Electoral punishment as signaling in subnational elections

Michael Kellermann*

September 30, 2008

Abstract

It is a well-established empirical regularity that parties in federal office suffer setbacks in state-level elections. Many authors attribute this to a desire on the part of voters to balance the policy preferences of the federal incumbent. In this paper, I consider an alternative explanation with a long tradition in the literature: voters punish the party of the federal incumbent in state elections in order to send a signal to the federal government. I construct a simple signaling model to formalize this intuition, which predicts that under most circumstances signaling can occur at only one level of government. I estimate a statistical model allowing for electoral punishment using data from German elections and find support punishment at the state level, rather than the punishment at both levels implied by balancing theories.

*Ph.D. candidate, Department of Government and Institute for Quantitative Social Science, Harvard University.

1 Introduction

On May 22, 2005, voters in the German state of Nordrhein-Westphalen defeated the incumbent Social Democratic (SPD)-Green coalition government, bringing the Christian Democrats (CDU) to power in the state for the first time in 39 years. The vote was widely reported in the media as a message to the federal SPD-Green coalition led by Gerhard Schroeder. He responded to the defeat by entering campaign mode, calling for new federal elections that provoked a minor constitutional crisis and led to Schroeder's resignation in the fall of 2005 (Nagel, 2006). While the scale of the defeat was perhaps surprising, the unpopularity of the federal government was not. National opinion polls taken before the election provided ample evidence that the governing coalition was in trouble, yet the message sent by voters in Nordrhein-Westphalen finally prompted the government to act.

The tendency for parties in power at the federal level to lose support in elections at the state level is one of the more robust stylized facts in comparative politics, but the mechanisms underpinning this regularity remain unclear. The 2005 election in Nordrhein-Westphalen certainly appears to be an example of voters sending a message to Berlin, tartly commenting on the performance of the Schroeder government. To paraphrase van der Eijk, Franklin and Marsh (1996), however, "why comment on national politics in [subnational] elections?" In principle, there are many ways for voters to express discontent; it is not clear why state elections provide an attractive venue. Relatively few efforts have been made to formalize the intuition behind voting as a means of communication (Piketty, 1999, 2000). In contrast, several authors attribute punishment in subnational elections to a desire on the part of voters to balance the governments at various levels against each other in an attempt to achieve more moderate policy outcomes; the formalization of this logic is more advanced (Alesina and Rosenthal, 1996; Kedar, 2006).

In this paper, I formalize the "punishment as message-sending" argument and derive implications that in principle could allow this mechanism to be distinguished from "punishment as balancing". I briefly review the extensive empirical literature which demonstrates that federal governing parties suffer losses in subnational elections, as well as the formal literatures on balancing and signaling in elections. I then present a simple model of the strategic interaction between an incumbent government at one level and a representative voter with private information about her evaluation of the government's performance. The voter faces a choice between the parties of the

federal government and opposition in an election for some other level of government. Under some conditions, an informative equilibrium exists in which voters credibly signal their displeasure with an incumbent governing party at one level by voting against it in elections to other levels of government. The sustainability of such an equilibrium depends on the relative importance of the two elections, and the costs of voting for a less-preferred party at one level in exchange for better results from the incumbent party at the other level.

One implication of this signaling model is that, under most conditions, signaling that results in electoral punishment can only occur in one direction. This is in contrast to balancing models, which imply that the party of an incumbent government at one level should lose support in elections to the other level (although these losses would typically be greater in state elections). To evaluate these competing predictions, I develop a statistical model of party vote shares as a function of the party's normal vote and its status as incumbent or opposition at the other level of government. The results of this model applied to German data confirm that federal incumbents suffer in subnational elections, but find no evidence that state incumbents suffer in federal elections. This result argues in favor of a signaling-based understanding of electoral punishment in subnational elections.

2 What happens in subnational elections?

Parties that control the government at the federal level often perform poorly in other elections. This pattern has been noted in off-cycle elections for the U.S. Congress (Mebane and Sekhon, 2002) and other federal legislatures (Shugart, 1996). Incumbent parties at the federal level suffer relative losses in non-federal elections (Rodden and Wibbels, 2005); electoral punishment of governing parties has been seen in state elections in the United States (Chubb, 1988) and Germany (Lohmann, Brady and Rivers, 1997; Gaines and Crombez, 2004; Kedar, 2006; Hainmueller and Kern, 2006), and provincial elections in Canada (Erikson and Filippov, 2001; Gelineau and Belanger, 2005). Governing parties also tend to lose seats in elections to the European Parliament (Reif and Schmitt, 1980; Marsh, 1998) and in local elections.

A variety of explanations have been advanced to account for the empirical regularity of governing party losses in non-federal elections. Reif and Schmitt (1980) propose three hypotheses, suggesting (1) that the composition of the electorate differs in federal and state elections, (2) that voters who vote strategically in federal elections may be willing to vote sincerely in state elections because

the stakes are lower, and (3) that voters may use state elections to signal their displeasure to the governing federal parties. This third argument, that voters use state elections to send messages to federal incumbents, is the mechanism explored in this paper. As discussed below, other authors have suggested instead that governing parties lose in subnational elections because voters are attempting to moderate the policy outcomes advocated by the incumbent federal government. Within the formal literature, little attention has been given to the relative importance of these motivations, in part because the signaling mechanism has not been formalized to the same extent as the balancing mechanism.

2.1 Balancing theories of electoral punishment

At the heart of balancing theories is an assumption that policy outcomes represent compromises between various institutional actors, and that voters take this into account when choosing candidates or parties to support. Models incorporating this assumption have been proposed in various contexts by Alesina and Rosenthal (1996); Mebane and Sekhon (2002); Erikson and Filippov (2001) and Kedar (2006), among others. Policy outcomes are assumed to lie somewhere between the positions promoted by various institutional actors. As a result, voters do not choose between the policies offered by parties in an election, but instead choose between those positions as modified by the positions of incumbents not facing the voters. This gives moderate voters in particular an incentive to balance against the incumbents in order to obtain policies that are closer to their preferred outcomes than could be obtained from supporting the governing party, even if those voters would prefer the policies of the incumbent party to the policies of the opposition party if forced into an up-or-down choice.

Since balancing models will provide a contrast to the signaling model developed below, it is useful to review the intuition in some detail. In a two-party spatial voting model with sincere voters, voters can be ordered according to their policy preferences along a single dimension from left to right. Voters choose between the policy positions L and R offered by the two parties. Voters closer to L than R vote for L and vice versa. Thus, the continuum of voters is divided by a cutpoint located halfway between the two parties, as illustrated in the first line of Figure 1. Balancing models complicate this picture by assuming that policy outcomes are a weighted average of the positions advocated by the parties in office at each level. Thus, when a right party is in office at the

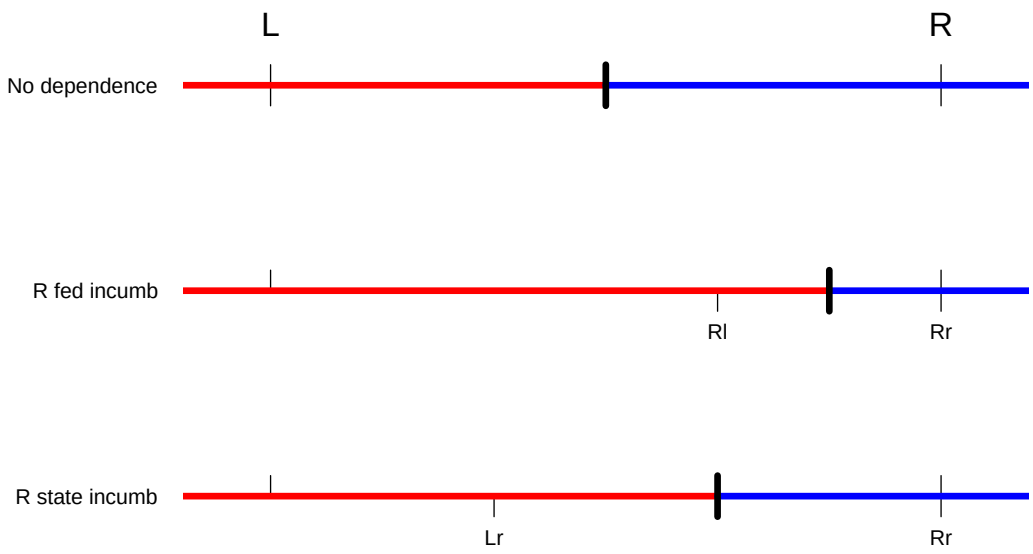

Figure 1: Intervals of voters for parties under a simple balancing model. Based on the model in Erikson and Filippov (2001), assuming that the federal election is twice as important as the state election.

federal level, voters in a state election are not choosing between L and R, the positions advocated by the parties at the state level. Instead, they are choosing between Rl and Rr, the net policy outcome of a right-federal, left-state combination versus a right-federal, right-state combination, as shown in the second line of Figure 1. These new alternatives shift the cutpoint between the right and left parties at the state level toward the party in office at the federal level. As a result, the parties of federal incumbents lose votes in state elections.

