 \text{ is an equilibrium} & \text{for } z = 3 \end{cases}$$

We establish several properties of $\lambda(z)$. By Lemma 7, $\lambda(z)$ is nonempty, $\forall z \in [0,3]$. In the proof of Lemma 6, we established that the difference in the voter's expected utility from re-electing versus removing the official, conditional on her beliefs after observing the first period policy choice and outcome, is a linear and hence monotonic function of α_V within the region $\left[\underline{\alpha}^2, \overline{\alpha}^2\right]$. Thus if ν is an equilibrium for α_{v1} and for α_{v2} it must be an equilibrium for all $\alpha_V \in (\alpha_{v1}, \alpha_{v2})$, i.e., $\lambda(z)$ is convex valued.

The last step is to show that $\lambda(z)$ is upper semi-continuous. Suppose $z_n \to \tilde{z}$ and $y_n \to \tilde{y}$ where $y_n \in \lambda(z_n), \forall n$. We need to show that $\tilde{y} \in \lambda(\tilde{z})$. To do this we define $\underline{\alpha}^1(\tilde{z})$ and $\overline{\alpha}^1(\tilde{z})$ to be the cutpoints for incumbent policy choice from Proposition 2 given voter behavior specified by $\tilde{z}$, and similarly define $\underline{\alpha}^1(z_n)$ and $\overline{\alpha}^1(z_n)$ based on z_n. We show that $\underline{\alpha}^1(z_n) \to \underline{\alpha}^1(\tilde{z})$ and $\overline{\alpha}^1(z_n) \to \overline{\alpha}^1(\tilde{z})$. We similarly use Lemma 6 to define cutpoints for voter behavior $\alpha^{AA}(\tilde{z}), \alpha^{AB}(\tilde{z}), \alpha^{B}(\tilde{z})$, $\alpha^{AA}(z_n), \alpha^{AB}(z_n)$, and $\alpha^{B}(z_n)$, then show that $\alpha^{AA}(z_n) \to \alpha^{AA}(\tilde{z})$, $\alpha^{AB}(z_n) \to \alpha^{AB}(\tilde{z})$, and $\alpha^{B}(z_n) \to \alpha^{B}(\tilde{z})$. Finally we show that convergence of the voter cutpoints ensures that $\tilde{y} \in \lambda(\tilde{z})$.

To show $\underline{\alpha}^1(z_n) \to \underline{\alpha}^1(\tilde{z})$ recall from the proof of Proposition 2 that $\underline{\alpha}^1(z) = \{\max \alpha_I \in [0,1] : U(A|A) - U(B|A) \leq 0\}$. From Lemma 1, $U(A|A) - U(B|A)$ is continuous and strictly increasing in α_I. We now note that it has a slope that is strictly bounded away from zero, since for $\alpha_I < \underline{\alpha}^2$, $\pi(\phi_R q + \phi_A) + (1-\pi)(\phi_R(1-q) + \phi_A) = \phi_R(\pi q + (1-\pi)(1-q)) + \phi_A < 1$ and for $\alpha_I \in \left(\underline{\alpha}^2, \overline{\alpha}^2\right)$, $\pi(\phi_B q - \phi_A(1-q)) + (1-\pi)(\phi_B(1-q) - \phi_A q) = \phi_B(\pi q + (1-\pi)(1-q)) - \phi_A(\pi(1-q) + (1-\pi)q) > -1$. Let $d > 0$ be a lower bound on the slope. From Eq (1) in the main text the re-election probability difference $[r_A(s) - r_B(s)]$ is a continuous function of the voter strategy ν, and ν_B, ν_{AA} and ν_{AB} are continuous functions of z. Thus the utility difference for z_n, which we denote as $U_n(A|A) - U_n(B|A)$ converges pointwise to the utility difference for $\tilde{z}$, which we denote as $\tilde{U}(A|A) - \tilde{U}(B|A)$. And, if we pick an $\epsilon > 0$ and let $\delta = \epsilon d$ there exists an N such that for all

$n > N$, at $\underline{\alpha}^1(\tilde{z})$, $U_n(A|A) - U_n(B|A) < \delta$ and thus $\left|\underline{\alpha}^1(z_n) - \underline{\alpha}^1(\tilde{z})\right| < \frac{\delta}{d} = \epsilon$.[20] By a similar argument $\overline{\alpha}^1(z_n) \to \overline{\alpha}^1(\tilde{z})$.

