Learning in a Complex World: How Multidimensionality Affects Policymaking

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

April 19, 2024

Abstract

Contemporary governments face a multitude of policy challenges. Choosing which issues to prioritize or leave unaddressed, and how to align their programs on different (and often connected) dimensions, present critical challenges for reelection-seeking policymakers. We develop an accountability model to study these decisions. We find that trailing incumbents embrace comprehensive policy programs, while leading ones prioritize fewer, independent policies. Our central result identifies a substitution effect among correlated dimensions: policymakers shift towards extreme or moderate positions when expanding their policy agenda, depending on their electoral prospects. This result generates novel empirical implications about the relationship between issue salience, policy extremism, and the incumbent's electoral prospects, for which we provide a preliminary empirical test using data from the US.

*UCSD, *Email*: fizzo@ucsd.edu

†Duke University, *Email*: giovanna.invernizzi@duke.edu

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

1 Introduction

Contemporary governments are confronted with a wide range of policy issues, from international affairs, to social security, to the economy. Two crucial decisions bear significant strategic importance in this context: First, determining which issues to prioritize and which to leave unaddressed, if any. Second, establishing the direction of policymaking on the various issues, and determining which positions should "go together" (Bawn et al., 2012). For policymakers, both sets of decisions represent strategic dilemmas with significant implications for electoral success.

Historically, policymakers have responded differently to this dilemma, adopting narrower or broader policy programs under different circumstances. Following a rocky start to his presidency, for example, President Harry Truman aggressively expanded his policymaking efforts to multiple issue areas in 1947, introducing a series of bills known as the "Fair Deal." In stark contrast, after cruising to victory in 1952 President Dwight Eisenhower focused his attention on the economy and spending, eschewing broader policy interventions on other issues.

Moreover, when policymaking encompasses multiple dimensions, extreme measures on some dimensions are often combined with more moderate policies ones on other issues. Again, depending on the features of the electoral environment, we see variation in what types of issues are associated with more extreme or more moderate policies. Looking outside of the United States, French President Sarkozy adopted a moderate stance on economic issues,[1] while taking extreme positions in other policy areas. In particular, his bold positions on immigration and social issues, and his risky choices on foreign policy, have widely been considered as a gamble amid a personal political crisis, with his approval ratings plummeting in the period preceding his reelection campaign.[2] In comparison, German Chancellor Angela Merkel, who enjoyed a consistent electoral lead, implemented austere economic measures but maintained a more moderate approach on immigration (De

[1] Sarkozy "has trumpeted the return of government intervention and soft-pedalled on sensitive economic and labour reforms" (Hall, 2022): https://www.ft.com/content/a4151446-5e8f-11de-91ad-00144feabdc0.

[2] Prime examples are the so-called Burqa-ban, and Sarkozy's interventionist approach in response to the Arab Spring. According to Politico, Sarkozy's approval ratings in March 2011 reached 31%, the lowest of any president in modern French history: https://www.politico.com/story/2011/03/sarkozys-war-051708.

La Baume, 2017) and climate change (Schwagerl, 2011).[3]

In this paper we propose a game-theoretic framework to examine the decision of policymakers within a multidimensional landscape, with the aim of shedding light on this variation and how it may depend on features of the electoral environment. We address several key questions: What drives policymakers to adopt either a broad or narrow policy focus? Which policy areas are likely to be neglected under a more focused agenda? What types of reforms do policymakers pursue when they have a broader versus narrower policy agenda?

In answering these questions, we begin from the premise that policymaking is complicated, leaving voters unsure about the expected consequences of various policies for their welfare. As we elaborate further below, in this uncertain world voters react to the results of policy choices—not simply the substance of the policies themselves— by assessing their personal well-being (Fiorina, 1981; Stimson, 2018). If a policy choice leads to an outcome that voters like, then their evaluation of the policy improves. Since the inferences voters draw when observing outcomes depend on the exact policies implemented (as in Izzo, 2023), incumbent officeholders must then consider how their policy choices today influence their retention chances tomorrow.

In the model, when a policy remains at the status quo it does not provide any new information about that particular dimension. Instead, new policies offer opportunities for voters to learn, and more extreme reforms yield more informative outcomes. Consider a scenario where a voter experiences a favorable outcome from an extremely leftist reform. In such a case, it becomes evident that this policy aligns well with the voter's interests. Outcomes of moderate policies are however less informative. Even if the policy moves slightly in a direction that is not ideal for voters, random chance may still allow them to experience relatively high welfare.

Importantly, this learning process becomes more complex in a multidimensional world, where policy issues may be interconnected and learning spillovers may arise. For example, if a voter's experience suggests that liberal economic reforms are optimal for her, she may also become more inclined to believe that liberal reforms on other issues, such as healthcare, are beneficial. In other

[3] http://www.politico.eu/article/angela-merkel-defends-open-border-migration-refugee-policy-germany/ and https://e360.yale.edu/features/germanys-unlikely-champion-of-a-radical-green-energy-path.

words, voters can learn about how well a political program fits their preferences on one dimension by observing the policy outcome on another dimension.[4]

In such a complex world, wherein voters care about multiple dimensions that they perceive as connected, voter learning has a nuanced effect on policymaking choices by politicians. In this context, an 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 correlated dimensions that generate learning spillovers). Consequently, the incumbent faces a complex web of incentives as he considers the impact of policy choices on voters' well-being and, in turn, his own chances of reelection. These incentives also interact with the incumbent's own ideological preferences, determining his optimal choice with regards to which dimensions to address, and which policies to pursue.

To better elucidate these dynamics, we begin by analyzing a benchmark case where the voter only cares about a single dimension, say the economy, and she believes this issue to be unrelated to any other, so that learning spillovers are not possible. In this case, even if the incumbent cares about multiple dimensions, his strategic problem is unidimensional: his electoral chances are a function solely of his policy choice on the economic dimension. The results of this benchmark align with the classic intuition on policy gambles. An incumbent who is ex-ante leading always faces incentives to *prevent* voter learning, and in equilibrium he pursues an economic policy that is *more moderate* than his ideal one. This follows from the result that more extreme policies generate more information and thus higher electoral risk, whereas more moderate policies closer to the status quo hinder voter learning and thus are more likely to preserve the incumbent's initial advantage. The opposite holds for a trailing incumbent: such an incumbent has incentives to *gamble* and thus implements extreme policies that *facilitate* voter learning.

A similar logic intuitively explains under which conditions the incumbent has incentives to expand or contract the scope of policymaking in a multidimensional world (i.e., a world wherein the voter cares about multiple dimensions). A trailing incumbent is motivated to promote voter learning. As such, he has incentives to reform policy on all available dimensions, including the ones

[4]Section 3.1 discusses this assumption in more depth.

he has no ideological preferences over. In contrast, a leading incumbent wishes to avoid generating fresh policy information that could potentially undermine his advantage. Therefore, this incumbent has no incentives to pursue a broader policy agenda than his own ideological preferences dictate. Indeed, we show that under some conditions a leading incumbent leaves policy unchanged on some issues, even though statically (i.e., following his ideologial preferences) he would prefer to pursue reforms on all dimensions.

Having characterized conditions under which equilibrium policy is multidimensional, our central result then shows that this multidimensionality can influence the *nature* of policymaking in unexpected ways. Specifically, we uncover a *strategic substitution effect* between dimensions that emerges as a consequence of learning spillovers when policy dimensions are sufficiently correlated. This result also produces novel empirical implications which we further discuss below.

To understand the substitution effect, consider first a leading incumbent. As described above, this incumbent wants to prevent voter learning. In a unidimensional context, he consistently adopts moderate policies. In a multidimensional context, he maximizes his reelection chances by avoiding reforms on secondary issues to prevent harmful learning spillovers. However, if his ideological preferences on these secondary issues are strong enough, the incumbent will still choose to expand his policy agenda to include them. In this case, one might think that the goal of preventing voter learning would drive this incumbent to pursue moderate policies across *all* available issue domains. Instead, we show that expanding the scope of his policy agenda can actually lead the incumbent to implement extreme policies on the primary dimension. This happens because by adopting more extreme–and consequently more informative–policies on the primary issue, the leading incumbent can counteract the adverse effects of learning spillovers from the secondary dimensions.

Symmetric results hold when the incumbent is electorally disadvantaged. In contrast with intuition from the unidimensional case, a trailing incumbent may adopt more *moderate* stances on the primary dimension when he expands his policy initiatives to include other issues. This happens because implementing moderate policies on the primary dimension generates less information about this issue, thereby amplifying the impact of learning spillovers from extreme policies on secondary dimensions. In turn, this is strategically valuable if favorable outcomes are more likely to arise on secondary issues.

Our findings hold substantial implications for our understanding of policymaking. Both theoretically and empirically, it is common for scholars to assume that a one-dimensional world closely approximates the multidimensional reality. This assumption is rooted in the observation that preferences across various issues are often correlated (Converse, 1964; McMurray, 2014). However, our framework suggests that it is precisely this correlation that adds complexity to the scenario.

In situations where issues are orthogonal to each other, the existence of multiple dimensions need not distort the fundamental nature of policymakers' strategic challenges when addressing each issue individually. Consequently, a unidimensional model can serve as a suitable approximation of a multidimensional world. Yet, when issues are highly correlated (even if just in voters' minds), policymaking can take on a significantly distinct character due to the substitution effect described above. To truly understand policymaking, one must account for its inherent multidimensionality.

We conclude the paper by discussing in more detail our theory's empirical implications. Our central substitution-effect result points to a subtle relationship between the salience of different issues for voters, and the extremism of policies pursued by the office holder. In particular, this relationship is mediated by the incumbent's electoral prospects. A leading incumbent tends to pursue more extreme policies on primary, more salient, issues. Vice versa, for trailing incumbents our theory predicts a negative association between policy extremism and issue salience. A preliminary "reality check" using data from the US reveals results that, while noisy, are consistent with our directional predictions. While further analysis is obviously needed, we hope that our exercise can provide a useful road-map for empirical research.

2 Related Literature

Our paper contributes to the literature on multidimensional policymaking. Within the electoral accountability literature, Banks and Duggan (2008) are amongst the first to consider a multidimensional policy space. Other works study the incumbent's decision over how to allocate his budget between different tasks (e.g., Ashworth (2005); Ashworth and Bueno de Mesquita (2006); Ash, Morelli and Van Weelden (2017)). This literature, however, considers how politicians signal competence or their ideological preferences. This complements our approach, where we assume that the voter faces uncertainty about her own optimal policy program and learns by experience.

More recent related work is Buisseret and Van Weelden (2022). The paper 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 in a world where voters themselves are uncertain about the optimal policy. Thus, the incentives underlying policymakers' strategic choices in our setting pertain to manipulating voter preferences, rather than revealing information to the parties as in Buisseret and Van Weelden (2022).

Finally, our paper connects with the literature on policy experimentation and multi-armed bandit problems (e.g., Strumpf, 2002; Volden, Ting and Carpenter, 2008; Strulovici, 2010; Hirsch, 2016; Dewan and Hortala-Vallve, 2019; Gieczewski and Kosterina, 2020). Beyond the different substantive focus and research questions (our paper is the only one in this tradition to study multi-dimensional experimentation in an electoral accountability framework), most of this literature considers a binary policy space, with one risky option and one safe option.[5] As such, these works can only analyze a decision-maker's choice to experiment or not. Instead, we consider policy experimentation with a continuous space. Doing so allows us to analyze the intensity of the policymaker's dynamic incentives to take risks and study the equilibrium amount of policy experimentation. This is important because a binary policy choice may obfuscate much of the effect of multidimensionality on policymaking.

Callander (2011) and Callander and Hummel (2014) also study experimentation with a continuous of policies. However, Callander chooses to abstract from dynamic electoral considerations, by assuming either myopic players (Callander, 2011) or exogenous re-election probabilities (Callander

[5]Strumpf (2002) considers an extension with two experimental policies. Hirsch (2016) considers a binary policy space where one option is not inherently more risky than the other, but a correct policy succeeds only if a bureaucrat exerts sufficient effort in its implementation.

and Hummel, 2014). Instead, the focus of this paper is precisely on the incumbent's dynamic incentives to control information. Furthermore, Callander's framework considers a world in which voters know whether right-wing or left-wing policies tend to generate better outcomes, but experiment to learn about the exact consequences of each policy program. In this paper, we build on a different framework to think about policy experimentation, in which the nature of uncertainty is reversed. Voters aim to learn whether liberal or conservative reforms are optimal in expectation, even though the exact consequences of each policy are somewhat unpredictable. This allows us to think about policy experimentation in connection to ideology, and generates the result that extreme policies, rather than small incremental changes as in Callander, produce more information.