Balancing models thus depend crucially on the assumption that the policy positions considered by voters when deciding how to vote are weighted averages of the policies advocated at different levels of government or in different branches at the same level. Balancing is most plausible in off-cycle legislative elections in presidential or semi-presidential systems, such as mid-term elections for the U.S. Congress. In systems where implementing policy requires the agreement of multiple institutional actors, outcomes are collectively determined through some formal process. Punishment motivated by balancing is less plausible in elections where voters choose state governments,

but those governments are represented directly or indirectly at the federal level, as in the German Bundesrat or the U.S. Senate before 1920. In these cases, voters are simultaneously electing federal legislators and state legislatures. They have one instrument with which to accomplish two tasks: balancing the incumbent government at the federal level and choosing a government at the state level. Balancing is least plausible in state or local elections in those systems where lower-level governments have no formal representation at the federal level. Canada and Australia are two examples of such systems. In these cases, politicians elected at the state or provincial level literally do not have a vote on policies decided at the federal level. To view policy outcomes as a weighted average of the two levels, one must either assume a strong interaction between the policies implemented at each level or rely on informal institutions to produce policy moderation.[1]

Empirical evidence in support of balancing models is largely derived from the outcomes of elections believed to be less important to voters. As noted above, there is overwhelming evidence that party in office at the federal level suffers electoral losses at the subnational level. The logic of balancing models, however, operates in the other direction as well: state incumbents should affect federal election outcomes. Returning to Figure 1, the lowest line illustrates the decision facing voters in a federal election when their state government is controlled by the right party. In this case, they are not choosing between L and R. Instead, they must choose between Lr and Rr, a left-federal, right-state combination and a right-federal, right-state combination. This also shifts the cutpoint to the right, suggesting electoral losses for the party in office at the state level. If state elections are less important, the magnitude of the shift is smaller. Nevertheless, a shift should be observed if voters use a balancing logic. Evidence for such reverse balancing is difficult to find.

Other empirical implications of balancing models have not been fully investigated. First, if voters do seek to balance one elected body against others in order to achieve more moderate outcomes, one might expect that they would attempt to balance in simultaneous as well as sequential elections. One potential problem is that, to use the metaphor of Kedar (2006), neither end of the seesaw is fixed in simultaneous elections. This is true in elections that are expected to be close, but the balancing logic would seem to apply when the outcome of one of the elections is fairly certain.

[1]Possible examples of such informal institutions are the First Minister's meetings in Canada, which bring together the leaders of the federal and provincial governments for discussions on various issues.

Second, balancing models rely on the assumption that some voters lie in the interval between the policy positions advocated by the major parties. This will almost certainly be the case if voters are sincere, supporting the party that produces the weighted average policy closest to their ideal point. There will almost always be some moderate voters with ideal points in the interval between the major parties. Moderate voters, however, need not be pivotal in all regions. If voters were more instrumental, caring only about the policies actually implemented rather than the policies for which they are voting, one would not expect to see balancing behavior in subnational units where the median voter is extreme relative to the major parties. The median voter in the province of Alberta, known as a right-wing province in Canada, almost certainly thinks that a Conservative government in Ottawa is too moderate. She would have no inclination to balance against that government by electing a provincial government led by the Liberals or the NDP that would shift policy even further from her ideal point.

The lack of evidence for some implications of balancing models suggests it is worth looking for other models that predict electoral punishment in second-order elections. In particular, the absence of evidence that voters balance at both levels is problematic for standard balancing models. This paper attempts to address this problem by developing a signaling model for elections which implies that punishment will only occur at one level of government.

2.2 Signaling theories of electoral punishment

A small but growing literature exists on voting as signaling. As Piketty (2000) notes, the act of voting is not simply about electing candidates to office; it also allows the electorate to communicate with their elected officials. Voting lends itself naturally to the framework of signaling games. These have received significant attention in many areas of political science (Banks, 1991), but have been relatively unexplored in electoral politics.[2] In thinking about models of voting as signaling, one must consider the types of private information that voters might wish to communicate. Voters could send signals regarding their own unobserved policy preferences or about their private evaluation of government performance. These private evaluations might be thought of as corresponding to

[2]In a closely related area, however, Lohmann (1994) considers information transmission through other forms of costly political activity, such as signing petitions, writing letters, or engaging in street demonstrations.

valence aspects of electoral competition: attributes that are valued by all voters, regardless of their ideological preferences (Stokes, 1963). Models of voting as signaling also differ in their electoral context, with some focusing on signaling within a single election and others on signaling across elections. Finally, models differ in the identity of the actor(s) receiving the signal. In some cases, voters are sending signals to politicians, while in others they are sending signals to each other.

Razin (2003) models the process of signaling within a single election where voters imperfectly observe a common shock to the preferences of the electorate and politicians are responsive to information about the common shock that can be gleaned from aggregate election results. In certain circumstances, voters have an incentive to vote against their preferred candidate in order to promote policy moderation. This incentive is strongest when the election is likely to produce a lopsided winner. Fowler and Smirnov (2007) report evidence from survey experiments suggesting that such behavior does occur. When told that one party is likely to win an election by a large margin, respondents who described themselves as moderate were more likely to support the party expected to lose the election.

A larger number of models have explored the incentives to signal across repeated elections for the same office. Among models in which the distribution of preferences is unknown to politicians or other voters, Piketty (2000) presents a model in which some voters do not vote sincerely in a first election in order to aid coordination in a second, more important election in the future. Shotts (2006) develops a two-period model with private information about voter preferences in which some voters abstain in the first election in order to promote policy moderation in the second election. Callander (2007) also looks at sequential elections but develops a model in which voters have private information about the quality of candidates. Voters in early rounds use their votes to send signals about the electability of candidates to voters in future rounds.

Meirowitz and Tucker (2007), in contrast, focuses on signaling across different types of elections. The empirical motivation for their model is quite similar to the situation addressed here. They consider the Russian parliamentary and presidential elections of 1995-96, in which voters supported opposition parties for the Duma before re-electing Boris Yeltsin to the presidency. In their model, the distribution of preferences is unknown and presidential candidates compete over valence characteristics. Voters have an incentive to misrepresent their preferences in the parliamentary election in order to encourage their preferred candidates to work harder in the presidential election cam-

paign. Signaling is possible in this game because the outcome of the parliamentary election is less important to voters than the presidency; the equilibrium can be thought of as producing partial pooling on weaker parliamentary parties.

While the presidential candidates modeled by Meirowitz and Tucker (2007) exert effort in response to voter signals, the uncertainty which makes signaling necessary is fundamentally about the distribution of voter preferences. In the context of subnational elections in countries where the distribution of preferences is fairly stable and well known, it may be more useful to think of the private information available to voters as representing private evaluations of the perceived quality of incumbent governments rather than information about voter policy preferences. In other words, governments may not know exactly how discontent (or, perhaps less frequently, how satisfied) voters are with their performance. To the extent that government effectiveness is valued by all voters and can be increased through effort exerted by the government, there is an obvious incentive for voters to pretend that they are unhappy so long as they can do so at no cost; this would naturally lead governments to discount such claims as not credible. In the next section, I develop a simple signaling model to identify the conditions under which voter dissatisfaction can be credibly communicated to incumbent governments.

3 An extremely simple signaling model

3.1 Setup

To analyze the potential for signaling in subnational elections, I use an extremely stylized model of the strategic interaction between voters and incumbent governments. The basic logic behind this voting-as-signaling approach is that voters know how satisfied they are with the incumbent's performance. Voters will communicate their information at some cost if it affects the behavior of the incumbent. To formalize this intuition, I assume the following sequence of play in the game:

1. Nature reveals to the voter her type; that is, whether she is satisfied or dissatisfied with the federal incumbent's performance.

2. After observing her type, the voter chooses to support either the party of the federal incumbent or the federal opposition in a non-federal election.

3. After observing the results of this election, the federal incumbent either exerts additional effort or remains at the current level.

This sequence of play, represented in the game tree in Figure 2, implies a simple signaling game between the voter and the federal incumbent.[3]

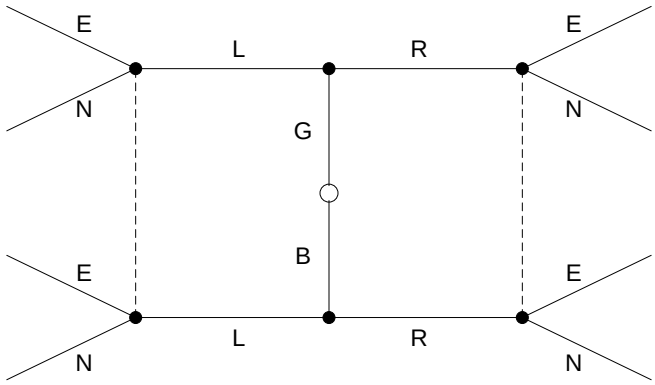

Figure 2: Game tree for signaling game. Nature chooses good or bad state, observed by voter. Voter votes left or right. Government only observes vote, either exerts effort or does nothing.

I assume that voters receive utility from both the policies enacted by governments at each level as well as the effort exerted by those governments. Voters have standard single-peaked preferences over policies and their utility is strictly increasing in governmental effort. Let m be the ideal point of the representative voter, $g \in \{l, r\}$ be the position of the government (such that l is the position of the left party and r is the position of the right party), and e be the effort exerted by the government. To distinguish between the utility from different levels of government, let superscript f indicate the federal level and s indicate the state level. Assuming quadratic-loss utility for simplicity, the utility to the voter from the federal government is

$$u^f(g^f, e^f) = (-(m - g^f)^2 + e^f) \tag{1}$$

[3]The game tree is identical, for example, to the Beer-Quiche game presented in Cho and Kreps (1987), although the payoffs (defined below) are different.

and analogously for the utility from the state government.

Including effort as an argument in the voter's utility function moves beyond purely policy-oriented conceptions of political competition. Effort in this model can be interpreted as a measure of government competence, with increasing effort implying that the government's policies are implemented more effectively whatever those policies happen to be. Effort in the model can also be interpreted as a measure of the government's ability to keep its nose clean, avoiding scandals and other embarrassments. Both interpretations fall under the broad heading of valence characteristics. Given this interpretation of effort, it is natural to include the term additively in the voter's utility function.[4] This formulation is consistent with standard models of party competition taking valence characteristics into account (Ansolabehere and Snyder, 2000; Groseclose, 2001).