We now establish convergence of the voter cutpoints, starting with $\alpha^B(z_n) \to \alpha^B(\tilde{z})$. From the proof of Lemma 6, for a given $\underline{\alpha}^1(\tilde{z})$ and $\overline{\alpha}^1(\tilde{z})$, $\tilde{U}(old|B) - \tilde{U}(new)$, the voter's utility difference for re-electing versus removing the incumbent when $x = B$, is strictly positive for $\alpha_V = \underline{\alpha}^2$. Likewise, given $\underline{\alpha}^1(\tilde{z})$ and $\overline{\alpha}^1(\tilde{z})$ for a voter with preference $\alpha_V = \overline{\alpha}^2$, $\tilde{U}(old|B) - \tilde{U}(new)$ is strictly negative. Let $d > 0$ be the slope of this utility difference function, and note that $U_n(old|B) - U_n(new)$ converges pointwise to $\tilde{U}(old|AA) - \tilde{U}(new)$ so we can pick an N such that for $n > N$ the absolute value of the slope of $U_n(old|B) - U_n(new)$ is greater than $\frac{d}{2}$. Then the same type of argument we used to show that $\underline{\alpha}^1(z_n) \to \underline{\alpha}^1(\tilde{z})$ can be used to show that $\alpha^B(z_n) \to \alpha^B(\tilde{z})$. The arguments for $\alpha^{AA}(z_n) \to \alpha^{AA}(\tilde{z})$, $\alpha^{AB}(z_n) \to \alpha^{AB}(\tilde{z})$ are similar, with the only difference being the fact that the slope of the utility difference is positive rather than negative.

To show that $\tilde{y} \in \lambda(\tilde{z})$ we consider two cases. First suppose $\tilde{z} \notin \{0, 1, 2, 3\}$, and consider the specific subcase $\tilde{z} \in (0, 1)$. Then $z_n \to \tilde{z}$ implies that there exists an N such that $\forall n > N$, $z_n \in (0, 1)$. For such $z_n \in (0, 1)$, by Lemma 7 and the definition of $\lambda(\cdot)$, $y_n = \alpha^{AA}(z_n)$ and $\alpha^{AA}(z_n) \to \alpha^{AA}(\tilde{z}) = \lambda(\tilde{z})$. The argument for $\tilde{z} \in (1, 2)$ or $\tilde{z} \in (2, 3)$ is similar, using α^B or α^{AB}.

Now suppose $\tilde{z} \in \{0, 1, 2, 3\}$, and consider the specific case $\tilde{z} = 1$. If $z_n < 1$ then $y_n = \alpha^{AA}(z_n)$, if $z_n > 1$ then $y_n = \alpha^B(z_n)$, and if $z_n = 1$ then $y_n \in \left[\alpha^{AA}(z_n), \alpha^B(z_n)\right]$. If an infinite number of elements of z_n are strictly less than 1 then if $y_n \to \tilde{y}$ we have $\tilde{y} = \alpha^{AA}(\tilde{z}) \in \lambda(\tilde{z})$. Likewise if there are an infinite number of elements of z_n that are strictly greater than 1 then if $y_n \to \tilde{y}$ we have $\tilde{y} = \alpha^B(\tilde{z}) \in \lambda(\tilde{z})$. If neither of these conditions holds then there exists an N such that $\forall n > N$, $z_n = 1$ so $y_n \in \lambda(\tilde{z})$, $\forall n > N$ and since $\lambda(\tilde{z}) = \left[\alpha^{AA}(\tilde{z}), \alpha^B(\tilde{z})\right]$ is a closed set if $y_n \to \tilde{y}$ then

[20] If $\underline{\alpha}^1(\tilde{z}) = 0$ and $\tilde{U}(A|A) - \tilde{U}(B|A)$ for $\alpha_V = 0$ then there exists an N such that for all $n > N$, $U_n(A|A) - U_n(B|A) > 0$ and thus $\underline{\alpha}^1(z_n) = 0$.

