Accountability with Multidimensional Policy Experimentation

Jesse M. Crosson * Giovanna M. Invernizzi † Federica Izzo ‡

September 8, 2022

Abstract

Politics in general and policymaking in particular are inherently multidimensional. We develop a model of accountability to study policymaking when multiple policy dimensions are available. In the model, an incumbent chooses whether and how to act on each of two correlated policy dimensions. While voters do not know what the optimal policies are, they can infer it by observing the incumbent's choices and the resulting outcomes. Thus, the incumbent influences voter learning by either focusing on a single issue or acting on multiple policy dimensions. We characterize the officeholder's decision to expand (or contract) the scope of policymaking, based on his ex-ante electoral strength and on the correlation across dimensions. Results also show how the possibility to act on multiple correlated dimensions influences policymakers' incentives to pursue moderate or extreme policies.

*Purdue University, *Email*: `jcrosson@purdue.edu`

†Collegio Carlo Alberto, *Email*: `giovanna.invernizzi@carloalberto.org`

‡UCSD, *Email*: `fizzo@ucsd.edu`

1 Introduction

Contemporary politics within a large number of Western democracies, most notably the United States, are frequently described as polarized along a dominant issue or ideological dimension. Often depicted as an economic or "scope of government" left-right spectrum, political scientists in particular have argued for several decades that a single preference dimension explains a substantial amount of variation in elite political behavior. Nevertheless, political pundits, practitioners, and even scholars acknowledge that politics and policymaking is inherently multidimensional, especially in polities with varied economies, racial diversity, and religious pluralism. More than just an important descriptive feature of politics, classic accounts of power and influence frequently contend that the decision over which of these dimensions to pursue—and which to avoid (Bachrach and Baratz, 1963)—is perhaps the most consequential factor of all in determining who wins and loses in politics. Indeed, as Schattschneider famously summarizes in his classic analysis of American democracy: "Whoever decides what the game is about decides also who can get into the game"

Multidimensionality has long played a role in models of bargaining and coalition-building within legislatures (e.g., Banks and Duggan 2000; 2006). However, as Bawn et al. (2012) and others have shown, the decision of whether and which issues and issue positions should "go together" is a strategic consideration of tremendous political consequence. Indeed, the expansion of a political appeal to include more than one issue area can dramatically influence the organization of a political party, the character of electoral campaigns, the dynamics of legislative bargaining, and the extent of policy output in a legislature. As a result, it stands to reason that political leaders must think carefully about which issues they will pursue.

More than deciding which issues—or how many—to pursue, political leaders must also contend with the fact that issues are both inherently and symbolically connected to one another. As Martin Luther King, Jr. eloquently argued in *Letter from a Birmingham Jail* (1963), "Injustice anywhere is a threat to justice everywhere. We are caught in an inescapable network of mutuality, tied in a single garment of destiny. Whatever affects one directly, affects all indirectly." More than solely a moral or ethical call to action, Dr. King's observations underscore a sincere belief that action or inaction in one policy area affects policy and outcomes in others. What is best for a citizen in one facet of public life and policy is anything but disconnected from other facets. Rather, if injustice obtains in one area, it is likely to redound to others, and citizens ought to adjust their behavior accordingly.

In this paper, we present a formal model to study policymaking in such a multidimensional world, within an electoral accountability framework. We investigate how the introduction of multiple, correlated issue dimensions into candidate-voter interactions influences how politicians approach policymaking. Generally, the inclusion of multiple issue spaces in models of policymaking—particularly those relating to legislative bargaining—creates opportunities for compromise or moderation (e.g., logrolling in appropriations negotiations). However, as we will show, the opportunity to expand policymaking and messaging to additional issue dimensions has, above all, a *distorting* effect on policymaking, under some conditions encouraging politicians to pursue more *extreme* policies than they otherwise might pursue. In the process, then, legislation and messaging along more than one dimension may worsen policy representation according to our model, rather than improving it.

Thus, our contribution is twofold. First, we study officeholders' decisions to expand (or contract) the scope of policymaking. Second, we characterize how the possibility to act on multiple correlated dimensions influences policymakers' incentives to pursue moderate or extreme programs.

1.1 Voter Learning, Multidimensionality and Accountability

Two features lie at the core of our theory. First, the voter faces uncertainty about her optimal policy. Second, when orienting herself in a multidimensional world the voter may obtain relevant information on one dimension by observing the content and outcome of policymaking in another.

Problems of voter information are a central challenge to popular rule. As a large body of both theoretical and empirical research has shown, voters face incentives to remain "rationally ignorant" of information useful in rendering voting decisions. As Downs (1957) puts it, gathering and processing information about candidates and issues makes little sense for any individual voter, particularly given the exceedingly low probably that she will be pivotal in an election.

Nevertheless, millions of voters cast ballots each election, and a large body of scholarship has argued that voters link their voting decisions not simply to policy information, but to their personal well-being. Although scholars disagree about the relative weight that voters place on various vote-influencing factors, they generally maintain that voters choose their optimal strategy based not only on the incumbent's actions, but also on how such actions impact observed outcomes. Indeed, some empirical research has shown that voters can and do react to the results of policy choices—not simply the substance of the policies themselves (e.g., Fiorina 1978, Alt, Bueno de Mesquita and Rose 2011).

While these answers to the problem of rational ignorance may help to explain how voter behavior can operate in modern democracy, they often do not contend with the reality that politics and policymaking—and outcomes themselves—are multidimensional in nature. Crucially, multidimensionality in politics presents voters—and parties and candidates—with important informational challenges. Even with available heuristics and policy feedback, voters face the problem of understanding new issue areas and how to orient themselves when navigating this multidimensional world (Izzo, Martin and Callander, Forthcoming).

In response to these challenges, political elites frequently make claims about how various political issues are related. Indeed, political elites commonly appeal to voters by underscoring that, if they like the party's stances on one facet of policymaking, they will also also like the party in a second or third facet. In many ways, this is precisely the function of party manifestos or platform. Each begins with a "preamble," articulating the party's overarching desired outcomes for society. In 2016, the Democratic Party in the U.S., for example, states a primary goal of *equality*—political, economic, and social—among Americans. This goal is then articulated in specific policy positions on a host of different dimensions, from education, to healthcare, to redistribution. Finally, the manifesto ends with the following statement: "What makes America great is our unerring belief that we can make it better. We can and we will build a more just economy, a more equal society, and a more perfect union—because we are stronger together." If a voter agrees with those outcomes, the authors imply, she should necessarily support the party's positions not just in one policy area, but in all those to follow.

Such claims are not confined, of course, to left-leaning causes, nor are they solely the purview of parties. Indeed, prior to the so-called Republican Revolution of the 1980s in the U.S., conservative thought leaders expended considerable energy making the case that the newly developed right-wing coalition of religious, social, and economic conservatives was *not* one of simple political convenience. Instead, right-leaning political theorists, economists, and some politicians branded the movement as the New Fusionism, wherein they argued that the societal goals of social and economic conservatives were *necessarily* intertwined.

As this discussion higlights, multidimensionality creates a rich strategic environment with which political leaders must contend. Nevertheless, the decision of whether and when to expand policy-making activities to new dimensions has remained largely unexplored in a wide variety of settings. In this paper, we offer among the first formal examinations of this strategic decision, within an accountability framework.[1]

1.2 Our Approach

In our model, a candidate or party decides whether to change policy in a primary area and, if so, where to move the policy. Beyond this baseline decision, however, the candidate must also determine whether to legislate in an *secondary* policy area, again determining where in ideological space the new policy should be set. Importantly, players face uncertainty about which policy is optimal for the voter on each dimension. For example, we can think about a world where the mapping from policies to outcomes in unknown (e.g., Callander 2011). The key difference between the primary and secondary dimensions is that voters and politicians alike have initially more information on the former. Complicating the policymaker's decision, the model builds on the aforementioned

[1] The most closely related paper is Buisseret and Van Weelden (2020), which analyzes the entry and platform positioning of outsider candidates in a multidimensional world, where parties face uncertainty about the distribution of voters. Buisseret and Van Weelden (2022) also looks at a multidimensional world, and analyzes an incumbent's decision to call a referendum on a secondary policy issue in order to reveal information about the distribution of voters and thus influence the equilibrium of the platform game in the following elections. In contrast, we consider multidimensional *policymaking* by officeholders, in a world where voters themselves are uncertain about the optimal policy.

observation that some policy areas are "correlated" with others, meaning that the voter can learn about how well a candidate's program fits their preferences in one dimension by observing the policy outcome in another. Voters thus respond in the model by updating their priors based on the policy changes adopted by the incumbent candidate/party and the resulting outcomes (as in the tradition of career concerns models, e.g., Holmström 1999 or Ashworth, Bueno de Mesquita and Friedenberg 2017), and by their beliefs regarding the correlation of policy areas.

In this setting the inferences voters draw when observing outcomes depend on the policies implemented by the officeholder (as in Izzo Forthcoming). On each dimension, extreme policies tend to generate more informative outcomes. Intuitively, if a voter obtains a good outcome from an extreme liberal policy, it must be the case that such policy is aligned with the voter's interests. In contrast, the outcome of a moderate policy is much less informative: favorable shocks may allow the voter to enjoy a relatively high welfare even under a wrong policy, if the policy is not too radical. Thus, the incumbent can control the amount of voter learning on each dimension, both directly (via the policy on that dimension) and indirectly (via his choices on the other dimensions that generate learning spillovers).

This setup enables us to examine politicians' decisionmaking over issue expansion, elucidating how the ability to expand policymaking efforts can fundamentally alter the nature of policy outcomes. Moreover, we are also able to show how these dynamics change in response to a host of relevant contextual variables—most notably in response to different levels of correlation between the policy areas the politician is considering.

1.3 Preview of Main Results

We begin by analyzing a baseline case where the voter only cares about the primary dimension and learning spillovers are not possible. In this case, even if the incumbent's ideological preferences are multidimensional his strategic problem is unidimensional: his electoral chances are only a function of his policy choice on the primary dimension. In this baseline, an incumbent who is ex-ante electorally leading always has incentives to prevent direct voter learning on the primary dimension.