Overall voter utility is assumed to be a weighted average of the utilities derived separately from each level of government. In other words, the overall utility is

$$u(g^f, g^s, e^f, e^s) = \alpha u^s(g^s, e^s) + (1-\alpha)u^f(g^f, e^f) \tag{2}$$

where $\alpha \in [0, 1]$ is the relative importance of the state government. In most federal countries, one would expect $\alpha < 0.5$, reflecting the greater importance of the federal level.

The value of effort is dependent on the realization of a state variable observed by voters but not by the government. The state variable is $\theta \in g, b$, where g represents a good state for the government (the voters are happy) and b represents a bad state (the voters are unhappy). Under this formulation, e_g represents the value of the government's existing level of effort in the good state and e_b is the value of the government's existing effort in the bad state. If the government exerts additional effort as a result of the election, the resulting payoff e_n does not depend on the state. The payoff given additional effort can be thought of as the expected value of whatever new initiatives the government chooses to implement.

[4]Complications may arise, however, if one extended the horizon beyond a single interaction between incumbents and voters. A voter who dislikes the policy position of the government might not want that government to exert more effort if that additional effort increases the probability that it will be re-elected in the future, even if those voters would be better off with more effort in the short term.

Turning to the incumbent government, I assume that it has the same utility function over policy and effort as the voters do. The only difference is that the government must bear a cost c if it chooses to exert additional effort after observing the results of the election, while it bears no cost if it makes no changes. To illustrate this framework, one can think of the incumbent government as being graded by the electorate on its valence characteristics. It doesn't know whether its current grade is a B- or a D, but it does know that if it tries harder (by launching a new initiative, reshuffling the cabinet, or some other action) it can expect to get a B+ on average. For some costs of effort, the incumbent would be better off moving from a D to a B+, but moving from a B- to a B+ would not be worthwhile.

This structure for the model implies several simplifications which enable a clear focus on the signaling aspects of the game. In particular, the policy position implemented by the incumbent is determined at the end of the game and so is chosen to maximize its own utility, unlike the models of signaling with responsive politicians discussed above. The candidates for government at the state level are not strategically relevant actors in the model at all; their choices of policy and effort are not influenced by the behavior of voters or the incumbents at the other level.

In summary, there are two key differences between the signaling model described above and standard balancing models. As discussed above, balancing models assume that policy is a weighted average of the positions of the state and federal governments. Moreover, balancing requires voters to choose between different weighted averages, which are themselves functions of the positions of the parties at the different levels. The model presented here, in contrast, assumes that voters know the locations of the parties at each level and evaluate those positions separately. This model also introduces a role for government effort in determining voter utility. Voters care not only about what policies are implemented but how well they are implemented. This model assumes that voters know whether they are satisfied or dissatisfied with incumbent performance. Both sets of assumptions have been widely used in the formal literature; in the end, it is simply a behavioral question best addressed by the data.

3.2 Conditions for separating equilibrium

As is typically the case in this class of signaling games, multiple equilibria may be sustainable for various combinations of parameter values. Rather than characterizing all such equilibria, in this

section I establish the conditions under which a separating equilibrium exists and information is revealed by voter behavior. Electoral punishment occurs in this equilibrium: voters vote against the party of the federal incumbent in some states of the world even though they prefer the incumbent's party on policy grounds. For concreteness, I assume that the incumbent government is right of center and that the median voter prefers the right-of-center party to the left-of-center party, such that $||m - r|| < ||m - l||$.

Under this assumption, a separating equilibrium with electoral punishment consists of the strategies (vote right if good state, vote left if bad state) for the median voter and (exert effort if left, no effort if right) for the incumbent. This equilibrium implies electoral punishment that reveals information to politicians. If voters are angry with an incumbent government, they can send a credible message by voting against its co-partisans at a different level.

The conditions under which such an equilibrium exists can be easily identified. Since both of the incumbent's information sets are reached in equilibrium, there is no need to specify off-path beliefs and the necessary conditions can be identified by backwards induction. It is also not necessary to specify the probabilities of the two states, since in a separating equilibrium there is no uncertainty conditional on reaching either one of the incumbent's information sets. Identifying the necessary conditions for the incumbent is straightforward; the following two inequalities must be satisfied:

$$\alpha(-(m-r)^2+e^s)+(1-\alpha)(-(m-r)^2+e_n^f)-c \leq \alpha(-(m-r)^2+e^s)+(1-\alpha)(-(m-r)^2+e_g^f) \quad (3)$$

$$\alpha(-(m-l)^2+e^s)+(1-\alpha)(-(m-r)^2+e_n^f)-c \geq \alpha(-(m-l)^2+e^s)+(1-\alpha)(-(m-r)^2+e_b^f) \quad (4)$$

Equation 3 ensures that the incumbent does not exert effort when its co-partisan wins election an election at another level of government, while equation 4 ensures that the incumbent does exert effort when its co-partisan loses. The policy position of the incumbent government could be treated as a choice variable, but since the government has the last move and policy utility is not state contingent, the government will choose the same policy regardless of the outcome of the election. Likewise, actions by state-level politicians are taken as given. Since only the effort level is state-contingent, equations 3 and 4 simplify to the following set of inequalities:

$$e_n^f - e_g^f \leq c/(1-\alpha) \leq e_n^f - e_b^f \quad (5)$$

This simply states that the (scaled) costs of effort are such that the incumbent would prefer to exert effort in the bad state but not in the good state. If $c/(1-\alpha) \geq e_n^f - e_b^f$ the incumbent would never exert effort, while if $e_n^f - e_g^f \geq c/(1-\alpha)$ the benefits of effort exceed the cost in both states. Equation 5 defines the only strategically interesting case in which voters and incumbents have conflicting interests. As such, I assume that these inequalities are satisfied.

Turning to the representative voter, the following equations define the incentive constraints in each state of the world:

$$\alpha(-(m-l)^2 + e^s) + (1-\alpha)(-(m-r)^2 + e_n^f) \leq \alpha(-(m-r)^2 + e^s) + (1-\alpha)(-(m-r)^2 + e_g^f) \quad (6)$$

$$\alpha(-(m-l)^2 + e^s) + (1-\alpha)(-(m-r)^2 + e_n^f) \geq \alpha(-(m-r)^2 + e^s) + (1-\alpha)(-(m-r)^2 + e_b^f) \quad (7)$$

If equation 6 is satisfied, the voter supports the incumbent's co-partisan in the good state, while equation 7 implies that the voter votes against the incumbent party in the bad state. These constraints simplify to

$$e_n^f - e_g^f \leq \alpha/(1-\alpha)(-(m-r)^2 + (m-l)^2)) \leq e_n^f - e_b^f \quad (8)$$

Again, the interpretation of equation 8 is simple; for the representative voter to reveal information in equilibrium, the (scaled) policy cost of punishment must be less than the benefit of effort by the incumbent in the bad state of the world but greater than the benefit of effort in the good state.

While there is nothing surprising in the incentive constraints necessary for the existence of this separating equilibrium, these constraints enable characterization of the conditions under which informative signaling can arise. To make these conditions more transparent, it is useful to simplify equation 8 through some substitutions. Let $p = (-(m-r)^2 + (m-l)^2)$ be the differential policy utility for the median voter from supporting the right party. Since I assume the median voter is to the right of the cutpoint between the two parties, it follows that $p \geq 0$. Likewise, let $q = e_n^f - e_b^f$, the benefit of effort when the incumbent is performing poorly. Finally, let $k = e_n^f - e_g^f$, the benefit of effort in the good state. By assumption, $0 \leq k \leq q$. Given these definitions, equation 8 can be rewritten as

$$k \leq \alpha/(1-\alpha)p \leq q \quad (9)$$

Or, by rearranging,

$$\frac{k}{p+k} \leq \alpha \leq \frac{q}{p+q} \tag{10}$$

The region of the parameter space that satisfies these inequalities is shown in Figure 3. In this figure, the policy differential is plotted on the horizontal axis (in units of q, the benefit of effort in the bad state) and the importance of the election is plotted on the vertical axis. The shaded area represents the set $\{p, \alpha\}$ in which a separating equilibrium is sustainable when $k = 0.5q$. Note that the upper boundary of the shaded area is fixed (for fixed q), but the lower boundary moves toward or away from the upper boundary as k increases or decreases, respectively.

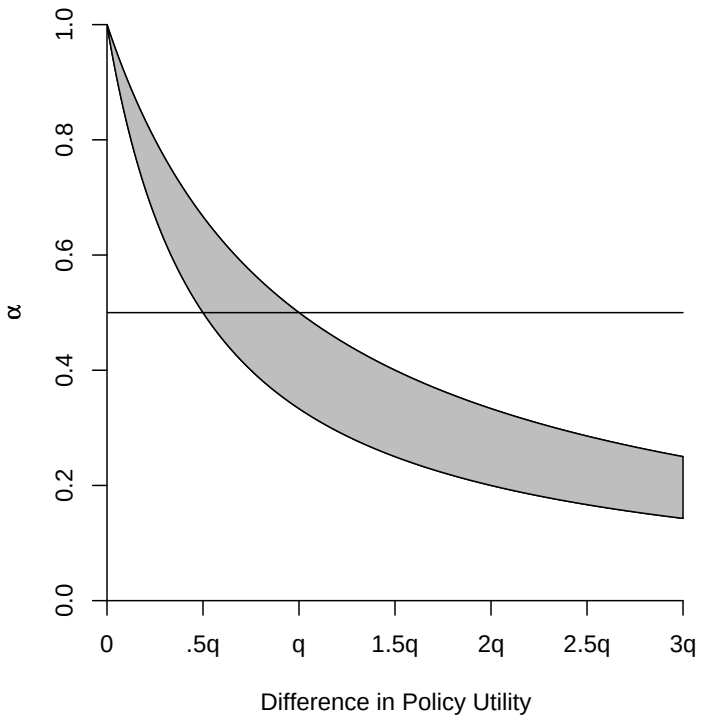

Figure 3: Conditions for electoral signaling. The grey area reflects the set of pairs of election importance (α) and differential policy utility (p) such that a fully informative separating equilibrium is sustainable. In the region above the grey area, signaling is too costly, while in the region below signaling is not credible.