Our learning technology relates more closely to the models introduced in Ashworth, Bueno de Mesquita and Friedenberg (2017) and Izzo (Forthcoming), which also share our assumption that voters must learn about the expected consequences of the various policy choices. However, both papers consider a unidimensional policy space. This contrasts with our focus on multidimensional problems.

3 The Model

Players and actions. Our model has three players: an incumbent, I, a challenger, C, and a voter, V. In each of the two periods in the model, the incumbent chooses whether to act on one or both of two policy dimensions, $D \in \{X, Z\}$.[6] 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. To avoid trivialities, we assume that if the officeholder is indifferent between acting or not on dimension D, then he chooses not to act.

[6]Considering a single voter allows us to streamline analysis and notation. However, even in our multidimensional setting, the qualitative insights of our theory would be similar in a world with multiple voters as long as these voters did not interpret new information in opposite ways. More specifically, our framework assumes that the same piece of evidence does not lead some voters to update, say, in favor of the incumbent while inducing others to update against him. Then, the incumbent's strategic incentives to control information would be qualitatively similar to the ones emerging here.

Information. Both the incumbent and the challenger's ideal points are common knowledge and, for simplicity, 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, the voter's optimal policy on each dimension, denoted by d_v, can take one of two values: $d_v \in \{-\alpha, \alpha\}$. Players share common prior beliefs that

$$\Pr(x_v = \alpha) = \pi_x \text{ and } \Pr(x_v = -\alpha) = 1 - \pi_x, \tag{1}$$

and

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

Thus, players believe the dimensions are positively correlated in a symmetric way. As the voter's initial beliefs on dimension X shift to the right, so does her prior on dimension Z. It follows that the ex-ante probability that $z_v = \alpha$ is

$$\Pr(z_v = \alpha) = \pi_z = \rho\pi_x + (1-\rho)(1-\pi_x). \tag{3}$$

Notice that in our setting players initially have more information about the voter's ideal policy on dimension X than on dimension Z, i.e., π_z is always closer to $\frac{1}{2}$ than π_x is. 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 payoff in period t is

$$u_{i,t} = -\lambda_i^x (x_t - x_i)^2 - \lambda_i^z (z_t - z_i)^2, \tag{4}$$

where $\lambda_i^d \geq 0$ captures how much i cares about dimension d, and d_i is i's ideal point.

On each dimension, the voter observes her true payoff plus a random shock, $\varepsilon_{d,t}$. The random shock is drawn in each period and for each dimension from a uniform distribution with support $\left[-\frac{1}{2\psi_d}, \frac{1}{2\psi_d}\right]$. The assumption that the noise $\varepsilon_{d,t}$ is uniformly distributed substantially simplifies the analysis, but is not necessary for our results. We briefly return to this point in Section 8.3 below.

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.

2. V observes I's choice and her realized payoff on each dimension.

3. V chooses whether to re-elect I or replace him with C.

4. The winner of the election takes office, then chooses whether to act on each dimension or instead keep the status quo from the previous period.[7]

The equilibrium concept is Perfect Bayesian Equilibrium. Before concluding this section, let us emphasize that in our setting there is no asymmetry of information between the voter and the politicians. In particular, the incumbent has no private information about his own ideological preferences. This allows us to assume away the possibility that the incumbent's policy choice directly provides electorally-relevant information to the voter and instead focus on what the voter learns from her lived experiences (i.e., the inferences she draws upon observing realized outcomes). We discuss how our results are robust to relaxing this assumption in Section 8.3.

In order to isolate the electoral incentives to expand the scope of policy-making, we begin by assuming that politicians do not intrinsically care about the secondary dimension Z:

Assumption 1. $\lambda_I^z = 0$ *and* $\lambda_C^z = 0$.

We relax this assumption in Section 9 where we consider politicians who have ideological preferences on both dimensions.

[7]Notice that, as it is common in electoral accountability models, the challenger plays no active role in the first period. While this is obviously a simplification, we believe it is relatively unproblematic in our setting. We model how voters learn by experiencing the results of different policies, and how this in turns influences the types of policies that are implemented in equilibrium. Since challengers have little policymaking authority, it seems reasonable to focus on the incumbent's strategic calculus.

3.1 Discussion of the Assumptions

Before concluding this section, we discuss in more details two ingredients that lie at the core of our theory.

Voter Uncertainty and Retrospective Voting. As described above, our model builds on the observation that voters often face substantial uncertainty about the possible consequences and relative virtue of the policy programs espoused by opposing parties (Callander, 2011; Tavits, 2007). In other words, the players are faced with what Tavits (2007) defines as pragmatic policy issues, and are unsure of 'what types of policies are related to what sorts of outcomes' (ibid: 155). For example, high taxation may be good for the voter, as it improves the provision of public goods, or bad for her, if it hampers economic growth. In our framework, voters respond to these informational challenges by looking at their personal well-being (Stimson, 2018). That is, voters look back at the incumbent's actions and how these impact observed outcomes. If the incumbent's past policy choices produced favorable outcomes, voters' evaluation and thus propensity to reelect him improve. Conversely, negative outcomes damage the incumbent's electoral chances. In this perspective, our theory builds on the retrospective voting framework, and the research arguing that voters form (and change) their preferences on the basis of their experiences (Fiorina, 1981; Achen, 1992).

Correlations and Multidimensionality. We build on this literature, but also observe that multidimensionality in politics presents voters—and parties and candidates—with additional informational challenges. Even with available heuristics (Sniderman, Brody and Tetlock, 1991; Lau and Redlawsk, 2001) and policy feedback (Pierson, 1993; Mettler and Soss, 2004), voters face the problem of understanding new issue areas and how to orient themselves when navigating a multidimensional world (Izzo, Martin and Callander, 2023).

Here, our argument is that voters may view different issue areas as interconnected, and therefore use experiences on one dimension to inform their beliefs and preferences on the others. We might interpret this assumption in two ways.

First, (certain) citizens may realize that different policy issues are actually correlated. While Converse (1964) suggested the public lacks political sophistication, more recent research indicates a rise in ideological consistency among voters (Levendusky, 2010). Increasingly, citizens recognize the connections across policy dimensions, from economics, to social and environmental issues, to minority rights (Hare, 2022). Notably, this trend extends across the electorate, even though it is more pronounced among politically active voters (Zingher and Flynn, 2019; Hare, 2022). In this perspective, voters thus tend to adopt consistent beliefs systems (Jewitt and Goren, 2016), a view captured in our model by the correlation parameter ρ.

Second, voters may not be as sophisticated as this perspective assumes, but might nonetheless use similar heuristics and respond to positive experiences in one dimension by adopting a more favorable attitude towards the same ideology in other dimensions. This view is consistent for example with the theory of motivated reasoning (Taber and Lodge, 2006; Kunda, 1987), according to which citizens tend to evaluate evidence in a way that aligns with their preexisting attitudes. Within our framework, an initially undecided voter who becomes convinced of the merits of, say, liberal economic policies, responds by adopting consistent views on other dimensions as well.

Both interpretations are consistent with our theory. In the analysis that follows, we take the connections across issue areas as given,[8] and study how they influence the inferences guiding voters' electoral decisions—and, in turn, policymakers' strategic incentives in a multidimensional world.

4 Voter Learning

Before moving to equilibrium analysis, it is useful to provide a characterization of the learning technology we adopt in our model, which builds on Izzo (Forthcoming). In this framework, the voter gains insights into her preferred policies through her real-life experiences. However, these experiences only represent a noisy signal of the genuine alignment between the voter's interests and the implemented policy, and this complicates her inference problem. Additionally (differently

[8]Thus, we assume that voters and politicians share the same beliefs about these connections. However, this is not necessary for our mechanism. All we need is that the incumbent correctly anticipates how his policy choices will influence voter's (potentially subjective) beliefs.

from Izzo (Forthcoming)), in our setup voter learning is twofold: direct and indirect. That is, the voter's realized utility on a given dimension provides her with information on her optimal platform on that dimension (*direct learning*) *and* on the other (*indirect learning*). Suppose that the economy and healthcare policies are connected in the voter's mind. Then, if the voter experiences positive outcomes in response to the incumbent's economic policy, she will infer not only that she likes the incumbent's economic policies but his healthcare policies as well.

We begin by considering the direct channel and characterizing the voter's *interim* posterior beliefs on each dimension D, i.e., her beliefs as a function of her realized utility on the given dimension only. We denote this interim posterior as $\tilde{\mu}^d$.

Lemma 1. *Direct voter learning satisfies the following properties:*

(i) The interim posterior $\tilde{\mu}^d$ *takes one of three values,* $\tilde{\mu}^d \in \{0, \pi_d, 1\}$*;*

(ii) If the incumbent does not act on dimension D*, then* $\tilde{\mu}^d = \pi_d$*;*

(iii) If the incumbent acts on dimension D*, the probability that* $\tilde{\mu}^d \neq \pi_d$ *is* $\min\{4\alpha\psi_d\lambda_d|d_1|, 1\}$.

Lemma 1 shows that, upon observing outcomes on a certain dimension, the voter either learns everything or nothing about her optimal platform on that dimension. If the policy remains at the status quo, the voter never receives new (direct) information on that dimension. If instead a new policy is implemented, she is more likely to discover her ideal point as the implemented platform becomes more extreme. While the starkness of the learning process follows from the assumption that shocks are uniformly distributed, the result that new and more extreme policies are more informative holds more generally, and is at the core of our mechanism. We further discuss this point in section 8.3 below.

The intuition behind this results is as follows. In *expectation*, a reform that moves policy in the optimal direction for the voter generates a higher payoff than one that moves in the wrong direction. However, the voter's *realized* payoff on each dimension is also a function of a random period-specific shock. Thus, the voter has a hard time distinguishing information from noise: did the reform produce a high (low) payoff because it moved in the right (wrong) direction, or purely as a consequence of the shock? Crucially, this inference problem is easier to solve when the implemented policy is more extreme. Consider for example a policy on the economic dimension. If the incumbent raises or lowers taxes in a radical fashion, changes to the voter's economic welfare are increasingly likely to be the result of the incumbent's chosen policy, rather than being the result of random fluctuations in the economy. As a consequence, the voter learns more about whether the policy was moved in the optimal direction or not. In contrast, if the incumbent adjusts tax policy modestly, observed changes to welfare may plausibly be attributed to chance and do not provide strong information about the desirability of the incumbent's policies. The voter is then unable to learn.

Figure 1 provides a graphical illustration of the results in Lemma 1. The blue and orange functions represent the conditional outcome distributions (i.e., the distributions of the voter's realized utility), under state of the world $x_v = \alpha$ and $x_v = -\alpha$, respectively. In the left panel, a moderate right-wing policy produces a partial overlap in the conditional distributions. If the realized utility falls within this range of overlap, the voter learns nothing and her beliefs remain at the prior (due to the uniformly distributed shock). If the realized utility is outside the region of overlap, the voter learns everything and her posterior moves to 0 (for low realizations) or 1 (for high realizations). As the policy becomes more extreme, the two distributions are pulled further apart, and the region of overlap shrinks. In the right panel, the policy is sufficiently extreme that there is no overlap in the conditional distributions and the voter always learns the true value of d_v.

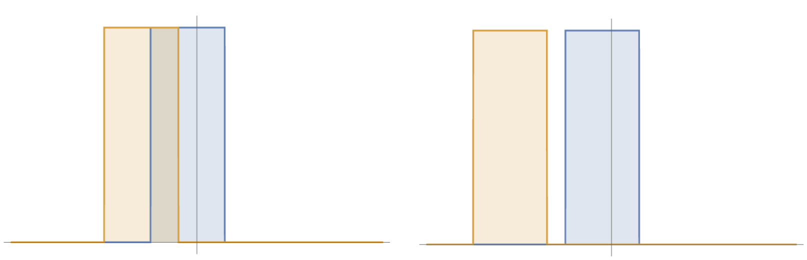

Figure 1: **Voter Learning.** The two plots display the realized voter utility on dimension D under a positive (blue function) and negative (orange function) state of the world. The policy extremism increases from left to right.