As a consequence, he always pursues a policy more moderate than his true ideological preferences. The opposite holds for a leading incumbent, who has incentives to gamble for resurrection and thus implement extreme policies that facilitate voter learning.

Moving to the multidimensional setting, we find that trailing incumbents in our model have strategic incentives to pursue policymaking in the secondary issue, even when they do not have any ideological preferences over it. Electorally disadvantaged incumbents exploit the secondary dimension to expand opportunities for voter learning, and therefore improve their electoral chances. Thus, even if he only cares about a single issue, policymaking under a trailing incumbent will inevitably be multidimensional when the different dimensions are sufficiently correlated. Furthermore, the incumbent will always pursue extreme policies on the secondary dimension in order to facilitate voter learning. In contrast, leading incumbents face a trade-off in multiple dimensions. Here, the correlation between the policy dimension is paramount, as high correlation actually *discourages* policymaking on a secondary dimension even if the incumbent has ideological preference over it. Thus, high correlation pushes leading incumbent to contract the scope of policymaking. By contrast, low correlation allows leading incumbents to legislate on the secondary dimension without risking their electoral advantage with voters.

Second, regardless of correlation, we also find that the presence of a secondary dimension has an impact on the policies the incumbent pursues on the primary one. When his ideological preferences push him to act on the secondary dimension, a leading incumbent tends to implement primary-dimension policies that are more extreme than he would pursue otherwise. He does so to facilitate direct information generation on the main dimension, and thus counteract the negative electoral effect of learning spillovers. In contrast, the second dimension pushes a trailing incumbent toward more moderate policymaking on the primary dimension, in order to magnify the positive electoral impact of learning spillovers.

Together, these results show how multidimensionality itself can influence the extremity of politics, and how a disconnect between policymaking areas can actually encourage more policymaking outside of the primary issue dimension.

Finally, our last result analyzes which features of a policy dimensios make it more likely to be activated by the officeholder. We show that leading incumbents will tend to expand the scope of policymaking to incorporate dimensions over which their ideological preferences are more moderate and for which voter's inference problem is more complex. In contrast, trailing incumbents will tend to open new dimensions they have extreme ideological taste on, and for which the voter's information challenges are easier to overcome.

2 The Model

Players and actions. There is an incumbent, I, a challenger, C, and a representative voter, V. In each period, the incumbent chooses whether to act on each of two policy dimensions, $D \in \{X, Z\}$. If he chooses to act on dimension D, then he selects a policy $d_t \in \mathbb{R}$ to be implemented. If he chooses not to act on dimension D, then the status quo d_{sq} remains in place in that period. For simplicity, we normalize the status quo on each dimension to 0, $x_{sq} = z_{sq} = 0$.

Information. Politicians' ideal points are common knowledge and, to streamline the analysis, symmetric around 0: $x_I = -x_C > 0$ and $z_I = -z_C > 0$.

Conversely, the policy that maximizes voter's welfare is unknown. Specifically, on each dimension d the voter's optimal policy d_v can take one of two values: $d_v \in \{-\alpha, \alpha\}$. Players share common prior beliefs that

$$\Pr(x_v = \alpha) = \pi$$

and

$$\Pr(z_v = \alpha | x_v = \alpha) = \Pr(z_v = -\alpha | x_v = -\alpha) = \rho \geq \frac{1}{2}$$

Thus, players believe the dimensions are positively correlated in a symmetric way and the ex-ante probability that $z_v = \alpha$, which we denote as β, is given by $\rho\,\pi + (1 - \rho)\,(1 - \pi)$.

Notice that in our setting players initially have more information about the voter's ideal policy on dimension X than on dimension Z. To reflect this, we will refer to X as the primary policy dimension, and Z as the secondary one.

Payoffs. Player $i \in \{I, V, C\}$'s per-period utility on dimension D is

$$u^i_{t,D}(d_t) = -\lambda^d_i (d_t - d_i)^2 + \varepsilon_{d,t},$$

where $\varepsilon_{d,t} \sim U \in \left[-\frac{1}{2\psi_d}, \frac{1}{2\psi_d}\right]$, and $\lambda^x_i = \lambda_i = 1 - \lambda^z_i$. Thus, λ_i is the weight player i puts on dimension X.

Timing. The timing is as follows.

1. For each dimension $D \in \{X, Z\}$, I decides whether to act by choosing a policy $d_1 \in \mathbb{R}$, or instead keep the status quo d_{sq}.
2. V observes I's choice and her realized utility on each dimension.
3. V chooses whether to re-elect I or replace her with C.
4. The winner of the election takes office, then chooses whether to act on each dimension or instead keep the status quo.

3 Equilibrium Analysis

We proceed by backward induction. In the second period, both incumbent and challenger implement their preferred policies on each dimension if elected. Thus, the voter faces a selection problem, wanting to elect the office-holder who is more aligned with her own multidimensional ideal point. The voter, however, does not know what the optimal policy is on each dimension. Further, she could be more aligned with the incumbent on one dimension and with the challenger on the other dimension. Thus, her electoral decision depends on her beliefs over the optimal policy on both dimensions X and Z.

Formally, denote by μ^d the voter's posterior that her ideal policy on dimension d is a right-wing one, where $\mu^d = \Pr(d_v = \alpha)$, and recall that politicians' bliss points are symmetric around zero on each dimension. Then, the following holds:

Lemma 1. *In equilibrium, the voter reelects the right-wing incumbent if and only if*

$$\mu^x > \frac{1}{2} - \frac{(1-\lambda_v)z_I}{\lambda_v x_I}\frac{2\mu^z - 1}{2} \equiv \widehat{\mu}^x(\mu^z). \tag{1}$$

Proof. All Proofs are collected in the Appendix. □

When the voter only cares about the primary dimension ($\lambda_v = 1$), it follows from (1) that the right-wing incumbent is reelected as long as the voter believes her ideal point on dimension X is more likely to be a right-wing one ($\widehat{\mu}^x = \frac{1}{2}$). Instead, when the voter cares about both dimensions ($\lambda_v < 1$) she becomes more lenient with the incumbent on dimension Z the more she likes him on dimension X (and vice-versa). This effect is stronger the more (less) polarized candidates are on the secondary (primary) policy dimension.

To streamline the presentation of the results, we will assume that the voter cares sufficiently more about the primary dimension X. Specifically, if the voter believes her ideal point is right-wing on dimension dimension X (i.e., $\mu^x = 1$) but left-wing on Z ($\mu^z = 0$), she prefers to re-elect the right-wing incumbent. Formally:

Assumption 1. $\lambda_v > \frac{z_I}{x_I + z_I}$.

3.1 Voter Learning

Moving one step backwards, we now study how the voter forms her posterior beliefs μ^z and μ^x. Here, the voter observes her realized utility on each dimension, and updates her beliefs by applying Bayes rule (as in Izzo (Forthcoming)). The innovation in this model is that when policies span multiple *correlated* dimensions, the voter learning is twofold: direct, and indirect. The voter's realized utility on each dimension provides her with new information on her optimal platform on that dimension (*direct learning*), but also on the policy-relevant state of the world on the others (*indirect learning*). Thus, the voter's posterior belief on x_v, is a function of her realized utility on both dimensions X and Z, and similarly for μ^z.

We begin by considering the direct channel. We characterize the voter's *interim* posterior beliefs on each dimension D, i.e., her beliefs as a function of her realized utility on that dimension only. The statements below refer to dimension X, expressions for dimension Z are analogous. The key feature of the voter learning in this setup is that more extreme policies generate more information:

Lemma 2 (Direct Learning)**.** *Define $\tilde{\mu}^x$ as the voter's (interim) posterior upon observing the outcome on dimension X. We have:*

(i) The outcome on dimension X is either fully informative of x_v or fully uninfomative, i.e., $\tilde{\mu}^x \in \{0, \pi, 1\}$;

(ii) There exists a policy x' such that if $|x_1| \geq |x'|$, then $\mu^x \neq \pi$. Specifically, $x' = \frac{1}{4\alpha\psi_x\lambda_v}$.

(iii) Let $L_x = 1$ denote the event that the realized outcome on X is fully informative. Then:

$$\Pr(L_x = 1 || x_t| < |x'|) = 4\alpha |x_t| \lambda_v \psi_x. \tag{2}$$

Lemma 2 shows that, upon observing outcomes on each dimension, the voter either learns everything or nothing about her true preferences on that dimension. Furthermore, she is more likely to discover her true preferences as the implemented policy becomes more extreme.

The logic behind this result is intuitive. In expectation, the voter's payoff is different under the two states of the world (i.e., the two possible values of her ideal policy) for any policy other than 0. However, the voter's realized utility on each dimension is also a function of a random period-specific shock. This, in turn, creates a partial overlap in the support of the payoff *realization.* When the policy is sufficiently moderate ($x_1 \in (-x', x')$), there exists a range of payoffs that may realize (i.e., be actually observed) whether the voter's true bliss point takes a positive or a negative value.

Clearly, if the payoff realization falls outside this range, it constitutes a fully informative signal. There is only one state of the world that could have generated that specific realization: the voter simply likes the policy too much, or too little, for this to be justified as a consequence of the shock. Thus, upon observing her payoff, the voter discovers her true preferences (i.e., the value of x_v). Conversely, any payoff realization that falls inside the range of overlap is completely uninformative. Since the shock is uniformly distributed, any such realization has exactly the same probability of being observed under the two states of the world. Thus, the voter learns nothing and her interim posterior remains at her prior beliefs. As the implemented policy becomes more extreme, the range of overlap becomes smaller, and the voter is more likely to directly learn her true preferences.

Figure 1 provides a graphical illustration of this result. The blue and orange functions represent the conditional outcome distributions (i.e., the distributions of the voter's realized utility), under a positive and a negative state of the world, respectively. In the left panel, a moderate right-wing policy $x > 0$ produces a large overlap in the conditional distributions. In the central panel, a more extreme policy $\hat{x} > x$ produces a much smaller overlap. As a consequence, the voter's inference problem is much easier in the second case. Finally, in the third panel the policy is sufficiently extreme that there is no overlap in the conditional distributions, $\tilde{x} > x'$, and the voter always learns the true value of x_v.