$\tilde{y} \in \lambda(\tilde{z})$.

By Proposition 4, $\underline{\alpha}^2 \in \lambda(0)$ and $\overline{\alpha}^2 \in \lambda(3)$ so by an intermediate value theorem for correspondences (de Clippel n.d., Lemma 2) for any $\alpha \in \left(\underline{\alpha}^2, \overline{\alpha}^2\right)$ there is some z such that $\lambda(z) = \alpha$.■

References

[1] Ansolabehere, Stephen, and James M. Snyder. 2000. "Valence Politics and Equilibrium in Spatial Election Models." *Public Choice* 103(3): 327-336.

[2] Ansolabehere, Stephen, James M. Snyder, Jr., and Charles Stewart III. 2001. "Candidate Positioning in U.S. House Elections." *American Journal of Political Science* 45(January):136-159.

[3] Aragones, Enriqueta, and Thomas R. Palfrey. 2002. "Mixed Equilibrium in a Downsian Model with a Favored Candidate." *Journal of Economic Theory* 103(1): 131-161.

[4] Austen-Smith, David. 1992. "Explaining the Vote: Constituency Constraints on Sophisticated Voting." *American Journal of Political Science* 36(February): 68-95.

[5] Baron, David P. 1994. "Electoral Competition with Informed and Uninformed Voters." *American Political Science Review* 88(March):43-57.

[6] Besley, Timothy, and Stephen Coate. 2003. "Elected Versus Appointed Regulators: Theory and Evidence." *Journal of the European Economic Association* 1(September): 1176-1206.

[7] Callander, Steven. 2004. "Political Motivations." Northwestern typescript.

[8] Calvert, Randall L. 1985. "Robustness of the Multidimensional Voting Model: Candidate Motivations, Uncertainty, and Convergence." *American Journal of Political Science* 29(February): 69-95.

[9] Canes-Wrone, Brandice, Michael C. Herron, and Kenneth W. Shotts. 2001. "Leadership and Pandering: A Theory of Executive Policymaking." *American Journal of Political Science* 45(July):532-550.

[10] Coate, Stephen, and Stephen Morris. 1995. "On the Form of Transfers to Special Interests." *Journal of Political Economy* 103(December):1210-35.

[11] de Clippel, Geoffroy. n.d. "An Axiomatization of the Inner Core Using Appropriate Reduced Games." Brown University typescript.

[12] Downs, Anthony. 1957. *An Economic Theory of Democracy*. New York: HarperCollins.

[13] Downs, George W., and David M. Rocke. 1994. "Conflict, Agency, and Gambling for Resurrection: The Principal-Agent Problem Goes to War." *American Journal of Political Science* 38(May):362-380.

[14] Duggan, John. 2000. "Repeated Elections with Asymmetric Information." *Economics and Politics* 12(July):109-136.

[15] Enelow, James, and Melvin J. Hinich. 1981. "A New Approach to Voter Uncertainty in the Downsian Spatial Model." *American Journal of Political Science* 25(August): 483-93.

[16] Fearon, James D. 1999. "Electoral Accountability and the Control of Politicians: Selecting Good Types versus Sanctioning Poor Performance." In *Democracy, Accountability, and Representation,* ed. Bernard Manin, Susan Stokes, and Adam Przeworski. Cambridge: Cambridge University Press.

[17] Ferejohn, John. 1986. "Incumbent Performance in Office." *Public Choice* 50(July):1-26.

