Learning in a Complex World: How Multidimensionality Affects Policymaking

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

January 31, 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 dimensions, present critical challenges for reelection-seeking policymakers. We develop an accountability model to study these decisions. We address several questions: What motivates policymakers to adopt broad or narrow policy focuses? Which areas are neglected under narrow agendas? What types of reforms are pursued with different policy scopes? In the model, voters face uncertainty about their preferred policies, learn via experience, and may perceive different dimensions as correlated. We find that trailing incumbents embrace comprehensive policy programs, while leading ones prioritize fewer, independent dimensions. We also identify a substitution effect among correlated dimensions: policymakers shift towards extreme or moderate positions when expanding their policy agenda, depending on their electoral prospects. When issues are correlated, a unidimensional model is a poor representation of policymaking.

*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.

Historical examples show wide variation in the scope of policy programs officeholders pursue, as well as the nature of their decisions across various policy areas. Following a rocky start to his presidency, for example, President Harry Truman aggressively expanded his policymaking efforts in 1947, introducing a series of bills known as the "Fair Deal." This initiative included several progressive proposals, introduced amidst considerable Democratic losses in the 1946 midterms—and in spite of the fact that the President was not known to be especially progressive. In stark contrast to Truman's expansive Fair Deal, after cruising to victory in 1952, President Dwight Eisenhower focused his attention on the economy and spending, eschewing broader policy interventions on other issues. Furthermore, despite Republican majorities in Congress, he stuck to relatively moderate policies. Outside of the United States, French President Sarkozy adopted a moderate stance on economic issues while taking extreme positions on immigration, social issues, and foreign policy. In comparison, German Chancellor Angela Merkel implemented austere economic measures but maintained a more moderate approach on immigration and climate change.

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. 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? And finally, are narrower programs associated with more extreme, or more moderate, reforms?

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, Forthcoming), 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 policies are optimal for her, she may also become more inclined to believe that liberal policies on other issues, such as healthcare or immigration, are beneficial. In other words, voters can learn about how well a political program fits their preferences on one dimension by observing the policy outcome on another dimension.[1]

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 reelec-

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

tion. 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.

Using a similar logic, we show that the incumbent's electoral strength determines whether he faces electoral incentives to expand or contract the scope of policymaking in a multidimensional world (i.e., a world wherein the voter cares about multiple dimensions). Suppose the incumbent is trailing, and thus is motivated to promote voter learning. In equilibrium, his incentive to generate more information leads the incumbent to *expand* his policy agenda, even incorporating issues he has no ideological preferences over. This provides a potential rationale for Truman's decision to broaden his policy agenda after incurring losses in the 1946 midterm elections, after which he was widely viewed as a long shot for reelection in 1948. Due to learning spillovers this strategic behavior, we show, can extend even to issues that the voters may not prioritize. As long as these issues still provide valuable information about the more salient dimensions, secondary policy choices become strategically crucial for an incumbent who seeks reelection.

An electorally leading incumbent faces opposite incentives. As described above, this incumbent wishes to avoid generating fresh policy information that could potentially undermine his advantage.

This concern becomes particularly pronounced when dealing with issues that exhibit a strong correlation with the salient dimension, due to the possibility of learning spillovers that may erode even a strong lead. Consequently, a leading incumbent may opt to leave these policy matters unattended and focus *narrowly* on primary-dimension policy initiatives, as exemplified with Eisenhower's focus on the economy in the early years of his "Modern Republicanism."

Next, we study how multidimensionality influences the *nature* of the policies implemented in equilibrium on each issue. Interestingly, our central result uncovers a *strategic substitution effect* between dimensions that emerges as a consequence of learning spillovers when policy dimensions are sufficiently correlated.

Consider a trailing incumbent, whose need to generate information pushes him to expand his policy agenda to a secondary dimension. As described above, in the context of a unidimensional environment this incumbent consistently adopts extreme policies to encourage voter learning. In a multidimensional world, one might intuitively think that his goal of promoting voter learning would drive this trailing incumbent to pursue extreme policies across *all* available issue domains. Instead, we show that the possibility to expand the scope of his policy agenda can actually lead this incumbent to moderate on the primary dimension. This happens because by implementing moderate, and thus less informative, policies on the primary issue, the incumbent can amplify the impact of learning spillovers of extreme secondary policies. In turn, this is strategically valuable if favorable outcomes are more likely to arise on these secondary issues. In line with this logic, Sarkozy's 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 personal approval ratings plummeting in the months preceding his reelection campaign.[2] These extreme policy positions went in tandem with more centrist stances on the primary economic dimension.[3]

Symmetric results hold when the incumbent is electorally advantaged. In contrast with intuition

[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.