3.3 Implications

Figure 3 illustrates the conditions under which voters can use elections to signal their displeasure with the performance of incumbents at other levels of government. The shaded area represents the region in which signals can be sent in a state election to the incumbent party at the federal level. In the region above the shaded area, sending a signal by voting against the incumbent party is too expensive in terms of the policy cost imposed. The upper boundary of the shaded region represents points at which the state-level policy costs are equal to the benefits from effort at the federal level, taking into account the relative importance of the two levels. The horizontal line at $\alpha = 0.5$ separates cases where the state election is more important than the federal election (above the line) from those where the state election is less important than the federal election (below the line).

In the region below the shaded area, signaling is too cheap to be credible. Voting against the incumbent party in a state election is a low-cost proposition in this region, because the state-level election is not very important and the policy differences between the parties are relatively small. In this region, voters in the good state have an incentive to mimic voters in the bad state by voting against the government to get it to expend effort; as a consequence, the government can no longer treat a vote against it as an unambiguous sign that it is worthwhile to exert effort.

This general division of the $\{p, \alpha\}$ space into three regions has several implications. First, for any fixed difference in policy utility p given the positions of the parties, there are some elections where the importance of the election makes signaling too expensive and other elections that are too insignificant to voters to support a signaling equilibrium. For any $p > 0$ there are non-empty intervals of α in each of the three regions. In particular, since the shaded region is bounded away from $\alpha = 0$, trivial elections can never be used to send informative signals to the incumbent party. While one may wonder whether α ever equals zero in a real election, public opinion polls can be viewed as elections with no direct consequences for voter utility. As a result, voters could not credibly signal dissatisfaction with the incumbent simply by responding to a survey. There is an incentive to misrepresent evaluations of the government, to exaggerate dissatisfaction, in order to spur the government to greater efforts.[5] Some cost is necessary to ensure credibility.

[5]Similar results regarding preference revelation in opinion polls are detailed in Meirowitz (2004); full information revelation does not occur in most cases.

Second, provided that the differential policy utility is greater than the benefit associated with incumbent effort when voters are dissatisfied, informative signaling can only occur in the less-important election. For $p > q$, the values of α for which the separating equilibrium is sustainable are always below 0.5. This implies that one would only expect to see punishment at the less important level of government under such conditions. Under the plausible assumption that the federal level is the most important, voters can signal their disapproval towards the federal incumbent but not toward state incumbents.[6]

Third, even when policy utility differences are small relative to the potential benefits of exerting additional effort, the region in which punishment can be sustained at both levels in equilibrium is quite small. For signaling to be feasible at both levels (assuming that parties have the same positions at the state and federal level), it must be the case that both α and $1 - \alpha$ are contained in the shaded area for a fixed p. This area can be found by reflecting the shaded area over the line $\alpha = 0.5$. The intersection between the shaded area and its reflection, as shown by the approximately diamond-shaped area in black in Figure 4, is the region where bilateral signaling is sustainable. The small size of this region suggests that it would be unusual to observe voters sending signals at both levels in the same political system.

Finally, and more speculatively, the policy differential needed to sustain a signaling equilibrium increases as elections become less important, holding the benefits of incumbent effort constant. Looking outside of the model to a certain extent, this might imply that voters have to vote for a more extreme option in relatively unimportant elections when a vote for the traditional opposition would suffice in a more important election. A Labour government at Westminster, for example, may be far more likely to react to victories by the far-right British National Party in local elections than they would to Conservative or Liberal Democratic wins in the same areas.

[6]Note that much of the shaded area lies above $\alpha = 0.5$ when policy differences are small, implying that signals could be sent in the more important election to affect the behavior of incumbents at the less important level of government. This result is driven by the fact that the representative voter doesn't much care who wins the election at the more important level when policy differences are small; faced with a choice between Tweedle-dee and Tweedle-dum, she is better off trying to get more effort out of the incumbent at the other level.

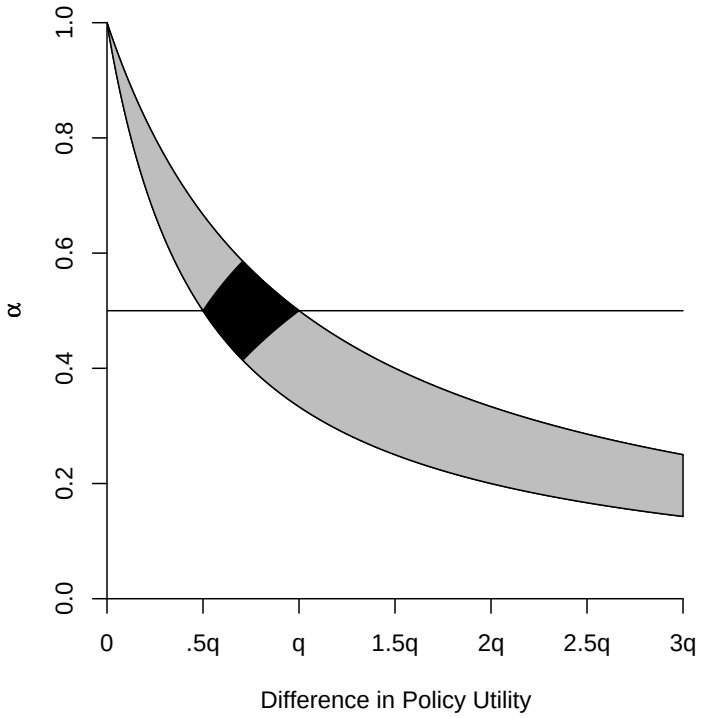

Figure 4: Conditions for signaling at both levels. The black area reflects the region in which voters could signal in either federal or state elections.

4 Building an empirical model

In order to evaluate these two explanations for electoral punishment, it is useful to construct an empirical model in which the predictions of both approaches can be embedded. As is standard in this literature, this paper uses aggregate vote results from both federal and state elections. A general empirical model of aggregate votes share for a particular party is

$$y_{i,t,l} = \nu_{i,t} + \beta_f x_f + \beta_s x_s + \beta_r x_r + \epsilon_{i,t,l} \tag{11}$$

where $y_{i,t,l}$ is the observed vote share for the party in state i at time t for level l (state or federal); $\nu_{i,t}$ is the baseline vote share for the party in state i at time t (commonly referred to as the normal vote); x_f is an indicator equal to 1 for a state election when the party is the federal incumbent; x_r is an indicator equal to 1 for a state election when the party is the federal opposition; x_s is an

indicator equal to 1 for a federal election when the party is the state incumbent. In this formulation, $\nu_{i,t}$ represents the normal vote in a federal election when the party is in opposition at the state level. Finally, $\epsilon_{i,t,l}$ is the realization of a symmetric, mean-zero disturbance.

Unfortunately, this general model is not particularly useful, not least because the normal vote $\nu_{i,t}$ is not observable. Researchers have used a variety of methods to address this problem. One approach used in several studies is to look at federal and state election pairs. The outcome is defined as the change in support for a party between the federal election and state election that immediately follows it, and regressing this on federal incumbency typically produces large estimates of electoral punishment at the state level. Unfortunately, if equation 11 is a reasonable approximation of the process generating observed vote shares, much of the alleged electoral punishment is simply a function of regression to the mean. Another approach looks at changes in vote shares between consecutive elections at the same level, but this approach will difference out examples of consecutive punishment, potentially reducing the estimated effect of federal incumbency.

As an alternative, it is possible to estimate $\nu_{i,t}$ from the observed data. In order to make the problem more tractable given the amount of data available, I assume that the normal vote can be decomposed into a smoothly evolving component due to time θ_t, and an additive mean shift γ_i for each state.

$$y_{i,t,l} = \theta_t + \gamma_i + \beta_f x_f + \beta_s x_s + \beta_r x_r + \epsilon_{i,t,l} \tag{12}$$

Allowing a different mean for each state accounts for long-run differences in the distribution of preferences from state to state. When combined with the common time-varying component θ_t, this is essentially equivalent to assuming a uniform national swing across units, an assumption often used in electoral studies. To the extent that the distribution of preferences is relatively flat over the range in which the cutpoints between parties are located, such an assumption is unlikely to produce major problems.