We now consider *indirect* learning, which emerges due to learning spillovers across correlated dimensions. We show that, even if the primary-dimension outcome is uninformative, the voter can learn something about her optimal policy on X if she observes an informative outcome on the secondary dimension Z, and vice versa.

Denote by $\mu^x(\tilde{\mu}^z)$ the voter's final posterior on x_v, as a function of her realized utility on *both* dimensions X and Z (via her interim beliefs $\tilde{\mu}^z \in \{0, \pi_z, 1\}$). Then, we have:

Lemma 2.

- *If $\tilde{\mu}_x \neq \pi_x$, then $\mu^x(0) = \mu^x(\pi_z) = \mu^x(1) = \tilde{\mu}^x$.*
- *If instead $\tilde{\mu}_x = \pi_x$, then $\mu^x(0) < \mu^x(\pi_z) = \tilde{\mu}^x < \mu^x(1)$.*

If the voter directly learns by observing an informative outcome on the primary dimension (i.e., $\tilde{\mu}^x \in \{0, 1\}$), then learning spillovers are irrelevant, and $\mu^x = \tilde{\mu}^x$. Instead, when no direct learning occurs on X (i.e., $\tilde{\mu}^x = \pi_x$), learning spillovers determine the voter's posterior. In this case, if the voter observes an uninformative outcome on Z ($\tilde{\mu}^z = \pi_z$), then $\mu^x = \tilde{\mu}^x$. Suppose instead that the voter learns that her ideal point on Z is a right-wing one, $\mu^z = 1$. Then, she becomes more convinced that her optimal policy on X is right-wing as well. Vice versa, if she learns that $z_v = -\alpha$ (i.e., $\mu^z = 0$), she updates negatively on the primary dimension as well. Intuitively, if the outcomes from economic and healthcare policies are correlated (or at least perceived as such by the voter), then a positive experience with a liberal economic policy will predispose the voter toward liberal policies on healthcare as well. The higher the correlation across dimensions ρ, the stronger these learning spillovers will be.

5 Voter Decision

Moving to equilibrium analysis, we proceed by backward induction. In the second period, both the incumbent and the challenger implement their preferred policy on each dimension if elected (if indifferent, they leave the status quo unchanged). Under Assumption 1, politicians have no incentives to reform policy on the secondary dimension Z in the second period, and the voter always anticipates that the status quo inherited from the first period will remain in place. Thus, even if the voter cares about both policy issues, her electoral choice is simply a function of her posterior beliefs on the primary dimension X. Recall that we denote by μ^d the voter's posterior that her ideal policy on dimension d is a right-wing one, $\mu^d = \Pr(d_v = \alpha)$, and that we impose $d_I = -d_C > 0$. Then:

Remark 1. *Let $\lambda_I^z = \lambda_C^z = 0$. In equilibrium, the voter reelects the incumbent if and only if $\mu^x > \frac{1}{2}$.*

Proof. All Proofs are collected in the Appendix. $\square$

As established in the previous section, the voter's posterior μ^x is a function of her realized utility on *both* policy dimensions, due to the emergence of learning spillovers. This, in turn, generates strategic incentives for the incumbent to manipulate voter learning using policymaking on both X and Z. The following sections characterize such incentives. For this purpose, and building on Remark 1, we introduce the following definition:

Definition 1. *If $\pi_x > \frac{1}{2}$, the incumbent is ex-ante leading. If $\pi_x \leq \frac{1}{2}$, the incumbent is ex-ante trailing.*

When $\pi_x > \frac{1}{2}$, the voter ex-ante prefers to retain the incumbent. For $\pi_x \leq \frac{1}{2}$ she instead prefers the challenger absent new information.[9]

6 The Incumbent's Problem

The findings in the previous sections shed light on how the incumbent's decisions impact his expected payoff. There are two key effects at play. The first is a *static* ideological effect, which is relatively straightforward. When the incumbent's implemented policy aligns more closely with his own ideological stance, his first-period payoff increases. The second is a *dynamic* information effect, which is more intricate. This effect revolves around how the incumbent's choices in the first period

[9] The indifference breaking assumption is without loss of generality.

affect his retention chances and thus his expected second-period payoff. This, in turn, depends on voter learning: The policy implemented on each dimension influences the likelihood of the voter *directly* learning her optimal policy for that specific dimension. In addition, the correlation across dimensions generates learning spillovers, so that the implemented policy on X can also *indirectly* influence the voter's beliefs on Z (and vice versa).

These two effects, ideological and informational, generate a potential trade-off for the incumbent. On 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, to encourage the optimal level of voter learning from the incumbent's perspective.[10]

Recall that, from Lemma 1, more extreme policies that move further from the status quo are more likely to generate informative outcomes. Thus, depending on whether information is electorally beneficial or not, the incumbent will have incentives to distort his choice either to the extreme or towards the status quo $d_{sq} = 0$. As we will see, whether one or the other distortion emerges in equilibrium is contingent upon two key factors: the incumbent's initial electoral prospects and his decision to either focus on a single policy dimension or broaden the scope of policymaking.

We will proceed by analyzing different versions of the model, shutting down each strategic force in turn before presenting the general model. To facilitate the comparison between these different benchmarks, we will impose the following assumption:

Assumption 2. $x_I < \frac{1}{4\alpha\psi_x\lambda_x}$.

[10]This trade-off clearly appears in the incumbent maximization problem, which we can express as follows:

$$\max_{x_1,z_1} \; -\lambda_I^x(x_1-x_I)^2 - \lambda_I^z(z_1-z_I)^2 - (1-\mathbb{P}(x_1,z_1))\left[\lambda_I^x(x_I-x_C)^2+\lambda_I^z(z_I-z_C)^2\right],$$

where $\mathbb{P}(x_1, z_1)$ denotes the incumbent's retention probability, which is a function of his policy choices. Recall that we are imposing $\lambda_I^z = 0$ in this baseline model. Then, the first-order necessary conditions for an interior maximum are, respectively:

$$(x_1) \quad -2\lambda_I^x(x_1-x_I) + \frac{\partial\mathbb{P}(x_1,z_1)}{\partial x_1}\lambda_I^x(x_I-x_C)^2 = 0,$$

$$(z_1) \quad \frac{\partial\mathbb{P}(x_1,z_1)}{\partial z_1}\lambda_I^x(x_I-x_C)^2 = 0.$$

This implies that the equilibrium policy on Z will be solely a function of the electoral incentives, i.e., the sign of $\frac{\partial\mathbb{P}(x_1,z_1)}{\partial z_1}$.

Recall that $\frac{1}{4\alpha\psi_x\lambda_x}$ is the smallest (positive) policy that guarantees direct learning on dimension d (from Lemma 1). This assumptions guarantees that the incumbent 's electoral incentives influence his policy choice,[11] and that the incumbent always chooses to act on the primary dimension X.

7 A Unidimensional Benchmark

To better understand the strategic incentives within our environment, it is useful to begin by analyzing a unidimensional benchmark. For this purpose, Suppose that $\lambda_v^z = 0$, so that the voter only cares about the primary dimension X. Crucially, this implies that there are no learning spillovers across dimensions.[12] As a consequence, the voter's posterior and, in turn, her electoral choice, only depend on the outcome of X. Denote by d_u the incumbent's optimal policy on dimension d in this unidimensional benchmark. We have:

Proposition 1. *Let $\lambda_I^z = \lambda_C^z = \lambda_v^z = 0$. In equilibrium:*
The incumbent does not act on the secondary dimension Z. On the primary dimension X:
(i) A leading incumbent implements a policy more moderate than his bliss point, $x_u < x_I$;
(ii) A trailing incumbent implements a policy more extreme than his bliss point, $x_u > x_I$.

As described above, when the voter cares solely about dimension X, the absence of learning spillovers implies that the incumbent's policy choice on Z is inconsequential for his retention chances. Since the incumbent does not ideologically care about the secondary dimension, he is indifferent between acting on it and keeping the status quo and (by assumption) chooses not to act. In contrast, the implemented policy on X determines the probability of the voter observing an informative outcome and, thus, the incumbent being reelected. As a consequence, the incumbent's policy choice on the primary dimension is distorted away from his ideological preference. The direction of this distortion depends on whether the incumbent is ex-ante leading or trailing.

Suppose the incumbent is ex-ante *trailing*. By definition, if the voter receives no new information, she will choose to oust him. The incumbent is then reelected if and only if the voter observes an

[11] Absent this assumption, a trailing incumbent's *statically* and *dynamically* optimal policies sometimes coincide.
[12] This follows from the fact that, if $\lambda_v^z = 0$, the voter's observed payoff on dimension Z is pure noise.

informative policy outcome and learns that right-wing policies are optimal for her. Recall that more extreme policies facilitate learning. Then, the incumbent has incentives to gamble and distorts policy to the extreme, away from both his static optimum and the status quo (normalized to 0).

In contrast, a *leading* incumbent can only be damaged by information. If the voter learns nothing new, this incumbent will be reelected for sure. Thus, leading incumbents want to avoid risks, and distort 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 1), the right-wing incumbent never implements a policy to the left of 0.

8 Multidimensional World

We now analyze the incumbent's policy choices in the multidimensional case, i.e., when the voter cares about both X and Z, $\lambda_v^z > 0$. Here, learning spillovers imply that the incumbent's retention chances are a function of his policy choices on both X and Y. That is, the incumbent's strategic problem becomes multidimensional even though his ideological preferences are unidimensional ($\lambda_I^z = 0$). Our objective is then to study the conditions under which the incumbent has strategic incentives to act on the secondary policy dimension, and how this influences his optimal choice on the primary one. For this purpose, it is useful to introduce the following:

Definition 2. *Let*

$$\widehat{\rho} \equiv \begin{cases} \pi_x & \text{if } \pi_x > \frac{1}{2} \\ 1 - \pi_x & \text{if } \pi_x < \frac{1}{2}. \end{cases}$$

We say that the correlation across dimensions X and Z is high if $\rho > \widehat{\rho}$, and low if $\rho < \widehat{\rho}$.

To begin, we characterize the incumbent's probability of winning under $\lambda_v^z > 0$. In what follows, we say that the policy outcome on dimension d is *uninformative* if it induces $\tilde{\mu}^d = \pi_d$, *favorable* if it induces $\tilde{\mu}^d = 1$, and *unfavorable* if it induces $\tilde{\mu}^d = 0$. Then, we have:

Lemma 3. *Let $\lambda_I^z = \lambda_C^z = 0$ and $\lambda_v^z > 0$. If the correlation across dimensions (ρ) is low, the incumbent's probability of winning is the same as in the unidimensional benchmark:*

- *A trailing incumbent is reelected if and only if the outcome on X is favorable*

- *A leading incumbent is reelected unless the outcome on X is unfavorable.*

If ρ is high, the incumbent's probability of winning is a function of his policy choice on both dimensions. Specifically:

- *A trailing incumbent is reelected if the outcome on X is favorable, or if the outcome on X is uninformative and the outcome on Z is favorable.*

- *A leading incumbent is reelected unless the outcome on X is unfavorable, or the outcome on X is uninformative and the outcome on Z is unfavorable.*

To understand this result, suppose first that π_x is relatively high and $\rho = \frac{1}{2}$. Recall that, even though the voter cares about both dimensions, her electoral choice only depends on her beliefs on X, since she expects both the incumbent and the challenger to keep the status quo on dimension Z if elected in the second period. Then, outcomes on Z matter for the voter's decision only insofar as they provide information about the optimal policy on X. Under $\rho = \frac{1}{2}$, learning spillovers can never emerge, therefore outcomes on Z are electorally irrelevant.

Suppose now that the two dimensions are correlated ($\rho > \frac{1}{2}$). In this case, indirect learning can affect the voter's electoral decision. That is, observing a negative (positive) outcome on healthcare leads the voter to adjust her beliefs against (in favor of) the incumbent on the economy as well.