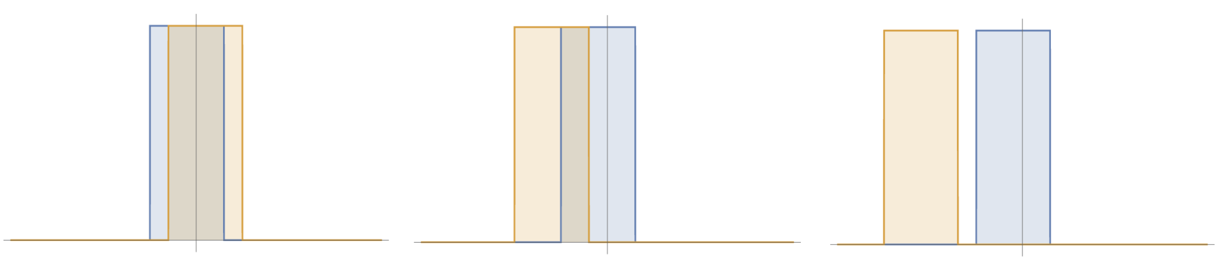

Figure 1: **Voter Learning.** The three plots display voter learning under a positive (blue function) and negative (orange function) state of the world. The policy extremism increases from the left to the right panel.

Lemma 2 indicates that the voter may learn her optimal policy on each dimension *directly* by observing how much she liked or disliked the implemented policy on that issue. Our next results indicates that such direct learning (through an extreme x_1) is sufficient but not necessary for the voter to obtain new information. Because dimensions are correlated, it could be the case that the voter observes an uninformative outcome on X and yet learns more about x_v via the second dimension, Z (and vice versa). Intuitively, if the voter finds that she likes liberal policies on social or economic issues, she may tend to acquire a more positive attitude towards liberal policies on healthcare as well.

Formally, the next result shows how the voter's posterior belief on X depends on the second dimension. Recall that μ^x is the voter's posterior on x_v, as a function of her realized utility on *both* dimensions X and Z. Instead, we denote $\tilde{\mu}^x$ the voter's *interim* posterior, as a function only of her realized utility on X. The next result derives the voter posterior on the primary dimension as a function of i) her interim posterior $\tilde{\mu}^x$ and ii) the possible outcomes on the secondary dimensions. The proof simply follows by applying Bayes rule, and is therefore omitted.

Lemma 3 (Full learning)**.** *Suppose that the voter observes an uninformative outcome on Z. Then:*

$$\mu^x = \tilde{\mu}^x \tag{3}$$

Suppose that the voter observes an informative outcome on Z and learns that $z_v = \alpha$. Then:

$$\mu^x(\tilde{\mu}^x, \alpha, \rho) = \frac{\tilde{\mu}^x \rho}{\tilde{\mu}^x \rho + (1 - \tilde{\mu}^x)(1 - \rho)} \tag{4}$$

Suppose instead that the voter observes an informative outcome on Z and learns $z_v = -\alpha$. Then:

$$\mu^x(\tilde{\mu}^x, -\alpha, \rho) = \frac{\tilde{\mu}^x(1 - \rho)}{\tilde{\mu}^x(1 - \rho) + (1 - \tilde{\mu}^x)\rho} \tag{5}$$

The previous Lemma highlights that when no direct learning occurs on X (i.e., $\tilde{\mu}^x = \pi$), learning spillovers determine the voter's posterior. If the voter learns that her ideal point on Z is a right-wing (left-wing) one, she becomes more convinced that her optimal policy on X is right-wing (left-wing) as well. The higher the correlation across dimensions ρ, the stronger these learning spillovers.

Notice that, under our assumption that $d_{sq} = 0$, if the officeholder chooses not to act on dimension d in the first period then the voter's realized utility on that dimension is always uninformative (i.e., there is full overlap in the conditional outcome distributions). Thus, the voter may only learn indirectly via the realized outcome on other correlated dimensions. We use this normalization to simplify notation, but our qualitative results below (in particular our results on when the incumbent chooses to act on the secondary dimension) simply require that if the policy remains at the status quo, then no new (direct) information is generated on that dimension. For example, we could assume that if $d_1 = d_{sq}$, then the voter does not observe a new realization of her utility on d, so that no direct learning ever occurs.

4 The Incumbent's Problem

The results of the above sections highlight that, in our setting, the policy that is implemented today influences voter learning and thus her optimal retention choice. As such, the incumbent's choice on each dimension has two effects on his expected payoff. First, a *static* ideological effect: the incumbent's first-period payoff increases as the implemented policy gets closer to his ideal point. Second, a *dynamic* information effect: the incumbent's expected second-period payoff depends on the first-period policy via voter learning. In turn, this information effect emerges via two channels. The implemented policy on each dimension influences the probability of the voter *directly* learning her true ideal point on that dimension. In addition, the correlation across dimensions generates learning spillovers, so that the implemented policy on X can also *indirectly* influence voters beliefs on Z (and vice versa).

These two effects, ideological and information, generate a potential trade-off for the incumbent: one the one hand, he wants to set a policy close to his ideal point, on the other, such policy might not generate enough information (or generate too little). This trade-off clearly appears in the incumbent maximization problem, which we can express as follows:

$$\max_{x_1, z_1} -\lambda_I(x_1 - x_I)^2 - (1 - \lambda_I)(z_1 - z_I)^2 - (1 - \mathbb{P}(x_1, z_1))(\lambda_I 4x_I^2 + (1 - \lambda_I)4z_I^2), \tag{6}$$

where $\mathbb{P}(x_1, z_1)$ denotes the incumbent's retention probability which is a function of the voter posterior and the incumbent policy choices.

The first-order necessary conditions for an interior maximum are, respectively:

$$(x_1) \quad -2\lambda_1(x_1 - x_I) + \frac{\partial \mathbb{P}(x_1, z_1)}{\partial x_1}\left(\lambda_I 4x_I^2 + (1-\lambda_I)4z_I^2\right) = 0 \tag{7}$$

$$(z_1) \quad -2(1-\lambda_1)(z_1 - z_I) + \frac{\partial \mathbb{P}(x_1, z_1)}{\partial z_1}\left(\lambda_I 4x_I^2 + (1-\lambda_I)4z_I^2\right) = 0 \tag{8}$$

Recall that, from Lemma 2, more extreme policies that move farther from the status quo are more likely to generate informative outcomes. Thus, depending on whether information is electorally beneficial (i.e., $\frac{\partial \mathbb{P}(x_1,z_1)}{\partial x_1} > 0$ for right-wing X policies and $\frac{\partial \mathbb{P}(x_1,z_1)}{\partial z_1} > 0$ for right-wing Z polices) or not (i.e., $\frac{\partial \mathbb{P}(x_1,z_1)}{\partial x_1} < 0$ and $\frac{\partial \mathbb{P}(x_1,z_1)}{\partial z_1} < 0$) the incumbent will have incentives to distort his choice either to the extreme or towards the status quo $d_{sq} = 0$.

In what follows, we will see that whether one or the other distortion emerges in equilibrium depends on the incubment's ex-ante prevailing electoral chances *and* whether he chooses to act only on a single dimension or expand the scope of policymaking.

4.1 Unidimensional Benchmark ($\lambda_v = 1$)

First, suppose that $\lambda_v = 1$, so that the voter only cares about the primary dimension X. This effectively makes the incumbent's problem unidimensional: the voter does not care about dimension Z directly, *and*, as the next result shows, no learning spillovers are possible.

Remark 1. *Suppose $\lambda_v = 1$. Then, the voter does not learn anything upon observing her realized utility on dimension Z and $\tilde{\mu}^x = \mu^x$.*

When $\lambda_v = 1$, the voter's realized utility on Z is pure noise ($u^v_{1,z} = \varepsilon_{z,1}$). Thus, the incumbent's retention chances are not a function of his choice on the secondary dimension: even though the incumbent's ideological preferences are multidimensional, his strategic problem is unidimemsional. In this world, we can easily characterize the incumbent's optimal policy on each dimension. Denotes this policy d_u, where the u subscript indicates that this is the optimal policy on dimension d in a world where the incumbent's strategic problem is unidimensional.

Straightforwardly, the incumbent always implements his ideal point on the secondary dimension Z, $z_u = z_I$. The assumption that $\lambda_v = 1$ implies that the implemented policy on the Z dimension has no information effect, therefore the optimal choice must be at the incumbent's ideologically preferred point. In contrast, the implemented policy on X influences his expected payoff via both the static and dynamic channels. The incumbent chooses x_1 to maximize:

$$-\lambda_I(x_1 - x_I)^2 - \Big(1 - \mathbb{P}(x_1)\Big)\Big(\lambda_I 4x_I^2 + (1-\lambda_I)4z_I^2\Big), \tag{9}$$

where, as a reminder, $\mathbb{P}$ is the probability of winnig reelection. Assume without loss of generality that the voter reelects the incumbent when indifferent. Recall that $L_x(x_1) = 1$ denotes the event that the voter discover her ideal point on dimension X. Then, using the results of the previous section, we obtain:

Remark 2. *Suppose that $\lambda_v = 1$. Then,*

- $\mathbb{P}(x_1) = 1 - \Pr(L_x(x_1) = 1)(1-\pi) = \min \in \{1, 4\alpha\psi_x(1-\pi)|x_1|\}$ *when* $\pi \geq \frac{1}{2}$*;*
- $\mathbb{P}(x_1) = \Pr(L_x(x_1) = 1)\pi = \max \in \{0, 1 - 4\alpha\psi_x\pi|x_1|\}$ *when* $\pi < \frac{1}{2}$*.*

Recall that, given our symmetry assumption $x_I = -x_C$, at $\pi = \frac{1}{2}$ the voter is ex-ante indifferent between reelecting the incumbent or replacing him with the challenger. Thus when $\pi < \frac{1}{2}$, the incumbent is ex-ante *trailing*: if the voter receives no new information, she will choose to oust him. This right-wing incumbent is reelected if and only if the voter discovers that her ideal point x_v is a right-wing policy. Recall that the probability that the voter discovers her true preferences on dimension X is (weakly) increasing as x_1 moves away from 0 in each direction. Thus, a trailing incumbent's probability of being reelected (weakly) increases as the implemented policy x_1 becomes more extreme.