The primary challenge in operationalizing this model arises from the estimation of the unobserved θ_t. While this can be done in a number of ways, a Bayesian approach maximizes flexibility in dealing with latent quantities. I use a Bayesian dynamic linear model (DLM) (West and Harrison, 1997) to estimate the evolution of the normal vote over time. Bayesian DLMs have been used for a variety of applications in political science, including estimating the evolution of judicial ideal points (Martin and Quinn, 2002) and presidential approval (Beck, Jackman and Rosenthal, 2006). Let

$\boldsymbol{\alpha} = [\boldsymbol{\gamma} \;\; \beta_f \;\; \beta_s \;\; \beta_r]'$ and $\mathbf{x_i}$ be a row vector of observed indicators. The observed data is assumed to be distributed

$$y_{i,t,l} \sim N(\theta_t + \mathbf{x_i}\boldsymbol{\alpha}, \sigma^2)$$

where σ^2 is assumed constant across elections.[7] If the θ_t were observed, this would be a simple Bayesian linear regression with $(y_{i,t,l} - \theta_t)$ as the response, which could be closed out with standard semi-conjugate priors:

$$\alpha \sim N(\mathbf{q_0}, \mathbf{V_0})$$

$$\sigma \sim InverseGamma(c_0, d_0)$$

The dynamic portion of the DLM arises from the model for the evolution of the normal vote θ_t. I assume a simple "local-level" random walk structure for the evolution of θ_t, making the following assumption:

$$\theta_t \sim N(\theta_{t-1}, s_0^2)$$

where s_0^2 represents the variability of the random walk, or equivalently, how quickly the normal vote is believed to change over time. This framework requires the specification of a prior for θ_0, assumed to be normal with mean m_0 and variance C_0.

This DLM as described above is quite similar to the model for presidential approval in Beck, Jackman and Rosenthal (2006), with election results taking the place of opinion polls and state-specific offsets replacing house effects.[8] As in that model, it is not feasible to estimate a mean shift γ_i for each of the subnational units in the data due to a lack of identification. The simplest approach to deal with this problem is to set one of the units as the baseline, so that the remaining γ_i estimate shifts relative to that unit. The results from this model provide estimates of the average shift in vote share for a party at one level of government when the same party is in office at the other level.[9]

[7] In particular, the observations are not weighted by the population of the various units. Given the number of votes cast in even the smallest states, uncertainty due to the number of votes cast is far outweighed by other sources of fundamental randomness

[8] The principal difference is that Beck, Jackman and Rosenthal (2006) allows for two different transition regimes while I assume the transition regime is stable over time.

[9] The results of this model are not strongly dependent on the Bayesian DLM structure. A similar

The statistical model described above provides a principled framework for estimating the average levels of electoral punishment experienced by parties. The formal model in section 3, on the other hand, implies that estimating a single effect for electoral punishment averages over two types of elections. In some elections, voters signal their dissatisfaction with the incumbent government at one level of government by voting against it at another level. In other elections, they vote as they normally would, without taking the incumbent party at the other level into account. The conditions under which these outcomes occur cannot be observed. Indeed, the logic of the model depends in part on the inability to condition on the state variable; presumably, if an analyst could observe the state *ex post* then politicians could observe it as well, rendering the signal superfluous. It is relatively simple in the DLM framework to estimate these unobserved states from the data. An observed value of 1 in the indicator variable for federal party incumbency during a state election can be replaced with an unobserved $x_{i,k}$ equal to 1 if the election is in the dissatisfied state and 0 otherwise. In other words, the mean is shifted by α_k when $x_{i,k} = 1$ while the mean is unchanged for $x_{i,k} = 0$.

Implementing this expanded model requires specifying a distribution for the $x_{i,k}$:

$$x_{i,k} \sim Bernoulli(\pi_k)$$

and a prior distribution for the probability π_k of being in the dissatisfied state:

$$\pi_k \sim Beta(a_k, b_k)$$

model can be fit using a generalized additive model (or GAM) (Hastie and Tibshirani, 1990; Beck and Jackman, 1998), which uses a non-parametric approach to estimate the mean support for a party as a smooth function of time while allowing fixed effects for each state and indicators for incumbency at the state and federal level to be entered parametrically. Since these are all indicator variables, the only structure imposed by this parametric assumption is that there is no interaction between these variables and the smoothed normal vote function. The estimates from the DLM and the GAM models are essentially identical for the quantities of interest. The motivation for choosing the DLM framework is the ease of estimating a mixture of punishment and non-punishment outcomes, as implied by the signaling model presented in section 3.

This formulation of the model can be interpreted as setting up a mixture of normal distributions with means separated by α_k and a common variance.

Estimation of these models is accomplished through Markov Chain Monte Carlo simulations to generate samples from the posterior distribution.[10] For each estimation, four chains of 100,000 iterations were run, saving every 100th draw for posterior inference. Point estimates were generated by estimating the posterior mean using samples from the posterior, and central credible intervals were constructed from the quantiles of the sample generated. Details of the estimation procedure are presented in the Appendix.

5 Data and Results

I compare the empirical implications of these two theories of electoral punishment using data from German elections to the Bundestag and the 16 Land governments. Germany is a useful starting point since evidence of electoral punishment in state elections has been found by many authors using a variety of statistical models (Lohmann, Brady and Rivers, 1997; Gaines and Crombez, 2004; Kedar, 2006; Hainmueller and Kern, 2006). In addition, state governments in Germany are represented at the federal level through the Bundesrat, the upper legislative chamber. Since the Bundesrat has some ability to influence legislation, one might expect balancing incentives to be stronger in German state elections than in federal systems where the constituent governments have no direct representation such as Canada or Australia.

Political competition in the Federal Republic of Germany has been dominated by two large parties, the Social Democratic Party (SPD) and the Christian Democratic Union. These parties contest elections at the federal level and in every state.[11] These two parties combined have averaged approximately 80% of the vote in federal elections. All post-war German federal governments have been led by chancellors from one of these two major parties. Likewise, nearly all governments at the state level have been led by politicians from these two parties. Information on the composition of state and federal incumbent governments at time of each election to the federal and state

[10] See Gelman et al. (2004) for an introduction to Bayesian simulation methods.

[11] For the purposes of this paper, the Christian Social Union in Bavaria is being treated as a part of the CDU, with which it has been in almost perpetual alliance. Barring exceptional circumstances, the CDU does not run candidates in Bavaria and the CSU does not compete outside of Bavaria.

legislatures, respectively, was collected from the websites of the state and federal legislatures.

In Germany, state and federal elections do not generally occur on the same day. State elections are staggered throughout the term of the federal legislature. Less than five percent of state elections since 1946 have coincided with federal elections. As a result, German voters typically know which party is in office at one level of government when casting votes for the other level. The results of German elections were obtained from Neu (2004); this source includes all Bundestag and Land elections from 1946 to February 2005. It reports the calendar day of each election and the vote share for each of the major parties, for a total of 359 observations. There are approximately 15 elections each to the Bundestag and to state governments in the states of the former West Germany and four elections each in the eastern states.

5.1 Average effects of cross-level incumbency

Since the realized vote shares for parties competing in a single election are clearly not independent, I estimate the model described in equation 12 using data from one party, focusing primarily on the vote share of the CDU. In particular, the CDU vote share is modeled as a smoothly evolving function of time (measured in years), with fixed effects for each state. The variables of interest are indicators for the partisan composition of the party in power at the other level of government. I define the baseline to be federal elections when the SPD is in power at the state level, and include indicators for CDU government at the state level during federal elections, SPD government at the federal level during a state election, and CDU government at the federal level during a state election.[12] Note that this model does not include variables to identify the incumbent government at each level going into elections at that level; since neither the balancing or signaling approaches make predictions about the effects of incumbency after conditioning on the party composition at the other level, any such effects are left to the error term.

The results of this model are best presented graphically. While not of primary interest, estimates of the normal vote produced by the DLM are useful in building intuition about the model. Figure 5 shows a scatterplot of election results over time; Bundestag elections are plotted as circles and state elections as triangles. The results of elections from Schleswig-Holstein, a relatively small state

[12]Including an interaction term for elections held when a grand coalition of the CDU and SPD was in power at the opposite level does not affect the results materially.

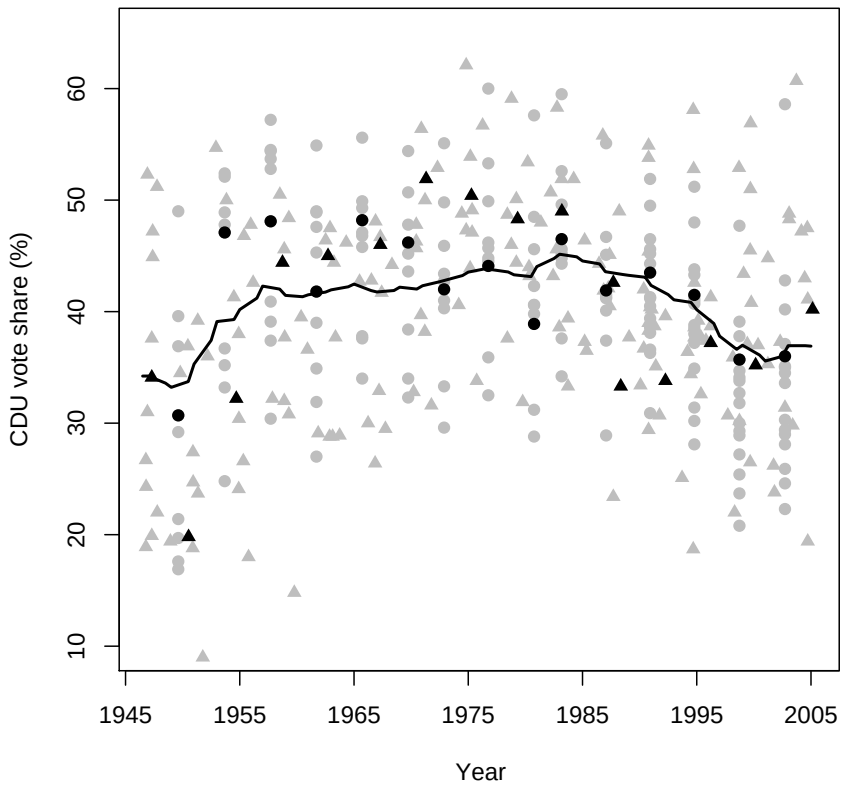

Figure 5: Vote share for the CDU. Elections from Schleswig-Holstein plotted in black, all others in grey. Bundestag elections are circles, state elections are triangles. The smoothed line represents the model prediction for CDU vote share in Schleswig-Holstein in federal elections when the CDU is not in state office.

bordering Denmark, are highlighted. This figure illustrates the significant heterogeneity in election results across the German states. The line in Figure 5 represents the estimated posterior mean for the CDU vote share in Schleswig-Holstein in federal elections when the SPD is in office at the state level. This fitted line implies that support for the CDU increased in the early years of the Federal Republic as the party system consolidated, remained fairly stable for about 35 years and then fell off during the 1990s. The posterior means for other states shift up and down by constants estimated in the model.