When the correlation is too low (relative to the prior π_x), learning spillovers are too weak and the secondary dimension continues to have no impact on the voter electoral decision. Instead, when the correlation across issue areas is high, learning spillovers are strong and outcomes on the secondary dimension hold significant electoral weight. In such a scenario, a leading incumbent may find himself in electoral jeopardy even in the absence of direct learning concerning the primary dimension if the outcome related to Z is both informative and unfavorable. Conversely, in a symmetrical fashion, if a trailing incumbent successfully generates favorable information regarding the secondary dimension, the spillover effects of learning become pivotal in propelling him toward re-election.

8.1 Policymaking in a Multidimensional World

Building on Lemma 3, we now characterize the incumbent's policy choice in this multidimensional world. First, we note that, as in the benchmark model, the incumbent always acts on the primary dimension X:

Remark 2. *Let $\lambda_I^z = \lambda_C^z = 0$ and $\lambda_v^z > 0$. In equilibrium, the incumbent always acts on X.*

Thus, if the incumbent chooses to act on Z as well, this always represents an expansion of the dimensionality of policymaking. The next result, which follows straightforwardly from Lemma 3, characterizes the conditions under which the incumbent chooses to open this secondary policy dimension, despite not having ideological preferences to do so.

Corollary 1. *Let $\lambda_I^z = \lambda_C^z = 0$ and $\lambda_v^z > 0$. In equilibrium, the incumbent acts on Z if and only if he is trailing and the correlation with the primary dimension is high.*

A leading incumbent never has strategic incentives to act on Z, since he wants to prevent the voter from obtaining any new information. Suppose instead that the incumbent is ex-ante trailing. Then, he *wants* to facilitate learning spillovers, in hopes of overcoming his initial disadvantage and jumping above the retention threshold. Even still, as highlighted above, outcomes on the secondary dimension remain electorally irrelevant if the correlation ρ is low (according to Definition 2). In this case, 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.

8.2 Multidimensionality and Extremism: the Substitution Effect

We now study how the multidimensionality of voter preferences influences the nature of the policies pursued by the incumbent. Our central result uncovers a strategic *substitution effect*, whereby a trailing incumbent, who pursues extreme policies on the secondary dimension, goes moderate on the primary one. Recall that x_u is the incumbent's optimal policy choice in the unidimensional benchmark. Then, we have:

Proposition 2 (**Substitution Effect**)**.** *Let* $\lambda_I^z = \lambda_C^z = 0$ *and* $\lambda_v^z > 0$*. Suppose that the correlation between the two dimensions is high and the incumbent is trailing, so that in equilibrium he chooses to act on the secondary dimension* Z*. Then:*

(i) on the primary dimension, $x_1^* < x_I < x_u$*;*

(ii) on the secondary dimension, $z_1^* \geq \frac{1}{4\alpha\psi_z\lambda_v^z}$*.*

Recall that $\frac{1}{4\alpha\psi_z\lambda_v^z}$ is the smallest reform that guarantees full learning on Z. Then, Proposition 2(ii) is intuitive and follows straightforwardly from Lemma 3. 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. Interestingly, this implies that in our setting extreme policymaking need not follow from extreme ideological preferences. In fact, extreme policymaking can emerge on issues over which policymakers have no ideological preferences at all.

Following the gambling for resurrection logic, one might think that, in a multidimensional world, a trailing incumbent should have incentives to pursue extreme policies on both dimensions. Surprisingly, Proposition 2(i) shows that a high correlation between dimensions induces *moderation* on the primary dimension: $x_1^* < x_I$. This is in sharp contrast with our results from the unidimensional benchmark, where in equilibrium $x_u > x_I$. Here, the incumbent continues to distort his policy choice on X away from his static ideal point, as in the unidimensional benchmark, but the *direction* of this distortion changes.

To understand this result, recall that the outcome on Z can influence the voter's retention decision only if she does not learn about x_v directly (as otherwise she reaches a degenerate interim posterior $\tilde{\mu}^x$). That is, healthcare affects the voter's selection only if she fails to learn directly about the merits of the incumbent's economic policy. Thus, to benefit from the learning spillovers, the incumbent must avoid generating an informative outcome on X (in this case, economic policy).

In sharp contrast with the results of the unidimensional baseline, then, this scenario generates incentives for the trailing incumbent to pursue *moderate* policies on the primary dimension.

Notice that, in principle, the incumbent could find it optimal to forgo the learning spillovers and gamble on X instead. Put differently, a trailing incumbent might want to skip healthcare policymaking altogether and instead gamble on radical economic policy. However, we know from Corollary 1 that this *never* occurs in equilibrium. The reason is that, ex ante, players have less information on the secondary dimension than on the primary one—they know better about the possible consequences of specific economic policies than they do about policies on healthcare.

To see this, notice that the incumbent is trailing if and only if $\pi_x < \frac{1}{2}$. Thus, even though a 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. To be clear, this is true on both dimensions (i.e., $\pi_x < \frac{1}{2}$ implies $\pi_z < \frac{1}{2}$ as well); however, because the players ex-ante have less information on Z than on X, when $\pi_x < \frac{1}{2}$ we have that $\pi_z = \pi_x \rho + (1 - \pi_x)(1 - \rho) > \pi_x$. In other words, when the incumbent is trailing, the ex-ante probability of generating a favorable outcome is higher on the secondary dimension than on the primary one. Consequently, a trailing incumbent prefers to gamble on Z, hoping to exploit a false positive. Given these incentives, in equilibrium the incumbent will never gamble on both dimensions. Rather, if the correlation is too low to exploit 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 prefer to avoid risky choices on X.

In concluding this section, we note that one may worry that the substitution effect arises solely because, under Assumption 1, the voter's beliefs about the secondary dimension do not directly influence her retention choice (i.e., outcomes on Z only matter via learning spillovers). In Section 9, we show that this is not the case: the effect emerges even when this assumption is relaxed.

8.3 Scope Conditions and Robustness

Voter information about incumbent's ideology. So far, we assumed that voters know politicians' ideological preferences. Consequently, the incumbent's choices influence the voter's choice solely by revealing information about her own preferences. Reassuringly, our main insights still hold in a setting where voters face uncertainty about both their optimal policy and the incumbent's ideology, as long as the incumbent's electoral advantage (or disadvantage) is sufficiently pronounced.

To see this, suppose that each politician may be one of two types, moderate or extreme. More precisely, the voter still knows that the incumbent is right-wing (i.e., his ideal policy is a positive one) and the challenger is left-wing. However, she does not know the exact location of the two politicians' ideal policies. In this setting, the voter may draw inferences about the incumbent's type by observing his policy choice. At the same time, as in our model presented above, the *outcome* of the implemented policy reveals information about the voter's own preferences.

Suppose then that the voter's prior belief that her optimal policy is right-wing is so low (compared to her beliefs about the challenger) that, even if she knew the incumbent was moderate, she would still prefer ex-ante to replace him with the challenger.[13] In this case, the incumbent's incentives are exactly as in our baseline model. This incumbent knows that, if the voter receives no new information about the location of her own ideal policy, she will choose to replace him. His only hope of being re-elected is to generate an informative outcome that induces the voter to learn that her own preferences are aligned with right-wing policies.[14] Then, the incumbent always has

[13]Formally, let γ^C be the voter's prior that the challenger is an extreme type, and denote as $x_C^m = -x_I^m$ and $x_C^e = -x_I^e$ the moderate-types' and extreme-types' ideal policies, respectively. Recall that π_x is the prior probability that the voter's ideal policy on the primary dimension is a right-wing one. Then, suppose that π_x is sufficiently low to satisfy:

$$-\pi_x(\alpha - x_I^m)^2 - (1-\pi_x)(-\alpha - x_I^m)^2 <$$
$$-\pi_x\Big(\gamma^C(\alpha - x_C^e)^2 + (1-\gamma^C)(\alpha - x_C^m)^2\Big) - (1-\pi_x)\Big(\gamma^C(-\alpha - x_C^e)^2 + (1-\gamma^C)(-\alpha - x_C^m)^2\Big).$$

[14]To avoid trivialities we assume that, if the voter learns that her optimal policy is a right-wing one, she will prefer to re-elect the incumbent even if he is an extreme type:

$$-(\alpha - x_e^I)^2 > -\gamma^c(\alpha - x_e^C)^2 - (1-\gamma^c)(\alpha - x_m^C)^2.$$

This assumption guarantees that the voter's beliefs over her own preferences are electorally relevant.

incentives to gamble by implementing extreme policies (on either the primary or the secondary dimension, depending on ρ), even if this leads the voter to infer that he is an extreme type.[15]

A similar but symmetric logic applies to a leading incumbent when π_x is so high that, even if the voter learns that the incumbent is a extreme type, she ex-ante prefers to re-elect him.

Thus, when the incumbent is sufficiently leading or sufficiently trailing the insights generated by our model remain valid. For intermediate values of π_x the incumbent potentially faces competing sets of incentives, deriving from the signalling and the gambling components of the model. Depending on the parameter values, one or the other may dominate. Fully solving this enriched model is outside the scope of this paper, but is a promising direction for future research.

Voter learning, time lags and endogenous attention. A crucial ingredient of our theory is that the voter observes the outcome of the implemented policy, and that this occurs prior to making her electoral choice. Here, we clarify the substantive scope conditions for these assumptions.

First of all, it is plausible that there may be a time lag between the incumbent's policy decision and the realization of its consequences. For example, a less efficient bureaucratic apparatus may generate delays in policy implementation. Alternatively, even once a policy is implemented, its effect on voter welfare may take some time to become visible. Our framework incorporates such frictions, which are captured by the shock term $\varepsilon_{d,t}$. Thus, a lower (higher) variance in the distribution of this shock may be interpreted as describing a situation where these time lags are less (more) significant.

Furthermore, one element outside the scope of our model, but that would reinforce our mechanism, is that the attention voters (and media) pay to a policy reform and its consequences is endogenous to the nature of the reform. In particular, it seems reasonable to assume that more radical (extreme) reforms generate higher scrutiny and attention. This would then reinforce the informational effect emerging in our framework, whereby incumbents anticipate that more extreme reforms facilitate voter learning and choose their optimal policy accordingly.

Voter Learning with Non-Uniform Shocks. As emphasized above, an essential element of our theory is that the amount of voter learning is endogenous to the implemented policy. On

[15]More precisely, the game has a unique equilibrium, in which the incumbent fully separates and both types act as if they had no private information, implementing their dynamically optimal policy.

each dimension, more extreme policies facilitate voter learning. In our model, the voter's learning technology is stark. Due to the assumption that the shock $\varepsilon_{d,t}$ is drawn from a uniform distribution, the voter's realized utility is either fully informative or completely uninformative. The more extreme the policy, the higher the likelihood that it generates an informative outcome realization.

It is important to emphasize that this information property of extreme policies holds more generally, beyond this specific distributional assumption. In particular, extreme policies are more informative whenever the distribution of $\varepsilon_{d,t}$ satisfies the monotone likelihood ratio property (MLRP). Suppose for example that the shock is drawn from a standard normal distribution. Under this assumption, the voter's learning process is smoother than the one emerging with a uniform shock. Any outcome realization reveals *some* information, but no realization allows the voter to reach a degenerate posterior. Nonetheless, the MLRP property guarantees that, as the implemented policy becomes more extreme, the distance in the expected outcomes as a function of the state increases. This, in turn, increases each signal's informativeness. Thus, extreme policies facilitate voter learning and induce a lower-variance posterior distribution (i.e., are Blackwell more informative). We provide a formal proof of this result in Lemma A-1 in the online Appendix.

9 General Model

The results of the previous section, where politicians have no ideological preferences over the secondary dimension, are useful to isolate the strategic incentives generated by the learning spillovers. In the Appendix, we study a general model where both the voter and the politicians care about all issue areas (i.e., $\lambda_i^z > 0$ for $i \in \{I, V, C\}$). We will assume that the voter cares more about the primary dimension than the secondary one (i.e., if the voter learns that $x_v = \alpha$ but $z_v = -\alpha$, she will prefer to reelect the incumbent; see Assumption 4 in the Appendix). In other words, we assume that the intensity of voter preferences on an issue (that is, the issue's salience) is positively correlated with the amount of information the voter initially has on that issue. As we demonstrate below, under this assumption, the substitution effect identified in Proposition 2 continues to arise in this general model. Indeed, it emerges for both trailing and leading incumbents.