The opposite holds when $\pi \geq \frac{1}{2}$, so that the incumbent is ex-ante electorally *leading*. If the voter learns nothing new, this incumbent will be reelected for sure. Thus, his probability of being retained is (weakly) decreasing as x_1 moves away from 0.

The next result follows straightforwardly. Here, we use x_u to denote the incumbent's equilibrium choice in this unidimensional benchmark:

Proposition 1. *Suppose that $\pi \geq \frac{1}{2}$. Then, in equilibrium:*

$$x_u = \max \in \left\{0, x_I - \frac{4\alpha\psi_x(1-\pi)}{\lambda_I}\Big(\lambda_I(x_I - x_C)^2 + (1-\lambda_I)(z_I - z_C)^2\Big)\right\} < x_I.$$

- *Suppose instead $\pi < \frac{1}{2}$. Then, in equilibrium:*

$$x_u = \min \in \left\{\max \in \left\{x_I, \frac{1}{4\alpha\psi_x\lambda_v}\right\}, x_I + \frac{4\alpha\psi_x\pi}{\lambda_I}\Big(\lambda_I(x_I - x_C)^2 + (1-\lambda_I)(z_I - z_C)^2\Big)\right\} \geq x_I.$$

In equilibrium, a trailing incumbent distorts policy to the extreme, away from both his static optimum and the status quo (normalized to 0), in order to facilitate voter learning.[2] In this case, we say that the incumbent *gambles* on this policy dimension. In contrast, a leading incumbent is risk averse, and distorts policy towards 0 so as to minimize information. Notice that, since any pair of policies x and $-x$ induces the same amount of learning (Lemma 2), the right-wing incumbent never implements a policy to the left of 0.

Having characterized equilibrium policy in this unidimensional benchmark, we now move to analyzing the incumbent's policy choices in the multidimensional case (i.e., when $\lambda_v < 1$). Our objective is to study the conditions under which the incumbent has strategic incentives to act on the secondary policy dimension, and characterize how this influences his optimal choice on the primary one.

[2] If $x_I > \frac{1}{4\alpha\psi_x\lambda_v}$, the incumbent's ideal point guarantees full learning. A trailing incumbent thus faces no trade-off, and in equilibrium $x_u = x_I$.

4.2 Multidimensional World ($\lambda_v < 1$)

So far, we have assumed that $\lambda_v = 1$ so that the voter only cares about dimension X, and no indirect learning is possible via dimension Z (Lemma 2). In other words, the incumbent's retention chance is not a function of his policy choice on Z. In this section, we relax this assumption to study the incumbent's multidimensional problem.

It is useful to begin by analyzing a baseline where $\lambda_v < 1$, but $\lambda_I = \lambda_C = 1$. In other words, the voter cares about both dimensions, but politicians only care about the primary one X. Further, we will assume that when indifferent an officeholder chooses not to act on dimension d.

Thus, even though she cares about both dimensions, the voter's retention decision does not *directly* depend on her beliefs over Z (since she anticipates that neither I nor C will act on Z in the second period). More specifically, the voter's optimal retention rule is exactly the same as in the unidimensional case: she retains the right-wing incumbent if and only if $\mu^x > \frac{1}{2}$. However, by Lemma 3, the voter's posterior on X is a function of her realized utility on Z. Therefore, even though the incumbent's ideological preferences are unidimensional, his strategic problem is *multidimensional*. These assumptions thus allow us to isolate the strategic incentives emerging solely due to the learning spillovers across dimensions.

First, we characterize the incumbent's probability of wining in this multidimensional world. Recall that $\beta = \pi\rho + (1-\pi)(1-\rho)$ is the prior probability that $z_v = \alpha$. Notice that $\beta > \frac{1}{2}$ if and only if $\pi > \frac{1}{2}$ (recall that $\rho > \frac{1}{2}$). Thus, from condition 1 we can verify that ex-ante the voter prefers to re-elect the right-wing incumbent if and only if $\pi > \frac{1}{2}$. Thus, as in the unidimensional world, we will refer to the incumbent as leading if $\pi > \frac{1}{2}$, and trailing otherwise. Remark 3 follows from this observation and Assumption 1:

Remark 3. *Suppose that $\lambda_v < 1$. Then, there exists unique $\widehat{\rho}_z(\pi)$ and $\tilde{\rho}_z(\pi)$ s.t.*

- $\mathbb{P}(x_1, z_1) = 1 - \Pr(L_x(x_1) = 1)(1-\pi)$ *when* $\pi \geq \frac{1}{2}$ *and* $\rho < \tilde{\rho}_z$*;*
- $\mathbb{P}(x_1, z_1) = 1 - \Pr(L_x(x_1) = 1)(1-\pi) - \Big(1 - \Pr(L_x(x_1) = 1)\Big) \Pr(L_z(z_1) = 1)(1-\beta)$ *when* $\pi \geq \frac{1}{2}$ *and* $\rho > \tilde{\rho}_z$*;*

- $\mathbb{P}(x_1, z_1) = \Pr(L_x(x_1) = 1)\pi$ *when* $\pi < \frac{1}{2}$ *and* $\rho < \widehat{\rho}_z$*;*
- $\mathbb{P}(x_1, z_1) = \Pr(L_x(x_1) = 1)\pi + \Big(1 - \Pr(L_x(x_1) = 1)\Big) \Pr(L_z(z_1) = 1)\beta$ *when* $\pi < \frac{1}{2}$ *and* $\rho > \widehat{\rho}_z$*.*

Recall that $\Pr(L_d(d_1) = 1)$ is the probability that the voter observes an informative outcome on dimension d, which increases as d_1 becomes more extreme. First, it is easy to see that Assumption 1 implies that when ρ is too low, outcomes on Z are electorally irrelevant (because the voter cares more about X than Z, and a low correlation implies that learning spillovers are too weak to dominate on the voter's prior π). Then, the incumbent's retention probability is as in the unidimensional $\lambda_v = 1$ case. Suppose instead the correlation is sufficiently strong. Then, in contrast to the unidimensional case, a leading incumbent will lose the election even when no direct learning occurs on X, if the outcome on Z is informative and unfavorable. Conversely, a trailing incumbent will be able to be reelected even when the outcome on the primary dimension is uninformative, if he generates favorable information on the secondary one. As an aside, we note that this Remark applies even in the case in which $\lambda_I < 1$, analyzed in the next section.

From this, we can easily identify conditions under which these spillovers create strategic incentives for the incumbent to act on the secondary policy dimension:

Proposition 2. *Suppose* $\lambda_v < 1$ *and* $\lambda_I = \lambda_C = 1$*. Then, the incumbent chooses to act on* Z *(i.e.,* $z_1 \neq z_{sq}$*) if and only if* $\pi < \frac{1}{2}$ *and* $\rho > \widehat{\rho}_z$*. Further, we have that* $\widehat{\rho}_z = 1 - \pi > \frac{1}{2}$*.*

Here, the incumbent acts on Z if and only if he has incentives to facilitate *indirect* learning on X. It follows from Remark 3 that a leading incumbent (i.e., $\pi > \frac{1}{2}$) never wants to act on Z, since he wants to prevent the voter from obtaining any new information. Suppose instead that $\pi < \frac{1}{2}$, so that the incumbent is ex-ante trailing. Then, he wants to facilitate indirect learning on X, in hopes of overcoming his initial disadvantage and jumping above the retention threshold. As highlighted above, however, outcomes on the secondary dimension remain electorally irrelevant if the correlation ρ is too small. Thus, a trailing incumbent is indifferent between acting on the secondary dimension and keeping the status quo and (by assumption) chooses not to act. If instead ρ is sufficiently large, the trailing incumbent can exploit learning spillovers to increase his probability of resurrecting himself. In equilibrium he will therefore always choose to expand the scope of policymaking to the secondary dimension, even if he has no ideological taste for it.

The next Corollary follows straightforwardly from the above discussion and Lemma 2:

Corollary 1. *Suppose that in equilibrium the incumbent chooses to act on Z. Then, he always implements a fully informative policy, i.e., $z_1^* \geq z'$.*

Even though the incumbent does not have ideological preferences over dimension Z, his strategic incentives to facilitate voter learning induce policy extremism on this secondary dimension. Thus, in equilibrium we either observe inaction on Z (i.e., the policy remains at the status quo), or we observe the incumbent pursuing extreme policies.

Next, we characterize how the possibility to exploit learning spillovers from the secondary dimension influences the incumbent's policy choice on the primary one. Recall that x_u is the incumbent's optimal policy choice in the unidimensional benchmark. Then, we have:

Proposition 3. *Suppose that in equilibrium the incumbent chooses to act on the secondary dimension Z. Then, his policy choice on the primary one x_1^* satisfies $x_1^* < x_I \leq x_u$.*

This Proposition highlights that the correlation ρ generates a strategic substitution effect between policy dimensions. When a trailing incumbent cannot exploit the secondary dimension (i.e., when $\lambda_v = 1$), then he always has strategic incentives to gamble on the primary one. Recall that extreme policies generate more information, therefore this incumbent always implements a policy more extreme than his ideological preference, $x_u > x_I$. When instead the incumbent has the possibility to exploit learning spillovers (i.e., $\lambda_v < 1$ and $\rho > 1 - \pi$), this induces moderation on the primary dimension: $x_1^* < x_I$.

To understand this, recall that the outcome on Z may influence the voter's retention decision only if she does not learn about x_v directly (as otherwise she reaches a degenerate interim posterior $\tilde{\mu}^x$).[3]

[3]Notice that this holds true even if the voter anticipates the second-period office holder to implement his own bliss point on Z. This is due to Assumption 1, which guarantees that everything else being equal dimension X has a larger relative weight for the voter. As a consequence, direct learning on X always determines the voter's optimal choice, even if she also obtains direct learning on Z.

Therefore, in order to benefit from the learning spillovers, the incumbent must avoid generating an informative outcome on X. When he can strategically exploit the correlation across dimensions, even a trailing incumbent then implements a policy more moderate than his static optimum on X, in sharp contrast with the results of the unidimensional baseline.