[18] Fiorina, Morris P. (with Samuel J. Abrams and Jeremy C. Pope). 2004. *Culture War? The Myth of a Polarized America*. New York: Pearson Longman.

[19] Fox, Justin. 2005. "Government Transparency and Policymaking." Yale typescript.

[20] Gordon, Sanford C., and Gregory A. Huber. 2002. "Citizen Oversight and the Electoral Incentives of Criminal Prosecutors." *American Journal of Political Science* 46(April):334-351.

[21] Grofman, Bernard. 2004. "Downs and Two-Party Convergence." *Annual Review of Political Science* 7(May): 25-46.

[22] Groseclose, Tim. 2001. "A Model of Candidate Location When One Candidate Has a Valence Advantage." *American Journal of Political Science* 45(October): 862-886.

[23] Harrington, Joseph E., Jr. 1993. "Economic Policy, Economic Performance, and Elections." *American Economic Review* 83(March):27-42.

[24] Maskin, Eric, and Jean Tirole. 2004. "The Politician and the Judge: Accountability in Government." *American Economic Review* 94(September):1034-1054.

[25] Meirowitz, Adam. 2005. "Probabilistic Voting and Accountability in Elections with Uncertain Policy Constraints." Princeton typescript.

[26] Morton, Rebecca B. Forthcoming. *Analyzing Elections*. New York: W. W. Norton.

[27] Prat, Andrea. 2005. "The Wrong Kind of Transparency." *American Economic Review* 95(June): 862-877.

[28] Rogoff, Kenneth, and Anne Sibert. 1988. "Elections and Macroeconomic Policy Cycles." *Review of Economic Studies* 50(January):1-16.

[29] Snyder, James M. Jr., and Michael M. Ting. 2002. "An Informational Rationale for Political Parties." *American Journal of Political Science*. 46(January): 90-110.