In contrast to the baseline model, the incumbent now holds strategic *and ideological* preferences over the secondary dimension Z. Specifically, a trailing incumbent has both strategic and ideological incentives to alter policy in this dimension, as doing so enhances voter learning opportunities, thereby improving his reelection prospects (as established in Lemma 3). By contrast, a leading incumbent faces a trade-off. While he would like to implement his ideologically preferred policy on the secondary dimension, the results from the previous section indicate that acting on this dimension could jeopardize his reelection chances–and consequently, his expected future payoff.

If ρ is low, this tradeoff does not bite. Because the voter cares more about the primary dimension, the leading incumbent can survive reelection even if the outcome on the secondary one reveals damaging information. He can therefore implement his preferred policy on the secondary dimension while avoiding the negative electoral consequences. Instead, if ρ is high, then a leading incumbent must balance dynamic electoral considerations and static ideological preferences. Proposition B-1 in the Appendix identifies sufficient conditions under which the incumbent solves this tradeoff by acting on the secondary dimension. Intuitively, if outcomes on this dimension are sufficiently noisy (i.e., ψ_z is sufficiently low), then the electoral consequences of changing policy are less worrisome for the incumbent. Thus, he chooses to act on both X and Z.

Generalizing our central result from the previous section, Proposition 3 then highlights that, if the incumbent acts on both issues in equilibrium, the substitution effect described in Proposition 2 continues to emerge. Confirming our earlier result, for a trailing incumbent this substitution effect implies that expanding policymaking to a secondary dimension induces *moderation* on the primary one. Conversely, for a leading incumbent, the substitution effect indicates that the strategic importance of dimension Z leads to increased *extremism* on dimension X.

Proposition 3 (**Substitution Effect, General Model**)**.** *Let* $\lambda_v^z, \lambda_I^z, \lambda_C^z > 0$. *Suppose that the incumbent chooses to act on the secondary dimension in equilibrium. When the correlation with the primary dimension is low, we have* $z_1^* = z_I$ *and* $x_1^* = x_u$. *Suppose instead the correlation is high. Then:*

- *when the incumbent is trailing, we have* $z_1^* > z_I$ *and* $x_1^* < x_u$*;*

- *when the incumbent is leading, we have* $z_1^* < z_I$ *and* $x_1^* > x_u$*.*

Recall that d_u is the equilibrium policy on dimension d in the unidimensional baseline, i.e., the world in which the voter only cares about the primary dimension X ($\lambda_v^z = 0$). As discussed previously, when the correlation is low, dimension Z is electorally irrelevant. Then, the incumbent implements his ideologically preferred policy on Z. On the primary dimension X, the equilibrium policy aligns with the incumbent's choice in the unidimensional baseline.

When instead the correlation is high, the intuition for the case in which the incumbent is trailing is exactly as described in the previous section. The trailing incumbent has incentives to exploit learning spillovers. He then gambles by distorting policy towards the extreme on Z, where a false positive is more likely, and moderates on the primary dimension X.

Suppose instead that the incumbent is leading. When his ideological preferences prompt him to act on the secondary dimension, he assumes greater electoral risk. By reforming policy on Z he may in fact generate an informative and unfavorable outcome on this dimension, which under a high ρ would hurt his reelection chances. Further, since $\pi_x > \frac{1}{2}$ implies $\pi_z < \pi_x$, an unfavourable outcome is ex-ante *more likely* on dimension Z than on X. Additionally, recall that since the voter prioritizes the primary dimension, direct learning on X makes outcomes on Z electorally irrelevant. These observations suggest that to mitigate the negative impacts of learning spillovers from the secondary dimension, the leading incumbent has incentives to *facilitate* direct learning on the primary one. In other words, a leading incumbent—who typically adopts a risk-averse (moderate) approach in a unidimensional world—becomes more risk-seeking on the primary dimension when he expands the scope of policymaking.

10 Empirical Implications

Our theoretical findings highlight that policymaking in a multidimensional world can appear crucially different than considering each issue in isolation. Perhaps counter-intuitively, when policy issues are sufficiently correlated, the incentives that a policymaker faces could even be opposite to those faced in a unidimensional setup.

Importantly, these findings generate specific observable implications that are unique to our theory. Proposition 3 uncovers a subtle relationship between issue salience and policy extremism, which depends crucially on the incumbent's electoral prospects. In particular, we should expect *trailing* incumbents to adopt moderate stances on primary (more salient) issues, and go extreme on less salient ones. In contrast, *leading* incumbents should moderate on secondary issues and pursue bold policies on salient, central ones.[16]

These predictions are broadly consistent with our motivating examples in the introduction. However, to our knowledge, the relationship between policy extremism, issue salience, and the incumbent's electoral prospects has not been systematically explored in empirical research.

Below, we outline a roadmap for empirical research, designed to serve as a guide for scholars interested in evaluating the main implications of our model. In doing so, we discuss findings from a preliminary "reality-check" we have pursued, combining existing available data from the US context.

An empirical roadmap. Here, we focus on the relationship between policymaking, issue salience, and the president's electoral prospects. Of course, the president does not have sole policymaking authority in the US. Nonetheless, presidents can, and do, set the tone of policymaking as leaders of their parties, and use institutions like the State of the Union to set policymaking agendas. They also can make use of expanded opportunities for unilateral actions, particularly in recent years. As such, we believe our theory should provide insights into their strategic incentives when managing policymaking in a multidimensional world.

The dependent variable related to our central result is the extremity of policy change on different policy dimensions. To measure it, we must compute the distance between the location of the status quo policy on each policy dimension, and the location of the implemented policy (if any). For this

[16]Proposition 3 identifies this effect across dimensions that the voter perceives as sufficiently correlated, so that learning spillovers are significant. Empirically, we believe this to be the case for most issues that are on the policy agenda: as discussed in section 3.1, voters increasingly see different issues are connected, and form their policy preferences accordingly.

purpose, we use data from Lowande et al. (2021), who uses legislative candidate responses to policy questions—indicating whether policies in specific issue areas was too liberal or conservative—to generate estimates of the status quo policy locations in each of the surveyed issue areas for the 103rd through 114th Congresses. Then, we compare the location of the status quo at the end of one congressional session with the status quo in that same issue area at the beginning of the session.[17] Matching our theoretical quantity, this implies that reforms are coded as more extreme the larger the distance from the inherited status quo.

The next key variable of our theory captures primary and secondary policy dimensions. Here, empirical tests should use measures of issue salience. We make use of Heffington, Park and Williams's "Most Important Problem" Dataset (2019). Finally, our model requires a measurement of whether an incumbent—in this case, an incumbent president—is "leading" or "trailing." Ideally, we would have a measure of incumbents' beliefs about whether they were likely to win or lose in the upcoming election. Such data do exists for U.S. presidents—the Iowa Electronic Markets provide a strong example—but these data do not extend beyond a few months prior to an election.

Lacking this data, we adopt two approaches. First, we use presidential approval ratings compiled by the Roper Center. We proceed by subtracting presidential disapproval rates from overall approval rates. Using this 'net approval' measure, we then code a president as "leading" if his approval rating is a net-positive, and trailing otherwise. While presidential approval is a good proxy of the president's own popularity, it does not capture broader features of the environment that would, according to our theory, influence the incumbent's strategic incentives. More specifically, the president may be concerned not only about his own popularity, but about the chances of his party obtaining a majority in congress (so that even term-limited incumbents may face incentives similar to the ones highlighted in our theory). To capture these concerns, we also provide a qualitative measurement that classifies presidents as "leading" or "trailing" also taking into account the popularity of their party for each Congress. This allows us to capture cases when the president retains fairly strong popularity due to outside factors like the honeymoon effect, even when his

[17]Our unit of analysis is then issue-president-congressional session.

party is faring quite poorly. During Pres. Obama's first term, for example, Democrats' political struggles eventually led to historic losses in the 2010 election—despite the fact that Obama retained relatively high approval ratings throughout.

One specific reason why the president may care about the composition of Congress is that an aligned legislature facilitates his action in office. Straightforwardly, these considerations are moot for a second-term president during his last two years in office, since no other congressional election will occur during his tenure. Thus, we consider the full sample of observations (N = 106) as well as restricted sample where we drop the last two year of a president's second term (N = 78). We expect our results to be stronger in our restricted sample, where strategic incentives to control information should be more pronounced.

To test our theory, we regress our measurement of the extremism of policy change on an interaction between issue salience and our indicator for leading status. We expect the coefficient for the interaction term to be positive. We report results from these regressions in Table 1. Across all specifications, results are consistent with our theory. In models considering the full sample ((1) and (3)), the interaction term is significant at the 10% level. As expected, the results are stronger in the restricted-sample models ((2) and (4)), where our coefficient of interest is larger in magnitude and significant at 5% level. The substantive magnitude of the effect remains comparable across models. As Figure 2 depicts graphically, for leading incumbents a shift from the minimum to maximum issue salience is associated with more than a 2-unit increase in extremity—on a scale that ranges roughly from -5 to 5. Conversely, for trailing incumbents a minimum-to-maximum shift in salience is associated with just under a one-unit *decrease.*

These results, of course, are noisy, and we do not aim to over-interpret the associations we depict. However, they do underscore how our theory may help empirical researchers better understand important outcomes like policy extremism, by disentangling interrelated factors like issue salience and politicians' electoral security. That being said, more research is doubtlessly needed to probe the robustness of these results, both within and outside the US.

Table 1: Policy Extremity, Issue Salience, and Electoral Expectations in the U.S.

	Dependent variable:			
	Extremity of Policy Change			
	(1)	(2)	(3)	(4)
Issue Salience	−0.065	−0.065	−0.074	−0.079
	(0.084)	(0.078)	(0.092)	(0.086)
Leading (Qualitative)	−1.150	−1.409		
	(0.803)	(0.876)		
Leading (Pres. Approval)			−0.582	−0.766
			(0.774)	(0.807)
Salience∗Leading Qual.	0.224^{*}	0.295^{**}		
	(0.114)	(0.124)		
Salience∗Leading Appr.			0.197^{*}	0.241^{**}
			(0.117)	(0.120)
Constant	1.689^{***}	1.661^{***}	1.552^{**}	1.660^{***}
	(0.630)	(0.594)	(0.642)	(0.615)
Observations	106	78	106	78
Presidential FEs	✓	✓	✓	✓
R^2	0.066	0.096	0.075	0.097
Note:		*p<0.1; **p<0.05; ***p<0.01		

Additional Implications. In concluding this section, we note that our theory's implications may help explain some puzzling empirical results on the relationship between issue salience and policy congruence (i.e., the alignment between policies preferred by a majority of voters and those implemented by policymakers). Intuitively, one would expect a positive relationship between congruence and salience: that is, the more salient a particular issue, the higher should be the efforts of a policymaker to match the voters' preferred policy on that issue. However, the empirical evidence is mixed (Canes-Wrone and Shotts, 2004; Rasmussen, Reher and Toshkov, 2019).