Further, notice that this result implies that the trailing incumbent finds it more profitable to gamble on the primary dimension when he can exploit the secondary one. Recall that the incumbent is trailing if and only if $\pi < \frac{1}{2}$. Thus, even though the trailing incumbent needs to generate information in order to be reelected, an informative outcome is more likely to reveal to the voter that she is aligned with the challenger's preferences. This is true on both dimensions, but since $\rho < 1$, we have that $\pi < \frac{1}{2}$ implies $\beta < \pi$. In other words, since the players have less accurate information on the secondary dimension, the state $z_v = \alpha$ is ex-ante more likely than $x_v = \alpha$, therefore the incumbent is more likely to generate a favorable outcome on Z. Thus, the trailing right-wing incumbent prefers to gamble on Z, hoping to exploit a false positive: generate a favorable outcome on Z and thus induce the voter to positively update on x_v, even if the state of the world on the primary dimension is actually a left-wing one.

These results therefore highlight that, in equilibrium, the incumbent will never gamble on both dimensions. Rather, if the correlation is too low to exploit the learning spillovers, he will have no strategic incentives to act on Z and will continue gambling on X. If instead the correlation is high, he will gamble on Z but become risk averse on X.

Finally, we characterize how the magnitude of the correlation ρ influences the incumbent's policy choice on the primary dimension:

Corollary 2. *Suppose that in equilibrium the incumbent chooses to act on the secondary dimension Z. Then, we have that $\frac{\partial x_1^*}{\partial \rho} > 0$.*

The stronger the correlation across dimensions, the weaker the substitution effect described above. Recall that in equilibrium the incumbent acts on Z only when $\pi < \frac{1}{2}$, i.e., the true state on X is more likely to be disfavorable for the right-wing incumbent. Thus, as the correlation increases it becomes more and more likely that $z_v = -\alpha$, so that if the voter observes an informative outcome on

Z she updates against the incumbent. Thus, as ρ increases it becomes less likely that the incumbent is able to resurrect his reelection chances by gambling on Z, and the incentives to prevent direct learning on X become weaker. As a consequence, whenever the incumbent chooses to act on Z in equilibrium, the policy on X becomes more extreme as ρ increases.

4.3 General Model

The results of the previous section are useful to isolate the strategic incentives generated by the learning spillovers. Here, we complete the analysis by studying the incumbent's policy choice on each dimension in the general model, where both the voter and the politicians care about both dimensions (i.e., $\lambda_i < 1$ for $i \in \{I, V, C\}$) In contrast with the analysis presented above, the incumbent now has both strategic *and ideological* preferences over the secondary dimension Z. As above, our goal is to characterize the conditions under which the incumbent chooses to open the secondary dimension, and study how this influences his policy on the primary one.

Proposition 4. *Suppose $\pi < \frac{1}{2}$. Then the incumbent always acts on the secondary dimension Z. Suppose instead $\pi > \frac{1}{2}$. Then the incumbent acts on Z if and only if*

- *The correlation ρ is sufficiently low, or*
- *The correlation ρ is high and λ_I is sufficiently small.*

A trailing incumbent has both ideological and strategic reasons to act on the secondary dimension. In equilibrium, he will therefore always choose to do so. In contrast, a leading incumbent faces a trade-off. On one hand, he has ideological preferences over the secondary dimension and would therefore statically find it optimal to act on it. On the other hand, as the results of the previous section demonstrate, acting on the secondary dimension hurts his retention chances and thus his expected future payoff. If the correlation between dimensions is sufficiently low, the leading incumbent can survive reelection even if the outcome on the secondary dimension reveals damaging information. He can therefore implement his preferred policy on the secondary dimension while avoiding the negative electoral consequences. Suppose instead the correlation ρ is high, then the trade-off discussed above is binding, and the incumbent only acts on Z if his ideological preferences are sufficiently strong (i.e., λ_I is low).

Our second result describes how strategic incentives to act on the secondary policy dimension influence the incumbent's policy choice on the primary one. As above, a trailing incumbent becomes more moderate on the primary dimension when he can act on a secondary one. However, the result is reversed for a leading incumbent: here, the presence of dimension Z generates more extremism on X:

Proposition 5. *Suppose that the incumbent chooses to act on the secondary dimension. Then*

- *When $\pi < \frac{1}{2}$, we have $z_1^* \geq z_u = z_I$ and $x_1^* \leq x_u$;*
- *when $\pi > \frac{1}{2}$, we have $z_1^* \leq z_u = z_I$ and $x_1^* \geq x_u$.*

Recall that d_u is the equilibrium policy on dimension d in the unidimensional world, i.e., the world in which the voter only cares about the primary dimension ($\lambda_v = 1$).

The intuition for the case in which $\pi < \frac{1}{2}$ is exactly as described in the previous section. The trailing incumbent has incentives to gamble on Z, where a false positive is more likely. Furthermore, the correlation across dimensions generates a substitution effect, whereby the incumbent has incentives to moderate on the primary dimension in order to exploit learning spillovers from the secondary one.

Suppose instead $\pi > \frac{1}{2}$, i.e., the incumbent is leading. As discussed above, the leading incumbent does not have strategic incentives to act on Z, since his retention chances are maximized when the voter learns nothing new. However, because $\lambda_I < 1$, the incumbent has ideological preferences over Z and therefore sometimes chooses to act on this secondary dimension. When he does so, he is undertaking more electoral risk, since he may generate an informative and unfavorable outcome on Z and hurt his retention chances. Further, since $\pi > \frac{1}{2}$ implies $\beta < \pi$, an unfavourable outcome is ex-ante more likely on dimension Z than on X. Also recall that, given our Assumption 1, direct learning on X renders outcomes on Z electorally irrelevant. Taken together, these two observations imply that in order to counteract the detrimental effects of learning on the secondary dimension, the leading incumbent has incentives to facilitate direct learning on the primary one. In other words, a leading incumbent - who is always risk-averse in a unidimensional world - becomes risk-loving on the primary dimension when he chooses to expand the scope of policymaking. As a consequence, the leading incumbent will distort policy towards 0 on the secondary dimension, while on the primary one will implement policies more extreme than in the unidimensional world.

Before concluding, we characterize the features that render a policy dimension more likely to be addressed by the incumbent. Suppose that multiple secondary dimensions are available, each characterized by a different correlation ρ_d with the primary dimension X. Suppose the incumbent is resource constrained, so that he cannot act on all available dimensions. To restrict our attention to the most interesting cases, we also assume that for any available dimension ρ_d is sufficiently large that learning spillovers are always electorally relevant. Then, we have:

Proposition 6. *All else equal, a leading incumbent ($\pi > \frac{1}{2}$) will (weakly) prefer to open the dimensions:*

- *He is less extreme on, or*
- *That have higher correlation with the primary dimension, or*
- *For which outcomes are more noisy.*

Suppose instead the incumbent is trailing ($\pi < \frac{1}{2}$). Then, all else equal, the incumbent (weakly) prefers to open the dimensions:

- *He is more extreme on, or*
- *That have lower correlation with the primary dimension, or*
- *For which outcomes are less noisy.*

The intuition is as follows. Consider first a leading incumbent. If he chooses to act on a secondary policy dimension, he will find it optimal to implement a moderate policy on this dimension so as to mitigate the electoral downsides. This is less costly if his ideologically preferred policy is a moderate one. At the same time, a dimension for which outcomes are more noisy carries less electoral risk and will therefore be more appealing. Finally, recall that an incumbent is leading when $\pi > \frac{1}{2}$, i.e., the state of the world on the primary dimension is likely to be in his favor. The higher the correlation with this primary dimension, the higher the likelihood that the policy outcome on the additional dimension would also be favorable to the incumbent. Again, this makes the dimension more appealing.[4] The opposite intuition underlies the results for a trailing incumbent.

5 Conclusion

As our analysis underscores, the introduction of multidimensionality within an accountability setting dramatically influences the incentives that incumbents face, as they make decisions about whether and how to change policy. Indeed, whereas incumbents in the unidimensional setting frequently pursue extreme policies in order to encourage voter learning, the presence of a secondary dimension complicates this decision. For the trailing incumbent, the possibility of policymaking in multiple dimensions presents greater opportunities for voter learning. As a result, such incumbents always expand their policymaking—and, at the same time, *moderate* on the primary dimension. For a leading incumbent, instead, the problem is reversed. This incumbent tends to pursue moderate policy in a unidimensional world, and multidimensionality generates incentives for extremism on the primary policy dimension.

Beyond multidimensionality's overall effect on policymaking and accountability, however, our incorporation of a correlation term between policymaking dimensions generates interesting nuances in policymaking and voter learning. Indeed, in order for trailing incumbents to leverage spillover learning—or for leading incumbents to avoid it—voter learning *between* issue areas is highly consequential. That is, if correlation between the issues is high, voter learning will spill over between the dimensions, altering the politician's willingness to pursue policy change.

Together, we believe these findings underscore the importance of incorporating multiple dimen-

[4] Let us emphasize again that this is due to the assumption that there is no available dimension for which the correlation is so low that outcomes become electorally irrelevant, i.e., $\rho_d > \max\{\tilde{\rho}_z, \widehat{\rho}_z\}$ for all available dimensions.

sions into models of policymaking, especially those involving voters. Given the rise of populism and the consequent interest in issue areas such as immigration and trade policy, both politicians and voters appear to clearly care about and pursue more issues than typical left-right economic ones. We believe models of policymaking should reflect these changes, even in cases when logrolling or bargaining is not involved. Voters can learn about policymaking in one area by observing activity in another, and they are often *encouraged* to do so by activists, partisan media, and other outlets in modern political life.

Moreover, in order to understand not only whether to expand policymaking but to *where*, we believe it is imperative for theoretical and empirical models to think carefully about this correlation between issue areas. Indeed, particularly in a polarized era in which parties and candidates are thought to be more consistent across issue areas than in eras past, the possibility for spillover learning is important both substantively and strategically. Political scientists have long pointed to the importance of understanding "what goes with what" and the extent to which voters relate issue areas in their mind (e.g., Converse 1964). We show that this association is consequential for policymaking, and we hope that similar parameters are incorporated into other models policymaking, accountability, and delegation.

References

Alt, James, Ethan Bueno de Mesquita and Shanna Rose. 2011. "Disentangling accountability and competence in elections: evidence from US term limits." *The Journal of Politics* 73(1):171–186.