[3]Despite a somewhat different rhetoric, 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.

from the unidimensional case, a leading incumbent adopts more *extreme* stances on the primary dimension when his ideological preferences push him to broaden his policy initiatives to include other issues. This happens because extreme primary policies are more likely to reveal favorable information, thus counteracting damaging learning spillovers from the secondary policy, where the incumbent tends to opt for more moderate policies. This logic is in line with the example of Angela Merkel's extreme austerity program coupled with more centrist stances on immigration (De La Baume, 2017) or climate change (Schwagerl, 2011).[4]

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 in the minds of voters, policymaking in a multidimensional world takes on a significantly distinct character due to the substitution effect described above. To truly understand policymaking, one must account for its inherent multidimensionality.

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,

[4] 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.

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 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

[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.

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 platforms 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.

The learning technology we use relates more closely to the models introduced in Ashworth, Bueno de Mesquita and Friedenberg (2017) and Izzo (2022), 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 representative 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.

Information. Both the incumbent and the challenger's ideal points are common knowledge. 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

[6] In this multidimensional setting, we can think about V as representative of a group of undecided voters who are, with some probability, electorally pivotal.

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 realized 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 4.1.1 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.

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. The incumbent does not have privileged information about what policy is optimal for the voter (or private information about his own ideological preferences). This allows us to assume away the possibility that the incumbent's policy choice directly provides information to the voter and instead, following the literature on retrospective evaluations, focus on what the voter learns from her lived experiences (i.e., the inferences she draws upon observing realized outcomes).

3.1 Discussion of the Model

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 a representative 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 objective 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 multidimensional 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,[7] and study how they influence the inferences that guide voters' electoral decisions—and, in turn, policymakers' strategic incentives in a multidimensional world.

4 Equilibrium Analysis

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

To reduce the number of cases under consideration, we will impose the following assumption:

Assumption 1. $d_I = -d_C > 0$ *and* $\lambda_I^z = 0$ *if and only if* $\lambda_C^z = 0$.

The assumption on λ_I^z and λ_C^z implies that the voter expects either both candidates, or neither of them, to reform policy on dimension Z in the *second* period.[8]

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 recall that politicians' ideal points are symmetric around zero on each

[7] 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.

[8] This in turns guarantees that the location of the first-period policy does not directly enter the voter retention strategy (as it instead would in our endogenous status quo setting in a case where only one of the two candidates is expected to reform this dimension in the second period). We purposefully shut down this channel as it is orthogonal to our logic and our argument on risk taking and policy experimentation.

dimension. Then, the following holds:

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

$$\mu^x > \frac{1}{2} - \mathbb{I}_z \frac{\lambda_v^z z_I}{\lambda_v^x x_I} \frac{2\mu^z - 1}{2} \equiv \widehat{\mu}^x(\mu^z), \tag{5}$$

where $\mathbb{I}_z = 0$ *if* $\lambda_I^z = \lambda_C^z = 0$ *and* $\mathbb{I}_z = 1$ *otherwise.*

Proof. All Proofs are collected in the Appendix. □

When the voter only cares about the primary dimension ($\lambda_v^z = 0$), it follows from (5) that the right-wing incumbent is reelected as long as the voter believes her optimal policy on dimension X is more likely to be a right-wing one ($\widehat{\mu}^x = \frac{1}{2}$). The same holds if the voter cares about both dimensions, but politicians only care about the primary one (i.e., $\lambda_I^z = \lambda_C^z = 0$), as the voter expects both the challenger and the incumbent to remain inactive on dimension Z in the second period.

Instead, when the voter and politicians care about both dimensions ($\lambda_v^z, \lambda_I^z, \lambda_C^z > 0$), the voter becomes more lenient with the incumbent on dimension X the more she likes him on dimension Z (and vice-versa). This effect is stronger the more (less) polarized candidates are on the secondary (primary) policy dimension.[9]

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 believes 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 2. $\lambda_v^x > \lambda_v^z \frac{z_I}{x_I}$.