Our model provides a potential explanation for why the observed correlation in the data does not

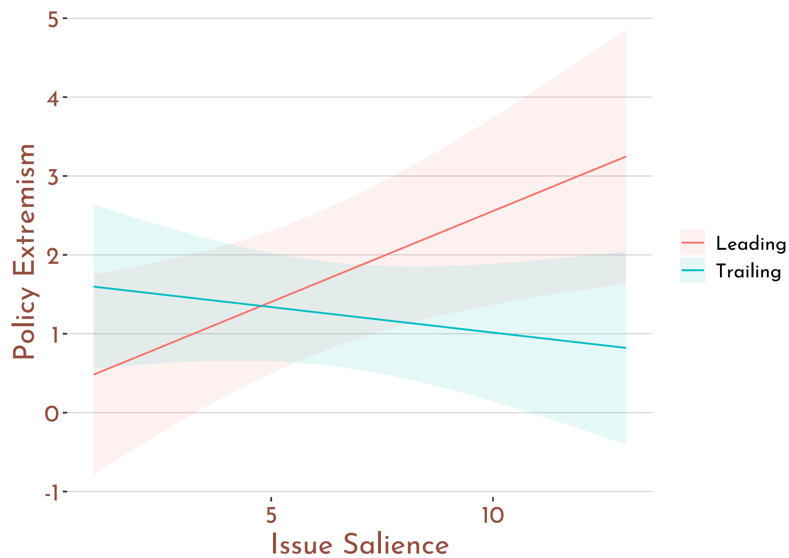

Figure 2: Conditional Relationship between Issue Salience and Policy Extremism

align with intuition. Our results on the substitution effect imply that the relationship between issue salience and policy congruence is, once again, mediated by the incumbent's electoral prospects. To see this, suppose for illustration purposes that the voter's preferred policy (given her prior beliefs) is sufficiently moderate. Proposition 3 then implies that for leading incumbents we should observe more policy congruence on secondary, less salient, dimensions, while for trailing ones the opposite should hold. Empirical scholars should then account for this interaction in order to obtain an accurate picture of the effect of issue salience on policy extremism.

11 Conclusions

In this paper, we explore policymaking in a multidimensional setting where voters face uncertainty about the outcomes of policy decisions and different issue areas are correlated. We identify conditions that incentivize officeholders to expand (or contract) the scope of policymaking. Additionally, we analyze how the existence of multiple correlated dimensions, which voters may prioritize differently, influences the decision to pursue moderate or extreme reforms.

Our central result uncovers a substitution effect across correlated dimensions. For a trailing incumbent, the possibility of policymaking in multiple dimensions presents greater opportunities for voter learning. As a result, such incumbents *expand* the scope of policymaking and, at the same time, *moderate* on the primary dimension. This substitution effect emerges precisely because the different dimensions are connected, and the resulting learning spillovers fundamentally alter the incumbent's strategic calculus. For a leading incumbent the opposite holds: While he always pursues moderate policy in a unidimensional world, in a multidimensional one he may choose more extreme policies on the primary dimension in order to mitigate the negative consequences of learning spillovers. When instead voters perceive various policy issues as essentially unconnected (the correlation is low), these dynamics remain dormant, and the incumbent's choices in the multidimensional world resemble those in the unidimensional baseline.

These findings significantly impact our understanding of the transition from a unidimensional to a multidimensional worldview in political analysis. Theoretically and empirically, the assumption that a unidimensional framework adequately represents a multidimensional reality is common. This perspective is often justified by noting that policy preferences across issues are correlated. When different issue areas are connected, the argument goes, a unidimensional model must be a good enough proxy for our multidimensional world (McMurray, 2014).

Our work identifies a framework where this logic breaks down. Indeed, it is precisely *because* issues are correlated that policymaking in the multidimensional world may be fundamentally different from the unidimensional case. When the correlation across dimensions is sufficiently strong, the substitution effect we identify suggests that the policies implemented in equilibrium in a multidimensional model might be qualitatively different from those we would expect if each issue were considered in isolation. Therefore, a model with a unidimensional policy space may fail to adequately capture the incentives and nature of policymaking in a multidimensional world.

Beyond this general point, our theory offers specific observable implications regarding the relationship between policy extremism and issue salience, and how this relationship is mediated by the incumbent's electoral prospects. An initial look at data from the US context yields promising, though noisy, results, and offers a guide for future empirical research.

References

Achen, Christopher H. 1992. "Social psychology, demographic variables, and linear regression: Breaking the iron triangle in voting research." *Political behavior* 14:195–211.

Ash, Elliott, Massimo Morelli and Richard Van Weelden. 2017. "Elections and divisiveness: Theory and evidence." *The Journal of Politics* 79(4):1268–1285.

Ashworth, Scott. 2005. "Reputational dynamics and political careers." *Journal of Law, Economics, and Organization* 21(2):441–466.

Ashworth, Scott and Ethan Bueno de Mesquita. 2006. "Delivering the goods: Legislative particularism in different electoral and institutional settings." *The Journal of Politics* 68(1):168–179.

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

Banks, Jeffrey S and John Duggan. 2008. "A dynamic model of democratic elections in multidimensional policy spaces." *Quarterly Journal of Political Science* 3(3):269–299.

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. 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.

Callander, Steven and Patrick Hummel. 2014. "Preemptive policy experimentation." *Econometrica* 82(4):1509–1528.

Canes-Wrone, Brandice and Kenneth W Shotts. 2004. "The conditional nature of presidential responsiveness to public opinion." *American Journal of Political Science* 48(4):690–706.

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

Dewan, Torun and Rafael Hortala-Vallve. 2019. "Electoral competition, control and learning." *British Journal of Political Science* 49(3):923–939.

Fiorina, Morris P. 1981. *Retrospective voting in American national elections*. Yale University Press.

Gieczewski, Germán and Svetlana Kosterina. 2020. "Endogenous Experimentation in Organizations." *Princeton University Typescript* .

Hare, Christopher. 2022. "Constrained Citizens? Ideological Structure and Conflict Extension in the US Electorate, 1980–2016." *British Journal of Political Science* 52(4):1602–1621.

Heffington, Colton, Brandon Beomseob Park and Laron K Williams. 2019. "The "Most Important Problem" Dataset (MIPD): a new dataset on American issue importance." *Conflict Management and Peace Science* 36(3):312–335.

Hirsch, Alexander V. 2016. "Experimentation and persuasion in political organizations." *American Political Science Review* 110(1):68–84.

Izzo, Federica. 2023. "Ideology for the Future." *American Political Science Review* .

Izzo, Federica. Forthcoming. "With friends like these, who needs enemies?" *JOP* .

Izzo, Federica, Gregory J Martin and Steven Callander. 2023. "Ideological Competition." *American Journal of Political Science* .

Jewitt, Caitlin E and Paul Goren. 2016. "Ideological structure and consistency in the age of polarization." *American Politics Research* 44(1):81–105.

Kunda, Ziva. 1987. “Motivated inference: Self-serving generation and evaluation of causal theories.” *Journal of personality and social psychology* 53(4):636.

Lau, Richard R and David P Redlawsk. 2001. “Advantages and disadvantages of cognitive heuristics in political decision making.” *American journal of political science* pp. 951–971.

Levendusky, Matthew S. 2010. “Clearer cues, more consistent voters: A benefit of elite polarization.” *Political Behavior* 32:111–131.

Lowande, Kenneth et al. 2021. “Presidents and the status quo.” *Quarterly Journal of Political Science* 16(2):215–244.

McMurray, Joseph. 2014. Why the political world is flat: an endogenous left-right spectrum in multidimensional political conflict. In *APSA 2014 Annual Meeting Paper*.

Mettler, Suzanne and Joe Soss. 2004. “The consequences of public policy for democratic citizenship: Bridging policy studies and mass politics.” *Perspectives on politics* 2(1):55–73.

Pierson, Paul. 1993. “When effect becomes cause: Policy feedback and political change.” *World politics* 45(4):595–628.

Rasmussen, Anne, Stefanie Reher and Dimiter Toshkov. 2019. “The opinion-policy nexus in Europe and the role of political institutions.” *European Journal of Political Research* 58(2):412–434.

Sniderman, Paul M, Richard A Brody and Phillip E Tetlock. 1991. *Reasoning and choice: Explorations in political psychology*. Cambridge University Press.

Stimson, James A. 2018. *Public opinion in America: Moods, cycles, and swings*. Routledge.

Strulovici, Bruno. 2010. “Learning while voting: Determinants of collective experimentation.” *Econometrica* 78(3):933–971.

Strumpf, Koleman S. 2002. “Does government decentralization increase policy innovation?” *Journal of Public Economic Theory* 4(2):207–241.

Taber, Charles S and Milton Lodge. 2006. “Motivated skepticism in the evaluation of political beliefs.” *American journal of political science* 50(3):755–769.

Tavits, Margit. 2007. “Principle vs. pragmatism: Policy shifts and political competition.” *American Journal of Political Science* 51(1):151–165.

Volden, Craig, Michael M Ting and Daniel P Carpenter. 2008. “A formal model of learning and policy diffusion.” *American Political Science Review* 102(3):319–332.

Zingher, Joshua N and Michael E Flynn. 2019. “Does polarization affect even the inattentive? Assessing the relationship between political sophistication, policy orientations, and elite cues.” *Electoral Studies* 57:131–142.

Online Appendix for "Learning in a Complex World: How Multidimensionality Affects Policymaking"

Table of Contents

Appendix A: Main Results

Proof of Lemma 1. 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), and conditioning on the true state x_v has to fall within:

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

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

The last case to consider is when u_v^t falls within the interval $\left[-\lambda_v^x(x_t-\alpha)^2 - \frac{1}{2\psi_x}, -\lambda_v^x(x_t+\alpha)^2 + \frac{1}{2\psi_x}\right]$. Denote by $f(\cdot)$ the PDF of the error term $\varepsilon_{x,t}$. When u_v^t falls within this interval we have that:

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

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

$$\Pr(x_v = \alpha | u_v^t) = \Pr(x_v = \alpha) = \pi_x.$$

(ii)-(iii) Now, denote by $L_x \in \{0,1\}$ 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 (A-1) is empty. This requires:

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

which rearranged yields:

$$x_t \geq \frac{1}{4\alpha\lambda_v^x\psi_x}. \tag{A-3}$$

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

$$\Pr(L_x = 1|\pi_x, 0 < x_t < x') = \pi_x \Pr\left(-\lambda_v^x(x_t-\alpha)^2 + \varepsilon_{x,t} > -\lambda_v^x(x_t+\alpha)^2 + \frac{1}{2\psi_x}\right) + (1-\pi_x)\Pr\left(-\lambda_v^x(x_t+\alpha)^2 + \varepsilon_{x,t} < -\lambda_v^x(x_t-\alpha)^2 - \frac{1}{2\psi_x}\right).$$

Since the two probabilities are symmetric, we have

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

The proof for dimension Z is analogous therefore omitted. □

Lemma A-1. *Suppose that $\epsilon_{d,t} \sim N(0,1)$. If $|d| > |d'|$ then policy experiment d is Blackwell more informative than d'.*

Proof. The noise term is distributed normally and thus satisfies the MLRP property. Furthermore, fixing a policy d on either side of zero, the policy choice and the state of the world (i.e., the true value of d_v) are strict complements. This can be verified by noting that, for any $d' > d > 0$, we have

$$-(d'-\alpha)^2 + (d'+\alpha)^2 > -(d-\alpha)^2 + (d+\alpha)^2,$$

with the symmetric result holding for $d' < d < 0$. Thus, Theorem 3.1 of Ashworth, Bueno de Mesquita and Friedenberg (2017) applies, and shows that outcomes are more Blackwell informative as d moves away from 0 in either direction. □

Proof of Lemma 2. If the voter observes an informative outcome on X, she immediately reaches a degenerate posterior and $\mu^x = \tilde{\mu}^x$. Suppose instead that the voter observes an uninformative outcome on X, so that $\tilde{\mu}^x = \pi_x$. To prove the first inequality, it suffices to apply Bayes' rule to derive the voter's posterior if the voter learns $z_v = -\alpha$:

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

Similarly, suppose that the voter learns that $z_v = \alpha$. By Bayes' rule, we have:

$$\mu^x(\tilde{\mu}^x, \alpha, \rho) = \frac{\pi_x\rho}{\pi_x\rho + (1-\pi_x)(1-\rho)} \geq \pi_x. \tag{A-6}$$

□

Proof of Remark 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. Recall that, under Assumption 1, the voter expects both challenger and incumbent to leave policy on dimension Z unchanged if elected in period 2. Then, V re-elects the incumbnent iff

$$-\lambda_v^x[\mu^x(x_I - \alpha)^2 + (1-\mu^x)(x_I + \alpha)^2] > -\lambda_v^x[\mu^x(x_C - \alpha)^2 + (1-\mu^x)(x_C + \alpha)^2]. \tag{A-7}$$

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

$$\mu^x > \frac{1}{2}. \tag{A-8}$$

□

Proof of Proposition 1. By Remark 1, the incumbent is re-elected if and only if $\mu_x > 1/2$. Suppose the incumbent is leading. By Definition 1, this means that $\pi_x > 1/2$. It follows that in the absence of information $\mu^x = \pi_x$ and the incumbent is re-elected. Similarly, if V learns that $x_v = \alpha$, then $\mu^x = 1$ and the incumbent is re-elected. It is only when V learns that $x_v = -\alpha$ that $\mu^x = 0$. In this case, the incumbent is ousted. Suppose now the incumbent is trailing ($\pi_x \geq 1/2$). Analogously to the argument above, it is only when when V learns that $x_v = \alpha$ that $\mu_x > 1/2$ and the incumbent is re-elected.