Ashworth, Scott, Ethan Bueno de Mesquita and Amanda Friedenberg. 2017. "Accountability and information in elections." *American Economic Journal: Microeconomics* 9(2):95–138.

Bachrach, Peter and Morton S Baratz. 1963. "Decisions and nondecisions: An analytical framework." *American political science review* 57(3):632–642.

Banks, Jeffrey s. and John Duggan. 2000. "A Bargaining Model of Collective Choice." *American Political Science Review* 94(1):73–88.

Banks, Jeffrey S and John Duggan. 2006. "A general bargaining model of legislative policy-making." *Quarterly Journal of Political Science* 1(1):49–85.

Bawn, Kathleen, Martin Cohen, David Karol, Seth Masket, Hans Noel and John Zaller. 2012. "A theory of political parties: Groups, policy demands and nominations in American politics." *Perspectives on Politics* 10(3):571–597.

Buisseret, Peter and Richard Van Weelden. 2020. "Crashing the party? Elites, outsiders, and elections." *American Journal of Political Science* 64(2):356–370.

Buisseret, Peter and Richard Van Weelden. 2022. "Pandora's Ballot Box: Electoral Politics of Direct Democracy." *arXiv preprint arXiv:2208.05535* .

Callander, Steven. 2011. "Searching for good policies." *American Political Science Review* 105(4):643–662.

Converse, Philip E. 1964. "The Nature of Belief Systems in Mass Publics." *Ideology and Discontent, ed. David Apter* .

Downs, Anthony. 1957. “An economic theory of political action in a democracy.” *Journal of political economy* 65(2):135–150.

Fiorina, Morris P. 1978. “Economic retrospective voting in American national elections: A micro-analysis.” *American Journal of political science* pp. 426–443.

Holmström, Bengt. 1999. “Managerial incentive problems: A dynamic perspective.” *The review of Economic studies* 66(1):169–182.

Izzo, Federica. Forthcoming. “Ideology for the Future.” *American Political Science Review* .

Izzo, Federica, Gregory J Martin and Steven Callander. Forthcoming. “Ideological Competition.” *American Journal of Political Science* .

Appendix

Main Results - Proofs

Proof of Lemma 1. The voter re-elects the right-wing incumbent if and only if the expected payoff from voting for I given the information received in $t = 1$ is greater than that of voting for C. Formally:

$$-\lambda_v[\mu^x(x_I - \alpha)^2 + (1 - \mu^x)(x_I + \alpha)^2] - (1 - \lambda_v)[\mu^z(z_I - \alpha)^2 + (1 - \mu^z)(z_I + \alpha)^2] > \\ -\lambda_v[\mu^x(x_C - \alpha)^2 + (1 - \mu^x)(x_C + \alpha)^2] - (1 - \lambda_v)[\mu^z(z_C - \alpha)^2 + (1 - \mu^z)(z_C + \alpha)^2]. \tag{10}$$

Plugging in the assumption that $d_I = -d_C$, the above reduces to

$$2\lambda_v\mu^x x_I\alpha - \lambda_v x_I\alpha + 2(1 - \lambda_v)\mu^z z_I\alpha - (1 - \lambda_v)z_I\alpha > 0$$

which rearranged yields:

$$\mu_v^x > \frac{1}{2} + \frac{(1 - \lambda_v)z_I}{\lambda_v x_I}\frac{(1 - 2\mu_v^z)}{2} \equiv \widehat{\mu_v^x}(\mu^z). \tag{11}$$

□

Proof of Lemma 2. We prove the statements for dimension X. Let $\mu^x \in [0, 1]$ denote V's posterior that the state of the world on dimension X is positive.

(i) A possible payoff realization for V given the incumbent's choice (x_t) has to fall within:

$$\left[-\lambda_v(x_t - \alpha)^2 - \frac{1}{2\psi_x}, -\lambda_v(x_t + \alpha)^2 + \frac{1}{2\psi_x}\right]. \tag{12}$$

Thus, if V observes $u_v^t > -\lambda_v(x_t + \alpha)^2 + \frac{1}{2\psi_x}$, she knows for sure that she likes the right policy, i.e., $\mu^x = 1$. Similarly, if V observes $u_v^t < -\lambda_v(x_t - \alpha)^2 - \frac{1}{2\psi_x}$, then $\mu^x = 0$.

The last case to consider is when u_v^t falls within the interval in Equation 12. Denote by $f(\cdot)$ the

PDF of the error term $\varepsilon_{x,t}$. We have:

$$\Pr(x_v = \alpha | u_v^t) = \frac{f\left(u_v^t + \lambda_v (x_t - \alpha)^2\right)\pi}{f\left(u_v^t + \lambda_v (x_t - \alpha)^2\right)\pi + f\left(u_v^t + \lambda_v (x_t + \alpha)^2\right)(1-\pi)}.$$

Since $\varepsilon_{x,t}$ is uniformly distributed, we have $f\left(u_v^t + \lambda_v (x_t + \alpha)^2\right) = f\left(u_v^t + \lambda_v (x_t - \alpha)^2\right)$, hence

$$\Pr(x_v = \alpha) = \pi.$$

(ii)-(iii) Now, denote by $L \in \{0, 1\}$ players' learning of x_v. There exists a value of policy x_t' such that, for any $x_t > x_t'$, the realization of u_v^t is fully informative, i.e., the interval (12) is empty. This requires:

$$-\lambda_v (x_t + \alpha)^2 + \frac{1}{2\psi_x} + \lambda_v (x_t - \alpha)^2 + \frac{1}{2\psi_x} \leq 0 \tag{13}$$

which rearranged yields:

$$x_t \geq \frac{1}{4\alpha\lambda_v\psi_x}. \tag{14}$$

Define $x' \equiv \frac{1}{4\alpha\lambda_v\psi_x}$, and assume $x_t \in [0, x']$. We have:

$$\begin{aligned}\Pr(L = 1 | \pi, 0 < x_t < x') = &\pi \Pr\left(-\lambda_v (x_t - \alpha)^2 + \varepsilon_{x,t} > -\lambda_v (x_t + \alpha)^2 + \frac{1}{2\psi_x}\right) \\ &+ (1-\pi)\Pr\left(-\lambda_v (x_t + \alpha)^2 + \varepsilon_{x,t} > -\lambda_v (x_t - \alpha)^2 - \frac{1}{2\psi_x}\right).\end{aligned}$$

Since the two probabilities are symmetric, we have

$$\begin{aligned}\Pr(L = 1 | \pi, 0 < x_t < x') = &\Pr\left(-\lambda_v (x_t - \alpha)^2 + \varepsilon_{x,t} > -\lambda_v (x_t + \alpha)^2 + \frac{1}{2\psi_x}\right) \\ = &\Pr\left(\varepsilon_{x,t} < 4\lambda_v \alpha x_t - \frac{1}{2\psi_x}\right) \\ = &4\alpha x_t \lambda_v \psi_x, \end{aligned} \tag{15}$$

where notice that the probability that V learns her true preference is increasing in x_t.

The proof for dimension Z is analogous therefore omitted. □

Proof of Remark 2. From Lemma 2 we know that $\Pr(L_x = 1|\pi, 0 < x_t < x') = 4\alpha\psi_x|x_1|$. It follows that, if $\pi \geq \frac{1}{2}$ (if $\pi < \frac{1}{2}$), $\mathbb{P}(x_1)$ is weakly decreasing (increasing) in x_1. $\square$

Proof of Proposition 1. When $\pi \geq \frac{1}{2}$ we can express I's problem as

$$-\lambda_I(x_1 - x_I)^2 - 4\alpha\psi_x x_1(1-\pi)\Big(\lambda_I(x_I - x_C)^2 + (1-\lambda_I)(z_I - z_C)^2\Big), \tag{16}$$

which yields the following FONC (which is also sufficient since the problem is concave):

$$-2\lambda_I(x_1 - x_I) - 4\alpha\psi_x(1-\pi)\left(\lambda_I(x_I - x_C)^2 + (1-\lambda_I)(z_I - z_C)^2\right) = 0, \tag{17}$$

Rearranging (17) yields:

$$x_1 = x_I - \frac{4\alpha\psi_x(1-\pi)}{\lambda_I}\Big(\lambda_I(x_I - x_C)^2 + (1-\lambda_I)(z_I - z_C)^2\Big).$$

It follows that

$$x_1 = \max\left\{0, x_I - \frac{4\alpha\psi_x(1-\pi)}{\lambda_I}\Big(\lambda_I(x_I - x_C)^2 + (1-\lambda_I)(z_I - z_C)^2\Big)\right\}. \tag{18}$$

When instead I is trailing, we can express I's problem as

$$-\lambda_I(x_1 - x_I)^2 - 4\alpha\psi_x x_1\pi\Big(\lambda_I(x_I - x_C)^2 + (1-\lambda_I)(z_I - z_C)^2\Big), \tag{19}$$

which yields the following FONC (which is also sufficient):

$$-2\lambda_I(x_1 - x_I) - 4\alpha\psi_x\pi\left(\lambda_I(x_I - x_C)^2 + (1-\lambda_I)(z_I - z_C)^2\right) = 0,$$

which rearranged yields:

$$x_1 = x_I + \frac{4\alpha\psi_x\pi}{\lambda_I}\left(\lambda_I(x_I - x_C)^2 + (1-\lambda_I)(z_I - z_C)^2\right).$$

It follows that

$$x_1 = \min\left\{x_I + \frac{4\alpha\psi_x\pi}{\lambda_I}\left(\lambda_I(x_I - x_C)^2 + (1-\lambda_I)(z_I - z_C)^2\right), \frac{1}{4\alpha\lambda_v\psi_x}\right\}. \tag{20}$$

□

Proof of Proposition 2. First, suppose $\pi > \frac{1}{2}$. Recall that this implies that the incumbent is always reelected if the voter receives no new information. Further, Assumption 1 implies that if she observes an informative outcome on X, then the outcome on Z is electorally irrelevant. Suppose that the voter observes an uninformative outcome on X, and an informative outcome on Z. If $z_v = \alpha$, I is always re-elected with $\pi > \frac{1}{2}$. Consider now the case in which the voter observes an uninformative outcome on X, but learns that $z_v = -\alpha$. Denote $\mu^x(\varnothing, -\alpha, \rho)$ the voter's posterior that the state of the world on dimension X is positive in this case, i.e, if she observes an uninformative outcome on dimension X but learns that the state on dimension Z is $-\alpha$. Then, we must consider two cases. If the prior π is sufficiently high relative to the correlation ρ so that $\mu^x(\varnothing, -\alpha, \rho) > \widehat{\mu_v^x}(0)$, then the incumbent is reelected. In this case, the incumbent's retention chances are not a function of the policy on dimension Z, therefore he is indifferent between acting and not acting and by assumption chooses not to. If instead the prior π is sufficiently low relative to the correlation ρ so that $\mu^x(\varnothing, -\alpha, \rho) < \widehat{\mu_v^x}(0)$, the incumbent's retention chances are hurt by information on Z, and he chooses not to act on this secondary dimension.