<u>Figure 1</u>: Second period behavior of official

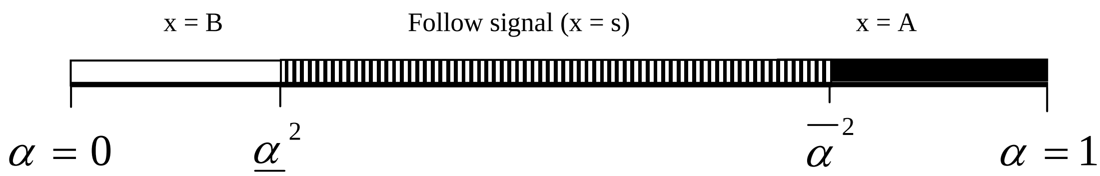

<u>Figure 2</u>: First period behavior of incumbent

(a) Official who is rewarded for choosing x = B

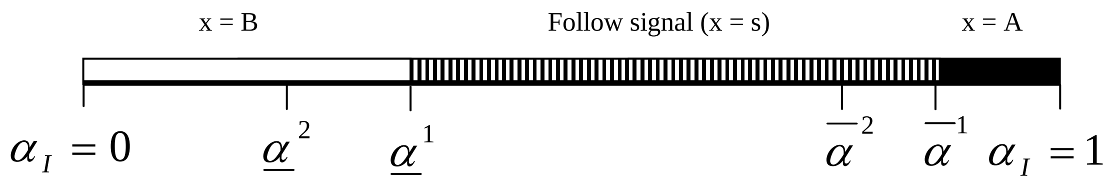

(b) Official who is rewarded for choosing x = A

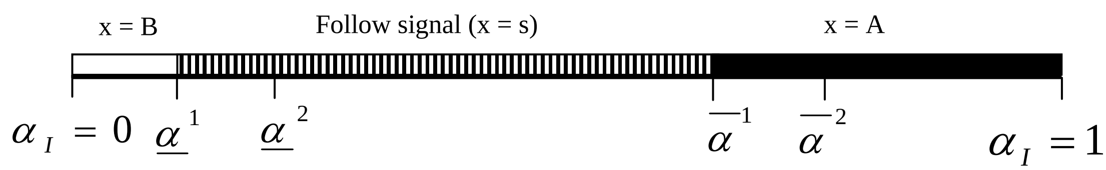

(c) Moderation by official who is rewarded for following signal

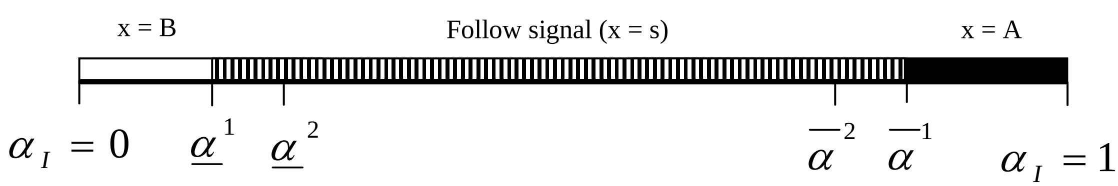

<u>Figure 3</u>: Equilibrium electoral incentives in Example 1, as a function of voter preferences α_v

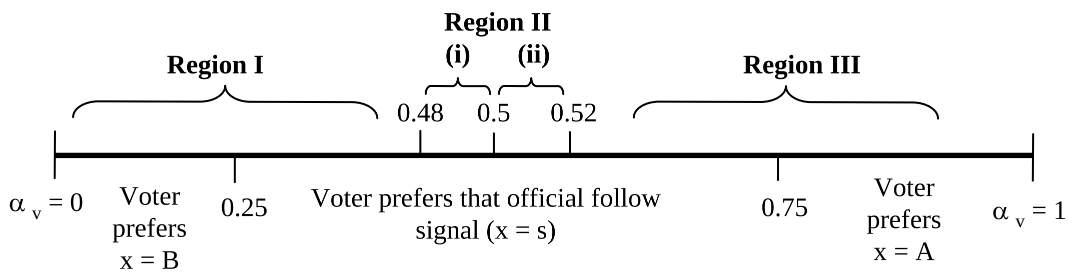

Region I: Official wins re-election if and only if x = B

Region II: Voter uses mixed strategy
- In Region II(i), official has electoral incentive to choose x = B
- In Region II(iii), official has electoral incentive to choose x = A

Region III: Official wins re-election if and only if x = A

<u>Figure 4</u>: Equilibrium electoral incentives in Example 2, as a function of voter preferences α_v

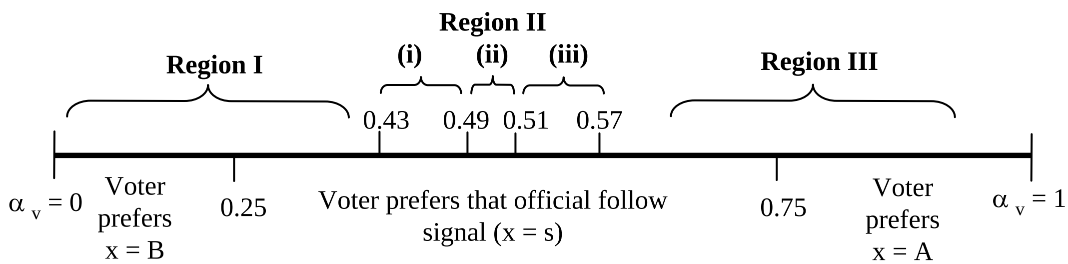

Region I: Official wins re-election if and only if x = B

Region II: Voter behavior when x = A may depend whether policy succeeds or fails
• Official is always more likely to win re-election after policy success than after policy failure
• In Region II(i), official has incentive to choose x = B regardless of private signal
• In Region II(ii), official has incentive to follow signal and hence moderation occurs
• In Region II(iii), official has incentive to choose x = A regardless of private signal

Region III: Official wins re-election if and only if x = A