Next, notice that given the assumption that $x_I < \frac{1}{4\alpha\psi_x\lambda_v^x} = x'$ and the fact that any policy $x \geq x'$ guarantees $L_x = 1$ with probability 1, in equilibrium the incumbent never implements a policy $x_I > x'$.

Suppose that the incumbent is leading ($\pi_x > \frac{1}{2}$). From above, we have that his probability of winning is $\mathbb{P}(x_1, z_1) = 1 - (1 - \pi_x)\Pr(L_x = 1|\pi_x, 0 < x_t \leq x')$. Using the proof of Lemma 1, we can substitute the value of this probability into the incumbent's maximization problem, which becomes:

$$-\lambda_I^x(x_1 - x_I)^2 - 4\alpha\psi_x\lambda_v^x x_1(1 - \pi_x)\Big(\lambda_I^x(x_I - x_C)^2 + \lambda_I^z(z_I - z_C)^2\Big). \tag{A-9}$$

Noting that $d_I = -d_C$ for $d \in \{x, z\}$, and letting $K = 4x_I^2 + 4z_I^2$, we can write the first-order necessary condition (which is also sufficient since the problem is concave) as:

$$-2\lambda_I^x(x_1 - x_I) - 4\alpha\psi_x\lambda_v^x(1 - \pi_x)K = 0. \tag{A-10}$$

Rearranging (A-10) yields:

$$x_1^* = x_I - \frac{2\alpha\psi_x\lambda_v^x(1 - \pi_x)}{\lambda_I^x}K < x_I.$$

Plugging in the value of K, assumption 2 then guarantees that this value is always positive, and thus x_1^* is in the feasible policy set. When instead I is trailing, we can express I's problem as

$$-\lambda_I^x(x_1 - x_I)^2 - (1 - 4\alpha\psi_x\lambda_v^x x_1\pi_x)K, \tag{A-11}$$

which yields the following first-order necessary condition:

$$-2\lambda_I^x(x_1 - x_I) + 4\alpha\psi_x\pi_x\lambda_v^x K = 0,$$

which rearranged yields:

$$x_1 = x_I + \frac{2\alpha\psi_x\pi_x\lambda_v^x}{\lambda_I^x}K.$$

It follows that

$$x_1 = \min\left\{x', x_I + \frac{2\alpha\psi_x\pi_x\lambda_v^x}{\lambda_I^x}K\right\} > x_I. \tag{A-12}$$

$\square$

Proof of Lemma 3. Suppose that the incumbent is leading ($\pi_x > \frac{1}{2}$). Then, $\widehat{\rho}$ solves:

$$\mu^x(\varnothing, -\alpha, \rho) = \frac{1}{2}, \tag{A-13}$$

where

$$\mu^x(\varnothing, -\alpha, \rho) = \frac{(1-\rho)\pi_x}{(1-\rho)\pi_x + \rho(1-\pi_x)}, \tag{A-14}$$

which yields:

$$\widehat{\rho} = \pi_x. \tag{A-15}$$

Since the RHS in B-4 is decreasing in ρ, under $\rho > \widehat{\rho}$ we have $\mu^x(\varnothing, -\alpha, \rho) < \frac{1}{2}$, and the leading incumbent is replaced if the voter observes an uninformative outcome on X, but learns that $z_v = -\alpha$. If instead $\pi_x < \frac{1}{2}$, $\widehat{\rho}$ satisfies:

$$\mu^x(\varnothing, \alpha, \rho) = \frac{1}{2}, \tag{A-16}$$

where

$$\mu^x(\varnothing, \alpha, \rho) = \frac{\pi_x\rho}{\pi_x\rho + (1-\pi_x)(1-\rho)}. \tag{A-17}$$

Combining the above, we have

$$\widehat{\rho} = 1 - \pi_x. \tag{A-18}$$

The RHS of B-7 is increasing in ρ, therefore under $\rho > \widehat{\rho}$ we have $\mu^x(\varnothing, \alpha, \rho) > \frac{1}{2}$, and a trailing incumbent is re-elected if the voter observes an uninformative outcome on X, but learns that $z_v = \alpha$. Thus, we have that the incumbent's probability of winning satisfies:

- $\mathbb{P}(x_1, z_1) = 1 - \Pr(L_x(x_1) = 1)(1 - \pi_x)$ when the incumbent is leading and $\rho < \widehat{\rho}$;

- $\mathbb{P}(x_1, z_1) = 1 - \Pr(L_x(x_1) = 1)(1 - \pi_x) - \Big(1 - \Pr(L_x(x_1) = 1)\Big) \Pr(L_z(z_1) = 1)(1 - \pi_z)$ when the incumbent is leading and $\rho > \widehat{\rho}$;
- $\mathbb{P}(x_1, z_1) = \Pr(L_x(x_1) = 1)\pi_x$ when the incumbent is trailing and $\rho < \widehat{\rho}$;
- $\mathbb{P}(x_1, z_1) = \Pr(L_x(x_1) = 1)\pi_x + \Big(1 - \Pr(L_x(x_1) = 1)\Big) \Pr(L_z(z_1) = 1)\pi_z$ when the incumbent is trailing and $\rho > \widehat{\rho}$.

□

We defer the proof of Remark 2, since it will be useful to reference the next two results in that proof.

Proof of Corollary 1. The proof follows from Lemma 3 and the assumption that when indifferent the incumbent prefers not to act on Z. □

Proof of Proposition 2. (i) Consider the incumbent's choice on X. Recall that in equilibrium we must have $x_1 \leq x' = \frac{1}{4\alpha\psi_x\lambda_v^x}$. Then, when I is trailing and $\rho > \widehat{\rho}$, we have

$$\mathbb{P} = 4\alpha\psi_x\lambda_v^x x_1\pi_x + (1 - 4\alpha\psi_x\lambda_v^x x_1)p^\dagger\pi_z, \tag{A-19}$$

where $p^\dagger = \min\{1, 4\alpha\psi_z\lambda_v^z z_1\}$. Since we are assuming that $\lambda_I^z = 0$, the incumbent will always find it optimal to implement a fully informative policy on the secondary dimension, $z_1^* \geq \frac{1}{4\alpha\psi_z\lambda_v^z}$. Then, the trailing incumbent's retention probability reduces to

$$\mathbb{P} = 4\alpha\psi_x\lambda_v^x\pi_x x_1 + (1 - 4\alpha\psi_x\lambda_v^x x_1)\pi_z. \tag{A-20}$$

Plugging (A-20) into the incumbent's problem yields the following first-order condition:

$$(x_1) \quad -2\lambda_I^x(x_1 - x_I) + 4\alpha\psi_x\lambda_v^x(\pi_x - \pi_z)K = 0. \tag{A-21}$$

Note that, given $\pi_x < \frac{1}{2}$, $\pi_z = \pi_x \rho + (1-\pi_x)(1-\rho) > \pi_x$, which implies that the LHS of (A-21) is negative at $x_1 \geq x_I$. Since the utility is concave, it follows that $x_1^* < x_I$.

(ii) Recall that under $\lambda_I^z = 0$ 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 $z_1^* \geq \frac{1}{4\alpha\psi_z\lambda_v^z}$. □

Proof of Remark 2. First, suppose the incumbent is leading. We know from the previous results that in equilibrium this incumbent will always set $z_1^* = 0$. Thus, the strategic problem resembles the unidimensional benchmark, and $\lambda_I^z = 0$ is enough to guarantee $x_1^* > 0$. Suppose instead the incumbent is trailing. Then, from inspection of A-21 we can verify that Assumption 2 guarantees that $x_1^* > 0$, since the incumbent's marginal utility is positive at $x_1 = 0$. □

Appendix B: General Model

Preliminaries. To facilitate comparison with the baseline model, we will impose the following assumptions (which represent a general version of Assumption 2 in the main body):

Assumption 3. *Let* $\underline{\lambda}_I^x$ *be the value that solves* $x_I - \frac{8\alpha\psi_x\lambda_v^x(1-\pi_x)}{\lambda_I^x}\left(\lambda_I^x x_I^2 + \lambda_I^z z_I^2\right) = 0$. *We will assume* $\lambda_I^x > \underline{\lambda}_I^x$ *and* $d_I < \frac{1}{4\alpha\psi_d\lambda_v^d}$ *for* $d \in \{x, z\}$.

We now characterize the voter retention rule in this general model. Because the voter cares about both issues, and expects both candidates to change policy on both X and Y, her beliefs about Z directly enter her decision rule. Formally, the voter reelects the incumbent if and only if:

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

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

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

which rearranged yields:

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

Using this result, we can see that even in this general model the incumbent is leading (i.e., the voter ex-ante prefers to re-elect him) if $\pi_x > \frac{1}{2}$, and trailing (i.e., the voter ex-ante prefers to oust him) otherwise.

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

Assumption 4. $\lambda_v^x > \lambda_v^z \frac{z_I}{x_I}$.

Notice that this implies that the secondary dimension is the one over which the voter has less intense preferences *and* has less information ex-ante. It seems substantively reasonable to assume that voters tend to be more informed over issues they care more about.

Next, we characterize the incumbent's probability of winning. Suppose that the incumbent is leading ($\pi_x > \frac{1}{2}$). Then, generalizing Definition 2, denote $\widehat{\rho}_g$ the value that solves:

$$\mu^x(\varnothing, -\alpha, \rho) = \widehat{\mu}_v^x(0), \tag{B-3}$$

where

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

which yields:

$$\widehat{\rho}_g = \frac{\big(1-\widehat{\mu}_v^x(0)\big)\pi_x}{\pi_x\big(1-2\widehat{\mu}_v^x(0)\big)+\widehat{\mu}_v^x(0)}. \tag{B-5}$$

Since the RHS in B-4 is decreasing in ρ, under $\rho > \widehat{\rho}_g$ we have $\mu^x(\varnothing, -\alpha, \rho) < \widehat{\mu}_v^x(0)$, and the leading incumbent is replaced if the voter observes an uninformative outcome on X, but learns that $z_v = -\alpha$.

If instead $\pi_x < \frac{1}{2}$, $\widehat{\rho}_g$ satisfies:

$$\mu^x(\varnothing, \alpha, \rho) = \widehat{\mu}_v^x(1), \tag{B-6}$$

where

$$\mu^x(\varnothing, \alpha, \rho) = \frac{\pi_x \rho}{\pi_x \rho + (1-\pi_x)(1-\rho)}. \tag{B-7}$$

Combining the above, we have

$$\widehat{\rho}_g = \frac{(1-\pi_x)\widehat{\mu}_v^x(1)}{\pi_x(1-2\widehat{\mu}_v^x(1))+\widehat{\mu}_v^x(1)}. \tag{B-8}$$

The RHS of B-7 is increasing in ρ, therefore under $\rho > \widehat{\rho}_g$ we have $\mu^x(\varnothing, \alpha, \rho) > \widehat{\mu}_v^x(1)$, and a trailing incumbent is re-elected if the voter observes an uninformative outcome on X, but learns that $z_v = \alpha$.

Then, denoting $\mathbb{P}(x_1, z_1)$ the probability of the incumbent being re-elected, as a function of the policy implemented on each dimension, we have:

- $\mathbb{P}(x_1, z_1) = 1 - \Pr(L_x(x_1) = 1)(1-\pi_x)$ when the incumbent is leading and $\rho < \widehat{\rho}_g$;
- $\mathbb{P}(x_1, z_1) = 1 - \Pr(L_x(x_1) = 1)(1-\pi_x) - \Big(1 - \Pr(L_x(x_1) = 1)\Big)\Pr(L_z(z_1) = 1)(1-\pi_z)$ when the incumbent is leading and $\rho > \widehat{\rho}_g$;
- $\mathbb{P}(x_1, z_1) = \Pr(L_x(x_1) = 1)\pi_x$ when the incumbent is trailing and $\rho < \widehat{\rho}_g$;
- $\mathbb{P}(x_1, z_1) = \Pr(L_x(x_1) = 1)\pi_x + \Big(1 - \Pr(L_x(x_1) = 1)\Big)\Pr(L_z(z_1) = 1)\pi_z$ when the incumbent is trailing and $\rho > \widehat{\rho}_g$.