Turning to the estimates of the quantities of interest, Figure 6 presents a coefficient plot with central credible interval estimates for the various possible combinations of election and incumbency. In this Bayesian framework, these estimates imply that the posterior probability of the parameter falling in the indicated interval is 95%, given the model and the observed data. For each quantity,

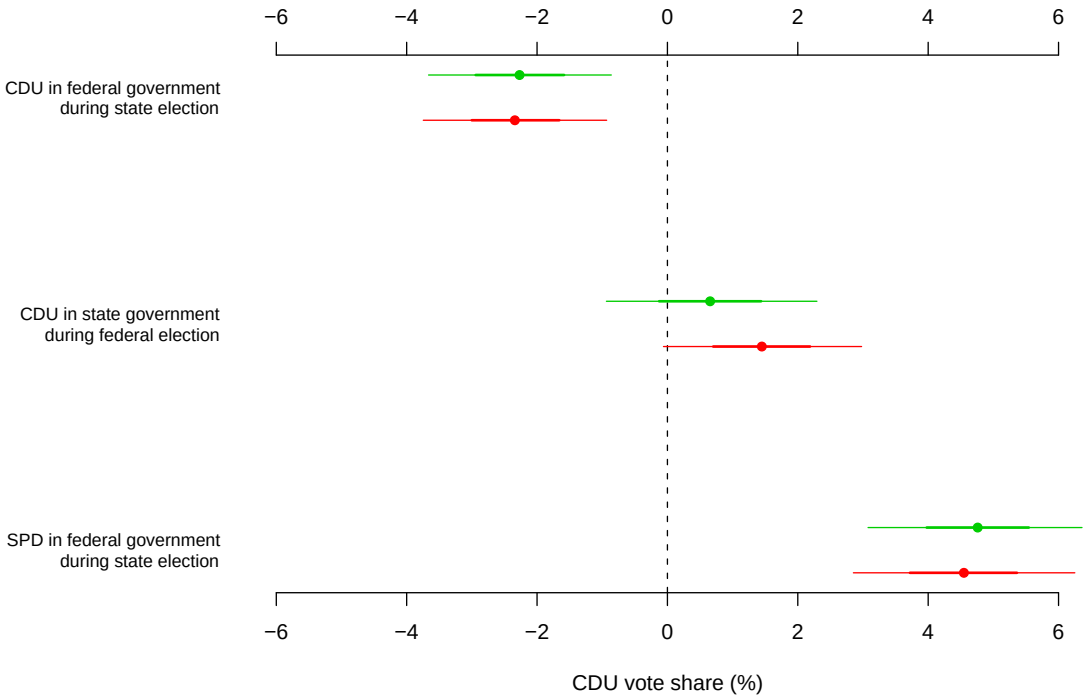

Figure 6: Effects of incumbency on CDU vote share in elections to other level of government. Interval estimates represent 66% (thick line) and 95% (thin line) central credible intervals. The green (upper) line estimated from full sample (N=359); the red (lower) line excludes eastern states (N=298).

the upper interval (green) is estimated from the entire dataset while the lower interval (red) is estimated without elections from the former eastern states.

The pair of intervals in the upper left corner represent the estimated difference in vote share for the CDU in state elections when the CDU is in office at the federal level. Both estimates suggest that the CDU does about two and a half percentage points worse in state elections when it is the federal incumbent, relative to the normal vote in federal elections. Likewise, the intervals in the lower right corner indicated that the CDU does a little over four percentage points better in state elections when the SPD is in office in Berlin.[13] The 95% interval estimates for both of these quantities are bounded well away from zero. These estimates provide strong evidence, consistent

[13]Note that the quantity usually estimated in this context is the difference between these two coefficients, since with the exception of the short grand coalition in the late 1960s, when the CDU is in office at the federal level the SPD is not.

with earlier analyses of the same data, that federal incumbents are punished at the state level on average, as would be expected under either the balancing or signaling models.

The estimates in the middle of Figure 6, on the other hand, may speak to the differences between the balancing and signaling approaches. In both the full dataset and the dataset restricted to the former West Germany, the estimated difference between the CDU's performance in federal elections when it is in office at the state level and its performance when the SPD holds state office is positive, not negative. These estimates are also quite close to zero. There is nothing in the data to suggest that voters attempt to moderate the positions advocated by their state governments by supporting the state opposition at the federal level. This is consistent with the signaling model presented here, which implies that under most plausible conditions voters would not use federal elections to signal their displeasure with the performance of incumbents at the other level. It is less consistent with some versions of balancing models, which imply that some voters have an incentive to vote against the party in power in their state, as discussed in section 2.

As a robustness check, Figure 7 presents the equivalent estimates from the model fit to SPD vote shares. The coefficients for state elections are essentially a mirror image of those obtained from the CDU data; the SPD does about four percentage points worse when it is in office at the federal level and about three percentage points better when the CDU is in office at the federal level.[14]. The estimated effect for SPD incumbency at the federal level is slightly more negative when data from the eastern states is included; this appears to be the result of voters in some of those states supporting smaller parties, in particular the (post-Communist) Party of Democratic Socialism, at the state level during the Schroeder government; this does not appear in the CDU vote shares. Turning to the coefficient on state incumbency, the estimated relationship is again positive rather than negative and close to zero.

5.2 Effects allowing for signaling

The results presented above demonstrate that state parties lose support on average when their co-partisans are in office at the federal level. The signaling model, however, implies that state parties should only lose support in certain states of the world; in particular, when voters are dissatisfied

[14]This highlights the need to fit the model to the results from one party; the same information is encoded in both subsets of the data

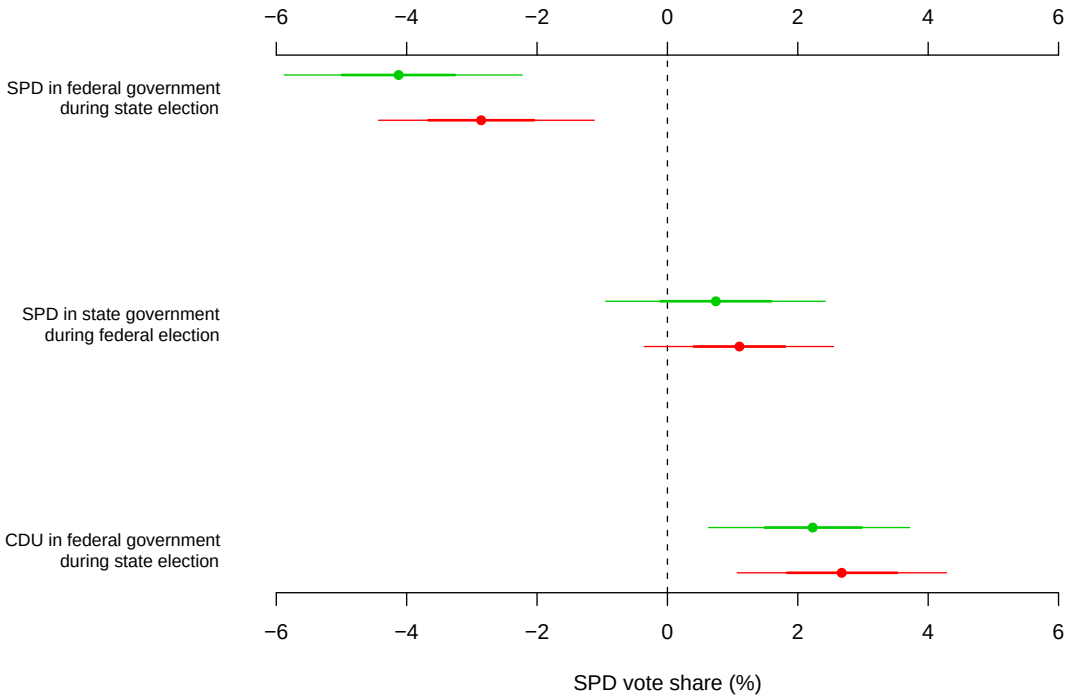

Figure 7: Effects of incumbency on SPD vote share in elections to other level of government. Interval estimates represent 66% (thick line) and 95% (thin line) central credible intervals. The green (upper) line estimated from full sample (N=359); the red (lower) line excludes eastern states (N=298).

with the incumbent federal government and need a credible way to communicate their discontent. It is not possible to condition directly on these unobserved state variables, but the Bayesian framework enables them to be estimated.

The posterior estimates allowing for unobserved states are presented in Figure 8. Not surprisingly, the magnitude of the electoral punishment inflicted on federal incumbent parties in state-level elections is much larger in the subset of elections where the model estimates that such punishment occurred. Instead of the roughly 2.5 percentage point shift in vote share for state CDU parties when the CDU is in office at the federal level, the extended model now estimates losses in excess of ten percentage points relative to the normal vote. Likewise, the CDU appears to benefit by about nine points when the SPD is in government at the federal level. Again, there is little evidence that the CDU suffers in federal elections when in office at the state level; while the posterior mean for these elections is now negative, it is even closer to zero.