Next, suppose $\pi < \frac{1}{2}$. Then, the incumbent is always ousted if the voter receives no new information. As above, if the voter observes an informative outcome on X, the outcome on Z is electorally irrelevant. Similarly, if the voter observes an uninformative outcome on X, and an informative outcome on Z such that $z_v = -\alpha$, I is always ousted with $\pi < \frac{1}{2}$. Suppose instead that the voter observes an uninformative outcome on X, but learns that $z_v = \alpha$. Again, we must consider two cases. If the correlation ρ is low, so that that $\mu^x(\varnothing, \alpha, \rho) < \widehat{\mu_v^x}(1)$, then the incumbent is ousted. Under this condition, the incumbent's ex-ante retention chances are not a function of the policy on dimension Z, therefore he is indifferent between acting and not acting and by assumption chooses not to. If instead the correlation ρ is sufficiently high that $\mu^x(\varnothing, \alpha, \rho) > \widehat{\mu^x_v}(1)$, generating an informative outcome on Z can only help the incumbent's retention chances. In other words, the incumbent's ex-ante retention chances increase as z_1 moves away from 0. Thus, he always chooses to act on Z.

Therefore, $\widehat{\rho}_T$ satisfies:

$$\mu^x(\varnothing, \alpha, \rho) = \widehat{\mu^x_v}(1), \tag{21}$$

where

$$\mu^x(\varnothing, \alpha, \rho) = \frac{\pi\rho}{\pi\rho + (1-\pi)(1-\rho)}. \tag{22}$$

Combining the above, we have

$$\widehat{\rho}_T = \frac{(1-\pi)\widehat{\mu}^x_v(1)}{\pi(1 - 2\widehat{\mu}^x_v(1)) + \widehat{\mu}^x_v(1)} \tag{23}$$

□

Proof of Corollary 1. Recall that under $\lambda_I = 1$ the incumbent's utility depends on z_1 only via the voter learning. Further, if the incumbent chooses to act on Z in equilibrium it must be the case that his probability of winning is increasing in the probability of generating an informative outcome on Z. This yields that in equilibrium the incumbent will always choose to implement a fully informative policy $z_1^* > z'$. □

Proof of Proposition 3. Consider the incumbent's choice on X. When I is trailing and $\rho > \widehat{\rho}_T$, we have $\mathbb{P} = 4\alpha\psi_x x_1 \pi + (1 - 4\alpha\psi_x x_1)4\alpha\psi_z z_1 \beta$. Plugging in $z_1^* = \frac{1}{4\alpha\psi_z(1-\lambda_v)}$, the trailing incumbent's retention probability reduces to

$$4\alpha\psi_x \pi x_1 + (1 - 4\alpha\psi_x x_1)\beta. \tag{24}$$

Note that, given $\pi < \frac{1}{2}$, $\beta = \pi\rho + (1-\pi)(1-\rho) > \pi$, therefore the incumbent's probability of winning is decreasing in x_1. It follows from Equation 7 that $x_1^* < x_I$. $\square$

Proof of Corollary 2. Applying the implicit function theorem, we have that

$$\frac{\partial x_1^*}{\partial \rho} = -\frac{\frac{\partial FOC}{\partial \rho}}{\frac{\partial FOC}{\partial x_1}}. \tag{25}$$

In equilibrium, we have that $\frac{\partial FOC}{\partial x_1} < 0$, therefore $\frac{\partial x_1^*}{\partial \rho} > 0$ if and only if $\frac{\partial FOC}{\partial \rho} > 0$:

$$\frac{\partial^2 \mathbb{P}(x_1, z_1)}{\partial x_1 \partial \rho} \left(\lambda_I 4x_I^2 + (1-\lambda_I)4z_I^2\right) > 0. \tag{26}$$

Recall that $\mathbb{P}(x_1, z_1) = 4\alpha\psi_x x_1 \lambda_v \pi + (1 - 4\alpha\psi_x x_1 \lambda_v)4\alpha\psi_z z_1(1-\lambda_v)\beta$, therefore the above reduces to

$$-16\alpha^2\psi_x\psi_z\lambda_v(1-\lambda_v)\frac{\partial \beta}{\partial \rho} > 0, \tag{27}$$

which is always true when $\pi < \frac{1}{2}$. $\square$

Proof of Proposition 4. Recall that $\widehat{\mu_v^x}(\mu^z)$ defines the value of μ_v^x such that the voter is indifferent between replacing and keeping the incumbent, for a given μ^z. We must consider four cases, which differ in whether the voter posterior upon observing an uninformative outcome on X and for a given value of ρ is above or below the retention threshold $\widehat{\mu_v^x}(\mu^z)$, given the outcome on the secondary dimension Z. The first two cases correspond to an ex-ante leading incumbent, the last two ones to a trailing one:

1. $\pi = \mu_v^x(\varnothing, \varnothing, \rho) > \mu_v^x(\varnothing, -\alpha, \rho) > \widehat{\mu_v^x}(0) > \frac{1}{2}$
2. $\pi = \mu_v^x(\varnothing, \varnothing, \rho) > \widehat{\mu_v^x}(0) > \mu_v^x(\varnothing, -\alpha, \rho)$ (which implies $\pi > \frac{1}{2}$)
3. $\pi = \mu_v^x(\varnothing, \varnothing, \rho) < \widehat{\mu_v^x}(1) < \mu_v^x(\varnothing, \alpha, \rho)$ (which implies $\pi < \frac{1}{2}$)
4. $\pi = \mu_v^x(\varnothing, \varnothing, \rho) < \mu_v^x(\varnothing, \alpha, \rho) < \widehat{\mu_v^x}(1) < \frac{1}{2}$.

Case 1: $\pi > \mu_v^x(\varnothing, -\alpha, \rho) > \widehat{\mu_v^x}(0) > \frac{1}{2}$.

First, suppose that the incumbent is leading and that the voter re-elects the incumbent even if she knew with certainty to be aligned with the challenger on dimension Z, i.e.:

$$\mu_v^x(\varnothing, -\alpha, \rho) > \widehat{\mu^x}_v(0). \tag{28}$$

Substituting $\mu_v^x(\varnothing, -\alpha, \rho) = \frac{\pi(1-\rho)}{\pi(1-\rho)+(1-\pi)\rho}$ yields the following condition on the correlation coefficient ρ:

$$\rho < \frac{\pi(1 - \widehat{\mu}_v^x(0))}{\pi + \widehat{\mu}_v^x(0)(1 - 2\pi)}, \tag{29}$$

where notice that the denominator $\pi + \widehat{\mu}_v^x(0)(1-2\pi) > 0$ (the RHS is linear in π, and the condition is always satisfied at $\pi = 0$ and $\pi = 1$). The condition implies that the incumbent is always reelected unless the voter observes an informative outcome on X and learns that $x_v = -\alpha$. Then, it is easy to see that the incumbent's probability of winning is not a function of the outcome on Z, therefore not a function of his policy choice z_1. This also implies that the incumbent's maximization problem reduces to the unidimensional one, and the incumbent sets his ideologically preferred policy on dimension Z, $z_1^* = z_I$.

Case 2: $\pi = \mu_v^x(\varnothing, \varnothing, \rho) > \widehat{\mu_v^x}(0) > \mu_v^x(\varnothing, -\alpha, \rho) > \frac{1}{2}$.

Suppose (as in the first case) that I is leading, and that—differently from the first case—the voter replaces the incumbent when she learns to be aligned with the challenger on dimension Z:

$$\widehat{\mu^x}_v(0) > \mu_v^x(\varnothing, -\alpha, \rho).$$

Substituting $\mu_v^x(\varnothing, -\alpha, \rho) = \frac{\pi(1-\rho)}{\pi(1-\rho)+(1-\pi)\rho}$ produces:

$$\rho > \frac{\pi(1 - \widehat{\mu}_v^x(0))}{\pi + \widehat{\mu}_v^x(0)(1 - 2\pi)}. \tag{30}$$

Straightforwardly, the incumbent chooses to open the secondary dimension if and only if his utility is increasing in z_1 at $z_1 = 0$.

Under the assumption on ρ, I's retention probability is given by:

$$\begin{aligned}\mathbb{P}(x_1, z_1) &= 1 - (1-\pi)\Pr(L=1|\pi, x_1) - (1-\Pr(L=1|\pi,x_1))(1-\beta)\Pr(L=1|\beta, z_1)\\ &= 1-(1-\pi)4\alpha x_1\lambda_v\psi_x - (1-4\alpha x_1\lambda_v\psi_x)(1-\beta)4\alpha z_1(1-\lambda_v)\psi_z\end{aligned}$$

Denote $K = 4\lambda_I x_I^2 + 4(1-\lambda_I)z_I^2$. Plugging the value of $\mathbb{P}(x_1, z_1)$ into I's objective and differentiating with respect to z_1, we get that I opens Z if and only if:

$$2(1-\lambda_I)z_I - (1-\beta)(1-4\alpha\widehat{x}\lambda_v\psi_x)4\alpha\psi_z(1-\lambda_v)\left[\lambda_I 4x_I^2 + (1-\lambda_I)4z_I^2\right] > 0, \tag{31}$$

where $\widehat{x}$ solves

$$-2\lambda_I(x_1 - x_I) - 4\alpha\psi_x\lambda_v\Big[1-\pi-4\alpha\psi_z(1-\lambda_v)z_1(1-\beta)\Big]K = 0. \tag{32}$$

and is equal to:

$$\widehat{x} = x_I - \frac{4\alpha\psi_x\lambda_v(1-\pi)\left[\lambda_I 4x_I^2 + (1-\lambda_I)4z_I^2\right]}{\lambda_I}. \tag{33}$$

Condition 31 is satisfied for $\lambda_I < \hat{\lambda}_I$. The expression for $\hat{\lambda}_I$ is lengthy therefore omitted. Intuitively, the incumbent opens the secondary dimension when he sufficiently cares about it.