Finally, given the probability of retention $\mathbb{P}(x_1, z_1)$, the Lagrangean associated with the incumbent's problem can be expressed as:

$$\mathcal{L}(x_1, z_1) = -\lambda_I^x(x_1-x_I)^2-\lambda_I^z(z_1-z_I)^2-[1-\mathbb{P}(x_1, z_1)]\,K+\chi_1(x_1)-\chi_2(x_1-x_I)+\chi_3(z_1)-\chi_4(z_1-z_I).$$

The optimization problem satisfies the constraint qualifications, hence we know that the solution of the incumbent's maximization problem must satisfy the following Karush-Kuhn-Tucker conditions:

$$-2\lambda_I^x(x_1 - x_I) + K\frac{\partial \mathbb{P}}{\partial x_1} + \chi_1 - \chi_2 = 0 \tag{B-9}$$

$$-2\lambda_I^z(z_1 - z_I) + K\frac{\partial \mathbb{P}}{\partial z_1} + \chi_3 - \chi_4 = 0 \tag{B-10}$$

$$x_1 \geq 0 \quad \wedge \quad \chi_1 x_1 = 0 \tag{B-11}$$

$$x_1 - x' \leq 0 \quad \wedge \quad \chi_2(x' - x_1) = 0 \tag{B-12}$$

$$z_1 \geq 0 \quad \wedge \quad \chi_3 z_1 = 0 \tag{B-13}$$

$$z_1 - z' \leq 0 \quad \wedge \quad \chi_4(z' - z_1) = 0 \tag{B-14}$$

$$\chi_1, \chi_2, \chi_3, \chi_4 \geq 0. \tag{B-15}$$

We can now use these preliminary results to characterize the equilibrium policies.

Characterization of Equilibrium Policies.

Remark B-1. *Let $\lambda_v^z, \lambda_I^z, \lambda_C^z > 0$. The incumbent always acts on X in equilibrium.*

Proof of Remark B-1. From the above, a necessary condition for an equilibrium with $x_1^* = 0$ is that $2\lambda_I^x x_I + K\frac{\partial \mathbb{P}}{\partial x_1} < 0$. First, suppose that ρ is low, $\rho < \widehat{\rho}_g$. Then, the problem is exactly identical to the unidimensional baseline and Assumption 3 implies $2\lambda_I^x x_I + K\frac{\partial \mathbb{P}}{\partial x_1} > 0$. Suppose instead ρ is high, $\rho > \widehat{\rho}_g$. Then, $\frac{\partial \mathbb{P}}{\partial x_1}$ is a function of z_1. For a leading incumbent, we have:

$$\frac{\partial \mathbb{P}}{\partial x_1} = 4\alpha\psi_x\lambda_v^x\Big[\pi_x - \Big(1-(1-\pi_z)4\alpha\psi_z z_1\lambda_v^z\Big)\Big] \tag{B-16}$$

Thus, our necessary condition for the incumbent to leave X unaddressed becomes:

$$2\lambda_I^x x_I + K4\alpha\psi_x\lambda_v^x\Big[\pi_x - \Big(1-(1-\pi_z)4\alpha\psi_z z_1\lambda_v^z\Big)\Big] < 0. \tag{B-17}$$

Notice that the LHS is increasing in z_1, and by Assumption 3 always strictly positive at $z_1 = 0$. Thus, the condition can never be satisfied.

If instead the incumbent is trailing, we have that:

$$\frac{\partial \mathbb{P}}{\partial x_1} = 4\alpha\psi_x\lambda_v^x\Big(\pi_x - \pi_z 4\alpha\psi_z\lambda_v^z z_1\Big), \tag{B-18}$$

and our necessary condition becomes:

$$2\lambda_I^x x_I + K4\alpha\psi_x\lambda_v^x\Big(\pi_x - \pi_z 4\alpha\psi_z\lambda_v^z z_1\Big) < 0. \tag{B-19}$$

Notice that the LHS is decreasing in z_1, and by Assumption 3 always strictly positive at $z_1 = z'$. Thus, the condition can never be satisfied. □

Proposition B-1. *Let* $\lambda_v^z, \lambda_I^z, \lambda_C^z > 0$.

1. *If the correlation across dimensions is low, then the incumbent always acts on the secondary dimension Z in equilibrium.*

2. *Suppose instead the correlation across dimensions is high.*

   - *A trailing incumbent always acts on the secondary dimension Z in equilibrium;*

   - *If the precision of the Z-dimension shock ε_z (ψ_z) is sufficiently low, then a leading incumbent acts on dimension Z. If instead ψ_z and λ_I^x are sufficiently high, then a leading incumbent does not act on Z.*

Proof of Proposition B-1. If $\rho < \widehat{\rho}$ then the incumbent's retention chances are not a function of his choice on the Z dimension, and $z^* = z_I$ whether the incumbent is leading or trailing.

Suppose instead that $\rho > \widehat{\rho}_g$. First, consider a trailing incumbent. From the KKT conditions, a necessary condition for an equilibrium with $z_1^* = 0$ is that $2\lambda_I^z z_I + K\frac{\partial \mathbb{P}}{\partial z_1} < 0$. Plugging in the value from of $\mathbb{P}$, this reduces to:

$$2\lambda_I^z z_I + K(1 - 4\alpha\psi_x\lambda_v^x x_1\pi_x)4\alpha\psi_z\lambda_v^z\pi_z < 0, \tag{B-20}$$

which can never be satisfied since $1 - 4\alpha\psi_x\lambda_v^x x_1\pi_x > 0$ in equilibrium. Thus a trailing incumbent must always set $z_1^* > 0$.

Finally, consider a leading incumbent. From the KKT conditions, a necesary condition for $z_1^* = 0$ is that

$$2\lambda_I^z z_I - (1 - \pi_z)(1 - 4\alpha\widehat{x}\psi_x\lambda_v^x)4\alpha\psi_z\lambda_v^z K < 0, \tag{B-21}$$

where $\widehat{x}$ is equal to:

$$\widehat{x} = x_I - \frac{2\alpha\psi_x\lambda_v^x(1 - \pi_x)K}{\lambda_I^x}. \tag{B-22}$$

We can immediately see that the LHS of (B-21) is continuous and decreasing in ψ_z, and never satisfied at $\psi_z = 0$. Thus, there exists a sufficiently low ψ_z that guarantees that the incumbent sets $z_1^* > 0$.

Next, notice that if ψ_z is at the upper bound $\frac{1}{4\alpha\lambda_v^z z_1}$, then (B-21) is satisfied for a sufficiently high λ_I^x (as the LHS is concave in λ_I^x). Suppose then ψ_z and λ_I^x are sufficiently high that (B-21) is satisfied. Then, a sufficient condition for $z_1^* = 0$ is that no other solution candidate satisfies the KKT conditions. From the previous result we know that $x_1^* > 0$. Furthermore, if x' then the the incumbent's retention probability is not a function of z_1, therefore it must be the case that $z_1 = z_I$. Finally, $\frac{\partial \mathbb{P}}{\partial z_1} < 0$, which implies that in equilibrium we must always have $z_1 \leq z_I$. This leaves us with three possible equilibrium candidates that may satisfy the KKT conditions: (x', z_I), $\Big(\widehat{x}(0), 0\Big)$ or $\Big(\widehat{x}(\widehat{z}), \widehat{z}(\widehat{x})\Big)$.

Consider first (x', z_I). For this candidate to satisfy the KKT conditions we need

$$-2\lambda_I^x(x' - x_I) + K4\alpha\psi_x\lambda_v^x\Big[\pi_x - \Big(1 - (1 - \pi_z)4\alpha\psi_z\lambda_v^z z_I\Big)\Big] > 0. \tag{B-23}$$

Straightforwardly, this can never be satisfied if λ_I^x is too high.

Consider instead $\left(\widehat{x}\left(\widehat{z}\right), \widehat{z}(\widehat{x})\right)$. For this candidate to be a possible solution, a necessary condition is that $\widehat{z}\left(\widehat{x}\right) > 0$ and $\widehat{x}\left(\widehat{z}\right) < x'$. This requires

$$\widehat{z} = z_I - \frac{2\alpha\psi_Z\lambda_v^z(1-\pi_z)K(1-4\alpha\psi_x\lambda_v^x\hat{x})}{\lambda_I^z} > 0. \tag{B-24}$$

In an interior solution $(1 - 4\alpha\psi_x\lambda_v^x\hat{x})$ must be strictly positive. Recall that K is continuously increasing in λ_I^x. Thus, again the condition fails for a sufficiently large λ_I^x. □

Proof of Proposition 3. If $\rho < \widehat{\rho}_g$ then the incumbent's strategic problem is identical to the unidimensional case on X. On Z, he simply implements his ideologically preferred policy since z_1 does not influence his retention chances.

Suppose instead, $\rho > \widehat{\rho}_g$. Denote by $\widehat{d}(\neg d)$ the possible interior solution on dimension d, given the policy choice on dimension $\neg d$. Further, recall that d_u is the optimal policy on dimension d in the unidimensional benchmark.

Consider first a trailing incumbent. From the previous results and inspection of the KKT conditions, we can verify that there are only four possible equilibrium candidates: (x', z_I), $\left(\widehat{x}(z'), z'\right)$, $\left(\widehat{x}\left(\widehat{z}\right), \widehat{z}(\widehat{x})\right)$ or (x', z'). As above, we can exclude the last case, since when $x = x'$ the incumbent's retention probability is not a function of z_1, therefore it must be the case that $z_1 = z_I$.

An inspection of the first-order conditions gives us that $\widehat{z}(\widehat{x}) > z_I$ (because $\frac{\partial \mathbb{P}}{\partial z_1} > 0$) and $\widehat{x}(z > 0) < x_u$ (since $\frac{\partial \mathbb{P}}{\partial x_1}$ is decreasing in z_1). Thus, a sufficient condition to ensure that $x_1^* \leq x_u$ is that $x_1^* = x'$ implies $x_u = x'$. A necessary condition for $x_1^* = x'$ is that the incumbent's utility is increasing in x_1 at $x_1 = x'$, given the optimal policy on the secondary dimension $\widehat{z}(x')$. Similarly, in the unidimensional world, a necessary and sufficient condition to ensure that $x_u = x'$ is that the incumbent's utility is increasing at $x = x'$. Then, the result follows from the fact that that if the incumbent's utility is increasing in x_1 at $x_1 = x'$ under $\lambda_v^z = 0$, then it must also be increasing under $\lambda_v^z > 0$:

$$-2\lambda_I^x(x' - x_I) + 4\alpha\psi_x\pi_x\lambda_v^x K \geq -2\lambda_I^x(x' - x_I) + 4\alpha\psi_x\lambda_v^x\Big[\pi_x - 4\alpha\psi_z\lambda_v^z z_1\pi_z\Big]K \tag{B-25}$$

which reduces to

$$\pi_x \geq \pi_x - 4\alpha\psi_z\lambda_v^z z_I\pi_z, \tag{B-26}$$

which is always satisfied.

Finally, consider a leading incumbent. As established in the proof of the previous result, there are only three possible equilibrium candidates: (x', z_I), $\Big(\widehat{x}(0), 0\Big)$ or $\Big(\widehat{x}(\widehat{z}), \widehat{z}(\widehat{x})\Big)$. An inspection of the first order conditions gives us that $\widehat{z}(\widehat{x}) \leq z_I < z_u$ (because $\frac{\partial\mathbb{P}}{\partial z_1} \leq 0$) and $\widehat{x}(z > 0) > x_u$ (because $\frac{\partial\mathbb{P}}{\partial x_1}$ is increasing in z_1). $\square$