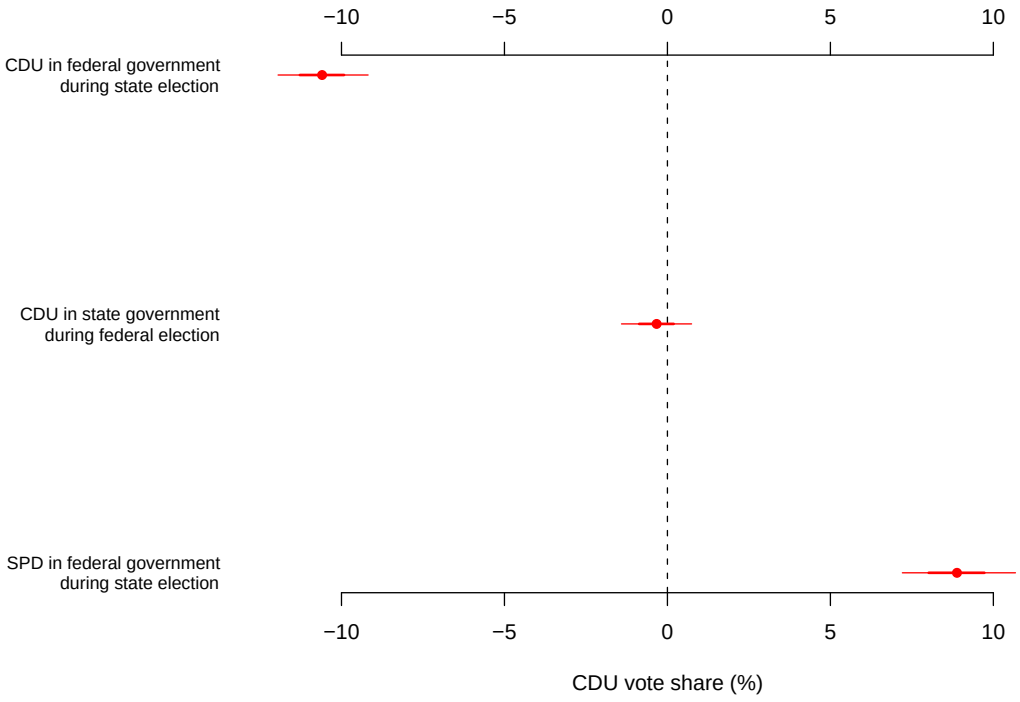

Figure 8: Effects of incumbency on CDU vote share in elections to other level of government, allowing for mix of punishment and non-punishment elections. Interval estimates represent 66% (thick line) and 95% (thin line) central credible intervals. Estimates from the full sample (N=359)

The Bayesian mixture model also produces estimates of the proportion of state elections in which punishment occurs. The posterior probability of punishment when the CDU is in office at the federal level is 0.366, with a 95% central credible interval of [0.28, 0.46]. The posterior probability that the state CDU benefits when the SPD is in office at the federal level is 0.45, with a 95% central credible interval of [0.30, 0.60]. While the SPD may be punished more often than the CDU, the greater uncertainly in the posterior probability implies that the separation of elections into punishment and non-punishment categories is not as clean when the SPD participates in the federal government.

This is confirmed by considering the posterior probability that each state election falls into the punishment category. The distributions of these posterior probabilities are shown in Figure 9, with elections under CDU federal governments on the left and under SPD governments on the right. Almost all state elections with CDU governments fall cleanly into one of the two categories, with few observations falling between the modes at 0 and 1. The separation is not quite as distinct

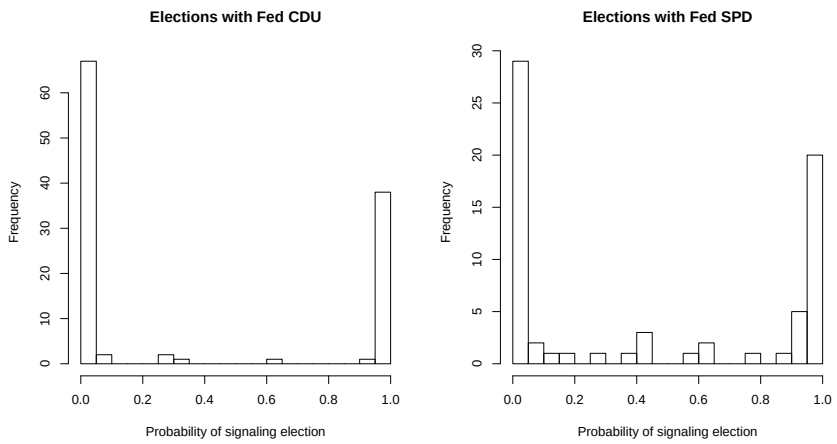

Figure 9: Histograms of posterior probabilities that punishment occurred in state elections under CDU and SPD governments at the federal level. Elections near 1 have a high posterior probability that punishment took place. There were 112 state elections while the CDU was in office at the federal level and 68 state elections while the SPD was in office at the federal level.

for elections under the SPD. Nevertheless, over 2/3 of the observations fall at the extremes of the distribution.

The results presented here are consistent with a signaling explanation for the losses that German parties suffer in state elections when they are in office at the federal level. The underlying logic of signaling-based approaches, however, depends crucially on the ability of incumbent parties to correctly identify the message being sent by voters. While the simple model presented in section 3 assumes that the voter's action is perfectly observable by the government, this is clearly a simplification that ignores the fundamental uncertainty associated with electoral outcomes. In light of this, it is worth considering whether it is plausible that politicians could decipher a signal given the estimates of the model. The relevant comparison here is between the estimated size of the effect and the residual uncertainty unexplained by the model. In this case, moving from the simple average shift model to a mixture model reduces the residual standard deviation from approximately 5.15 to 3.75, as would be expected given the increased flexibility of the model. With an estimated shift of roughly ten percentage points from the normal vote, the average punishment estimated by the model is almost three standard deviations away from the normal vote. This suggests that punishment, when it occurs, is severe enough to be quite informative to politicians viewing the election results.

6 Discussion

This paper has presented a simple model of signaling in subnational elections. In contrast to most balancing theories, it implies that under plausible conditions federal incumbent parties may be punished in state elections but state incumbent parties should not lose support in federal elections, provided that outcomes at the federal level are more important to voters. This model formalizes a longstanding argument in the literature on subnational elections, and evidence from German elections is consistent with its predictions. The signaling model suggests voters comment on federal politics in subnational elections because the cost of voting against their preferred party in a subnational election is precisely what makes their message credible to the incumbent government.

The empirical analysis presented in this paper focuses on the implications of the signaling model for aggregate election outcomes. The two main implications, that electoral punishment should only occur at one level and should only occur in some elections, both found support in the data. Another key empirical implication of signaling models remains untested: if voters are using state elections to send messages to federal governing parties, then one would expect the incumbent government to change its behavior in response to the information revealed by the signal. A systematic analysis of federal government response to electoral losses at the state level awaits future research.

Appendix

Estimating the Bayesian dynamic linear model described in section 4 requires that samples be drawn from the joint posterior distribution of the parameters $[\boldsymbol{\theta}, \boldsymbol{\alpha}, \boldsymbol{\pi}, \mathbf{X}, \sigma^2]$. This is most easily accomplished using Markov Chain Monte Carlo (MCMC) methods; the Markov Chain is constructed so that its stationary distribution is the posterior distribution of interest. As a result, running the chain for a sufficient number of iterations produces samples approximately from the joint posterior of the parameters.

The MCMC algorithm used to estimate the model presented in this paper is a Gibbs sampler, which iteratively produces draws from the full conditional distribution of a parameter or block of parameters given the current values of the other parameters. This approach is attractive because all of the parameters in this model have full conditionals from known parametric families, making it easy to draw the necessary samples. Given starting values for $[\boldsymbol{\alpha}, \boldsymbol{\pi}, \mathbf{X}, \sigma^2]$, the algorithm proceeds as follows:

1. Simulate $\boldsymbol{\theta}$ from $f(\boldsymbol{\theta}|\mathbf{y}, \mathbf{X}, \boldsymbol{\alpha}, \sigma^2)$. Simulating from this full conditional is the most difficult. While it is possible to sample from each of the θ_i separately, these parameters have high posterior correlations and so this implementation is inefficient. Instead, a forward-filtering backward-sampling algorithm is used, taking advantage of the fact that the joint full conditional $f(\boldsymbol{\theta}|\mathbf{y}, \mathbf{X}, \boldsymbol{\alpha}, \sigma^2)$ factors in a particularly convenient way for this class of models (West and Harrison, 1997; Martin and Quinn, 2002). The algorithm proceeds through these steps (following Martin and Quinn (2002)):

   - In the forward-filtering section, for each t in 1 to T:
     (a) Calculate the prior variance $R_t = C_{t-1} + s_0^2$.
     (b) Calculate the mean of the one-step-ahead forecast $\mathbf{f_t} = \mathbf{z_t} m_{t-1}$.
     (c) Calculate the variance of the one-step-ahead forecast $\mathbf{Q_t} = \mathbf{z_t} R_t \mathbf{z_t'} + \sigma^2 \mathbf{I}_{n_t}$.
     (d) Calculate the Kalman gain $\mathbf{A_t} = R_t \mathbf{z_t'} \mathbf{Q_t}^{-1}$.
     (e) Calculate the posterior mean $m_t = m_{t-1} + \mathbf{A_t}(\mathbf{y_t} - \mathbf{X_t}\boldsymbol{\alpha} - \mathbf{f_t})$.
     (f) Calculate the posterior variance $C_t = R_t - \mathbf{A_t Q_t A_t'}$
   - In the backward-sampling section, begin by sampling from $\theta_T \sim N(m_T, C_T)$. Then, for each t in $T-1$ to 1:

(a) Calculate $B_t = C_t + R_t^{-1}$.