Case 3: $\pi = \mu_v^x(\varnothing, \varnothing, \rho) < \widehat{\mu_v^x}(1) < \mu_v^x(\varnothing, \alpha, \rho) < \frac{1}{2}$**.**

Suppose now that the incumbent is trailing and

$$\widehat{\mu_v^x}(1) < \mu_v^x(\varnothing, \alpha, \rho), \tag{34}$$

which, plugging in $\mu_v^x(\varnothing, \alpha, \rho) = \frac{\pi\rho}{\pi\rho + (1-\pi)(1-\rho)}$, reduces to

$$\rho > \frac{\widehat{\mu}^x(1)(1-\pi)}{\widehat{\mu}^x(1)(1-2\pi)+\pi}. \tag{35}$$

Recall that $\pi < \frac{1}{2}$ implies the incumbent is always ousted if the voter learns nothing new. This also implies that the incumbent is ousted if the voter observes an uninformative outcome on X and learns $z_v = -\alpha$. Finally, $\pi = \mu_v^x(\varnothing, \varnothing, \rho) < \widehat{\mu_v^x}(1) < \mu_v^x(\varnothing, \alpha, \rho)$ implies that the incumbent is re-elected if the voter observes an uninformative outcome on X but learns $z_v = \alpha$.

Thus we have that in this case:

$$\mathbb{P} = 4\alpha\psi_x\lambda_v x_1 \pi + (1 - 4\alpha\psi_x\lambda_v x_1)4\alpha\psi_z(1-\lambda_v)z_1\beta. \tag{36}$$

Thus, the FOCs are

$$(x_1) - 2\lambda_I(x_1 - x_I) + 4\alpha\psi_x\lambda_V\Big[\pi - 4\alpha\psi_z(1-\lambda_V)z_1\beta\Big]\Big[4\lambda_I x_I^2 + 4(1-\lambda_I)z_I^2\Big] = 0 \tag{37}$$

$$(z_1) \qquad - 2(1-\lambda_I)(z_1 - z_I) + (1 - 4\alpha\psi_x\lambda_V x_1)4\alpha\psi_z(1-\lambda_V)\beta\Big[4\lambda_I x_I^2 + 4(1-\lambda_I)z_I^2\Big] = 0 \tag{38}$$

Recalling that in equilibrium $1 - 4\alpha\psi_x\lambda_I x_1 \geq 0$, we notice that the incumbent's utility is always increasing in z_1 at $z_1 = 0$. Thus, the incumbent z_1 is always strictly larger than 0 in equilibrium.

Case 4: $\mu_v^x(\varnothing, \varnothing, \rho) = \pi < \mu_v^x(\varnothing, \alpha, \rho) < \widehat{\mu_v^x}(1) < \frac{1}{2}$.

Lastly, suppose that the incumbent is trailing and that the voter ousts the incumbent even if she knew with certainty to be aligned with him on dimension Z, i.e.:

$$\mu_v^x(\varnothing, \alpha, \rho) < \widehat{\mu_v^x}(1). \tag{39}$$

The condition implies that the incumbent is always ousted unless the voter observes an informative outcome on X and learns $x_v = \alpha$. Then, analogously to case 1, the incumbent's probability of winning is not a function of the outcome on Z, therefore not a function of his policy choice z_1. Thus, in equilibrium the incumbent always sets $z_1^* = z_I > 0$. □

Proof of Proposition 5. From the proof of Proposition 4, we know that in Cases 1 and 4 $x_1^* = x_u$ and $z_1^* = z_I = z_u$.

Consider instead Case 2, i.e., a leading incumbent ($\pi > \frac{1}{2}$) under a high ρ. First, consider the equilibrium choice on Z. The incumbent's utility is always decreasing in z_1 at $z_1 \geq z_I$:

$$-2(1-\lambda_I)(z_1 - z_I) - (1-\beta)(1-4\alpha x_1 \lambda_v \psi_x)4\alpha\psi_z(1-\lambda_v)K \tag{40}$$

Therefore, in equilibrium it must be the case that $z_1 < z_I$. Next, consider the incumbent's choice on X. First, suppose that the problem is concave and the FOC is sufficient to identify the equilibrium policy. The result follows from inspection of 32. Suppose instead that the problem is not concave (or it is, but the equilibrium policy is at a corner). Then, depending on parameters, the equilibrium policy can take one of three values: $\{0, \hat{x}, x'\}$, where $\hat{x}$ is the interior solution. Then, to conclude the proof for Case 2 is sufficient to show that $x_1^* = 0 \implies x_u = 0$. This follows from the fact that if the incumbent's utility is decreasing in x_1 at $x_1 = 0$ under $\lambda_v < 1$, then it must also be decreasing under $\lambda_v = 1$:

$$2\lambda_I x_I - 4\alpha\psi_x(1-\pi)K \leq 2\lambda_I x_I - 4\alpha\psi_x\lambda_v\Big[1 - \pi - 4\alpha\psi_z(1-\lambda_v)z_1(1-\beta)\Big]K, \tag{41}$$

which reduces to

$$1 - \pi \geq \lambda_v\Big[1 - \pi - 4\alpha\psi_z(1-\lambda_v)z_1(1-\beta)\Big], \tag{42}$$

which is always satisfied.

Finally, consider Case 3, i.e., a trailing incumbent ($\pi < \frac{1}{2}$) under a high ρ. We proceed as above. Focus first on the equilibrium choice on Z. The incumbent's utility is always increasing in z_1 at $z_1 \leq z_I$ (follows from inspection of 38) therefore in equilibrium it must be the case that

$z_1 > z_I$. Next, consider the incumbent's choice on X. First, suppose that the problem is concave and the FOC is sufficient to identify the equilibrium policy. The result follows from inspection of 37. Suppose instead that the problem is not concave (or it is, but the equilibrium policy is at a corner). Then, depending on parameters, the equilibrium policy can take one of three values: $\{0, \hat{x}, x'\}$, where $\hat{x}$ is the interior solution. Then, to conclude the proof for Case 3 is sufficient to show that $x_1^* = x' \implies x_u = x'$. This follows from the fact that if the incumbent's utility is increasing in x_1 at $x_1 = x'$ under $\lambda_v < 1$, then it must also be increasing under $\lambda_v = 1$:

$$-2\lambda_I(x' - x_I) + 4\alpha\psi_x\pi K \geq -2\lambda_I(x' - x_I) + 4\alpha\psi_x\lambda_v\Big[\pi - 4\alpha\psi_z(1-\lambda_v)z_1\beta\Big]K \tag{43}$$

which reduces to

$$\pi \geq \lambda_v\Big[\pi - 4\alpha\psi_z(1-\lambda_v)z_1\beta\Big], \tag{44}$$

which is always satisfied. □

Proof of Proposition 6. Suppose that the incumbent has multiple secondary dimensions $\tilde{D}$ available to open, but can only choose one. Applying the envelope theorem, we can characterize how the incumbent's equilibrium utility changes if he chooses to open dimensions with different features in the first period. For simplicity, we will assume that in the second period the officeholder implements his ideologically preferred policy on all dimensions, and denote $\tilde{K}$ the cost of losing the election in this augmented multidimensional world. Further, we denote $\tilde{d}_I$ the incumbent's ideal point on dimension $\tilde{d}$, $\rho_{\tilde{d}}$ the correlation between X and $\tilde{D}$, and $\psi_{\tilde{d}}$ the precision of the shock term on dimension $\tilde{D}$. Then, we have

$$\frac{\partial U_I^*}{\partial \tilde{d}_I} = 2(d_1 - \tilde{d}_I). \tag{45}$$

From Proposition 5 we know that $d_1 \geq \tilde{d}_I$ iff $\pi < \frac{1}{2}$. Therefore $\frac{\partial U_I^*}{\partial \tilde{d}_I} \geq 0$ iff $\pi > \frac{1}{2}$. As an aside, note that here we are not treating $\tilde{K}$ as a function of $\tilde{d}_I$, since we are comparing utility across dimensions and the cost of losing does not depend on which dimension the incumbent chooses to open in the first period.

Looking at the correlation $\rho_{\tilde{d}}$, we have

$$\frac{\partial U_I^*}{\partial \rho_{\tilde{d}}} = \tilde{K}(1 - 4\alpha\psi_x x_1 \lambda_v) 4\alpha\psi_{\tilde{d}} \tilde{d}_1 (1 - \lambda_v) \frac{\partial \beta}{\partial \rho_{\tilde{d}}} > 0 \tag{46}$$

for a leading incumbent (β is increasing in $\rho_{\tilde{d}}$ under $\pi > \frac{1}{2}$), and

$$\frac{\partial U_I^*}{\partial \rho_{\tilde{d}}} = \tilde{K}(1 - 4\alpha\psi_x x_1 \lambda_v) 4\alpha\psi_{\tilde{d}} \tilde{d}_1 (1 - \lambda_v) \frac{\partial \beta}{\partial \rho_{\tilde{d}}} < 0 \tag{47}$$

for a trailing one (β is decreasing in $\rho_{\tilde{d}}$ under $\pi < \frac{1}{2}$).

Finally, consider the precision of the outcomes:

$$\frac{\partial U_I^*}{\partial \psi_{\tilde{d}}} = -(1 - 4\alpha\psi_x \lambda_v x_1)(1 - \beta) 4\alpha \tilde{d}_1 (1 - \lambda_v) < 0 \tag{48}$$

for a leading incumbent, and

$$\frac{\partial U_I^*}{\partial \psi_{\tilde{d}}} = (1 - 4\alpha\psi_x \lambda_v x_1)\beta 4\alpha \tilde{d}_1 (1 - \lambda_v) > 0 \tag{49}$$

for a trailing one. □