(b) Calculate $h_t = m_t + B_t(\theta_{t+1} - m_{t+1})$.

(c) Calculate $H_t = C_t - B_t R_t B_t'$.

(d) Sample from $\theta_t \sim N(h_t, H_t)$.

2. Simulate $\boldsymbol{\alpha}$ from $f(\boldsymbol{\alpha}|\mathbf{y}, \mathbf{X}, \boldsymbol{\theta}, \sigma^2) \sim N(\mu_\alpha, \boldsymbol{\Sigma}_\alpha)$ where:

- $\boldsymbol{\Sigma}_\alpha = \left[\sigma^{-2}\mathbf{X}'\mathbf{X} + \mathbf{V_0}^{-1}\right]^{-1}$
- $\mu_\alpha = \boldsymbol{\Sigma}_\alpha \left[\sigma^{-2}\mathbf{X}'(\mathbf{y} - \mathbf{Z}\boldsymbol{\theta}) + \mathbf{V_0}^{-1}\mathbf{g_0}\right]$

3. Simulate σ^2 from $f(\sigma^2|\mathbf{y}, \mathbf{X}, \boldsymbol{\theta}, \boldsymbol{\alpha}) \sim InvGamma((c_0 + N)/2, (d_0 + (\mathbf{y} - \mathbf{X}\boldsymbol{\alpha} - \mathbf{Z}\boldsymbol{\theta})'(\mathbf{y} - \mathbf{X}\boldsymbol{\alpha} - \mathbf{Z}\boldsymbol{\theta}))/2)$

When some of the indicators $x_{i,k}$ are unobserved, the model is extended as follows:

4. Simulate $x_{i,k}$ from $f(x_{i,k}|y_i, \mathbf{x_{i,-k}}, \boldsymbol{\theta}, \boldsymbol{\alpha}, \pi_k) \sim Bernoulli(\rho_{i,k})$ where:

- $\rho_{i,k} = \frac{\pi_k \exp\left\{-\frac{1}{2}(\alpha_k^2 - 2\alpha_k(y_i - \mathbf{z}_i\boldsymbol{\theta} - \mathbf{x}_{i,-k}\boldsymbol{\alpha}_{-k}))\right\}}{\pi_k \exp\left\{-\frac{1}{2}(\alpha_k^2 - 2\alpha_k(y_i - \mathbf{z}_i\boldsymbol{\theta} - \mathbf{x}_{i,-k}\boldsymbol{\alpha}_{-k}))\right\} + (1 - \pi_k)}$

5. Simulate π_k from $f(\pi_k|\mathbf{x_k}) \sim Beta(\sum x_{i,k} + a_0, n_k - \sum x_{i,k} + b_0)$.

The priors used in each model were as follows. For the prior on $\boldsymbol{\alpha}$, a normal distribution with prior mean $\mathbf{g_0} = \mathbf{0}$ and $\mathbf{V_0} = 1000\mathbf{I}_k$. For the prior on σ^2, an inverse gamma distribution with parameters $c_0 = 0.1, d_0 = 0.1$. For the evolution variance of the random walk, $s_0^2 = 0.04$, corresponding to a prior standard deviation of 0.2. This implies an expectation that most of the changes in the normal vote from year to year to be within plus or minus 0.4 percentage points. Finally, for the priors on the π_k, beta distributions with $a_k = b_k = 1$ such that the prior is uniform on the interval (0,1).

This model was estimated using the R statistical computing environment.[15] Each model was estimated by running four chains in parallel for 100,000 iterations and retaining every 100th draw. The retained draws from each chain were then combined and posterior inferences were made from those 4000 samples. Standard diagnostic tests (both single chain and multi-chain) were consistent with convergence.

[15]Replication code available from the author on request.

References

Alesina, Alberto and Howard Rosenthal. 1996. "A theory of divided government." *Econometrica* 64:1311–1341.

Ansolabehere, Stephen and James M. Snyder. 2000. "Valence politics and equilibrium in spatial election models." *Public Choice* 103:327–336.

Banks, Jeffrey. 1991. *Signaling Games in Political Science*. London: Routledge.

Beck, Nathaniel and Simon Jackman. 1998. "Beyond linearity by default: Generalized additive models." *American Journal of Political Science* 42:596–627.

Beck, Nathaniel, Simon Jackman and Howard Rosenthal. 2006. "Presidential approval: The case of George W. Bush." Presented at the 2006 Summer Methods Meeting, Davis, California.

Callander, Steven. 2007. "Bandwagons and momentum in sequential voting." *Review of Economic Studies* 74:653–684.

Cho, In-Koo and David Kreps. 1987. "Signaling games and stable equilibria." *Quarterly Journal of Economics* 102:179–222.

Chubb, John E. 1988. "Institutions, the economy, and the dynamics of state elections." *American Political Science Review* 82:133–154.

Erikson, Robert S. and Mikhail G. Filippov. 2001. "Electoral balancing in federal and sub-national elections: The case of Canada." *Constitutional Political Economy* 12:313–331.

Fowler, James H. and Oleg Smirnov. 2007. *Mandates, Parties, and Voters: How Elections Shape the Future*. Philadelphia: Temple University Press.

Gaines, Brian J. and Christophe Crombez. 2004. "Another look at connections across German elections." *Journal of Theoretical Politics* 16(3):289–319.

Gelineau, Francois and Eric Belanger. 2005. "Electoral accountability in a federal system: National and provincial economic voting in Canada." *Publius* 35:407–424.

Gelman, Andrew, John B. Carlin, Harold S. Stern and Donald B. Rubin. 2004. *Bayesian Data Analysis*. 2nd ed. Boca Raton: Chapman & Hall/CRC Press.

Groseclose, Tim. 2001. "A model of candidate location when one candidate has a valence advantage." *American Journal of Political Science* 45(4):862–886.

Hainmueller, Jens and Holger L. Kern. 2006. "Electoral balancing, divided government, and midterm loss in German state elections." *Journal of Legislative Studies* 12:127–149.

Hastie, Trevor and Robert Tibshirani. 1990. *Generalized Additive Models*. Boca Raton: Chapman and Hall/CRC Press.

Kedar, Orit. 2006. "How voters work around institutions: Policy balancing in staggered elections." *Electoral Studies* 25:509–527.

Lohmann, Susanne, David Brady and Douglas Rivers. 1997. "Party identification, retrospective voting, and moderating elections in a federal system: West Germany, 1961–1989." *Comparative Political Studies* 30:759–774.

Lohmann, Suzanne. 1994. "Information aggregation through costly political action." *American Economic Review* 84:518–530.

Marsh, Michael. 1998. "Testing the second-order election model after four European elections." *British Journal of Political Science* 28:591–607.

Martin, Andrew D. and Kevin M. Quinn. 2002. "Dynamic ideal point estimation via Markov Chain Monte Carlo for the U.S. Supreme Court, 1953-1999." *Political Analysis* 10(2):134–153.

Mebane, Walter and Jasjeet Sekhon. 2002. "Coordination and policy moderation at midterm." *American Political Science Review* 96:141–157.

Meirowitz, Adam. 2004. "Polling games and information revelation in the Downsian framework." *Games and Economic Behavior* 51:464–489.

Meirowitz, Adam and Joshua A. Tucker. 2007. "Run Boris Run: Strategic voting in sequential elections." *Journal of Politics* 69:88–99.

Nagel, Klaus-Jürgen. 2006. "North Rhine Westphalia: the Land election That dismissed a federal government." *Regional and Federal Studies* 16:347–354.

Neu, Viola. 2004. Wahllexikon 2004: Stand 23.02.2005 [Election encyclopedia 2004]. Working paper Konrad-Adenauer-Siftung Berlin: .

Piketty, Thomas. 1999. "The information-aggregation approach to political institutions." *European Economic Review* 43:791–800.

Piketty, Thomas. 2000. "Voting as communicating." *Review of Economic Studies* 67:169–191.

Razin, Ronny. 2003. "Signaling and election motivations in a voting model with common values and responsive candidates." *Econometrica* 71:1083–1119.

Reif, Karlheinz and Hermann Schmitt. 1980. "Nine second-order national elections – a conceptual framework for the analysis fo European election results." *European Journal of Political Research* 8:3–44.

Rodden, Jonathan and Erik Wibbels. 2005. "Retrospective voting, coattails, and accountability in regional elections." Presented at the 2005 Annual Meeting of the American Poltical Science Association, Washington, D.C.

Shotts, Kenneth W. 2006. "A signaling model of repeated elections." *Social Choice and Welfare* 27:251–261.

Shugart, Matthew S. 1996. "The electoral cycle and institutional sources of divided presidential government." *American Political Science Review* 89:1–17.

Stokes, Donald. 1963. "Spatial models of party competition." *American Political Science Review* 57:368–377.

van der Eijk, Cees, Mark Franklin and Michael Marsh. 1996. "What voters teach us about Europe-wide elections: What Europe-wide elections teach us about voters." *Electoral Studies* 15:149–166.

West, Michael and Jeffrey Harrison. 1997. *Bayesian forecasting and dynamic models*. New York: Springer.