Before continuing with the analysis, let us introduce some useful definitions. By plugging $\mu^x = \pi_x$ and $\mu^z = \pi_z = \pi_x\rho + (1 - \pi_x)(1 - \rho)$ into (5), we can verify that at $\pi_x = \frac{1}{2}$ the voter is ex-ante indifferent between retaining the right-wing incumbent and replacing him with the challenger. For any $\pi_x > \frac{1}{2}$ the voter ex-ante prefers the incumbent, and for $\pi_x < \frac{1}{2}$ she instead prefers the challenger. Thus, we will say that:

[9] Given our symmetry assumption $d_I = -d_c$, polarization on dimension D here is given simply by $2d_I$.

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

4.1 Voter Learning

In our framework, the voter gains insights into her preferred policies for each dimension through her real-life experiences (as in Izzo Forthcoming). However, these experiences only represent a noisy signal of the genuine alignment between the voter's interests and the implemented policy, and this complicates the voter's inference problem. Furthermore, when policies span multiple, correlated dimensions the 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. In what follows, we fully characterize these processes of direct and indirect learning in our framework.

4.1.1 Direct Learning

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 2. *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 amount of learning is a function of the implemented policy d_1, as more extreme policies increase the probability that $\tilde{\mu}^d \neq \pi_d$. Furthermore, there exists a unique policy d' such that if $|d_1| \geq |d'|$, then $\tilde{\mu}^d \neq \pi_d$ with probability 1.

Lemma 2 shows that, upon observing outcomes on each 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.

The logic behind this result is as follows. Suppose the incumbent acts on the primary dimension, perhaps by lowering taxes. In expectation the voter's payoff from this policy is different under the two states of the world, i.e., the two possible values of her optimal platform. Lowering taxes may be good for the voter, because it may spur economic growth, or bad, because it reduces redistribution and welfare spending. However, the voter's realized utility is also a function of a random period-specific shock $\varepsilon_{d,1}$—say, random fluctuations in the economy. This, in turn, creates a partial overlap in the support of the payoff *realization*.

Thus, when the policy is sufficiently moderate ($d_1 \in (-d', d')$), there exists a range of payoffs that may be realized (i.e., be actually observed) whether the voter's true bliss point takes a positive *or* a negative value.[10] If the payoff realization falls outside this range, it constitutes a fully informative signal. There is only one state of the world that could have generated that specific realization: the observed consequences of the policy are simply too good, or too bad, for this to be justified as a consequence of the shock. Conversely, any payoff realization that falls inside the range of overlap is completely uninformative. Since the shock is uniformly distributed, any such realization has exactly the same probability of being observed under the two states of the world. Thus, the voter learns nothing and her interim posterior remains at her prior beliefs. As the implemented policy becomes more extreme, the range of overlap becomes smaller, and the voter is more likely to directly learn her true preferences.

In other words, 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. As a consequence, the voter is likely able to learn whether the policy was moved in the optimal

[10]In the Appendix, we derive the value of the threshold d': $x' = \frac{1}{4\alpha\psi_x\lambda_v^x}$ and $z' = \frac{1}{4\alpha\psi_z\lambda_v^z}$.

direction or not. In contrast, if the incumbent adjusts tax policy modestly, observed differences in outcomes may plausibly be attributed to random shocks and do not provide strong information about the desirability of the incumbent's policies. The voter is then unable to learn.

For a similar logic, if the policy remains at the status quo then the voter can never learn anything new by observing the policy outcome. Formally, there is full overlap in the support of the payoff realizations, therefore the realized outcome is always uninformative.[11]

Figure 1 provides a graphical illustration of the results in Lemma 2. The blue and orange functions represent the conditional outcome distributions (i.e., the distributions of the voter's realized utility), under a positive and a negative state of the world, respectively. In the left panel, a moderate right-wing policy produces a partial overlap in the conditional distributions. 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.

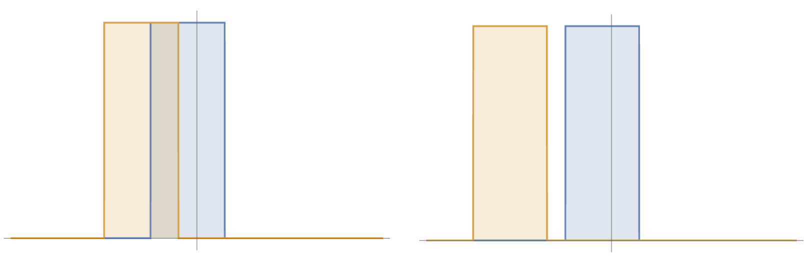

Figure 1: **Voter Learning.** The two plots display the realized voter utility on dimension D under a positive (blue function) and negative (orange function) state of the world. The policy extremism increases from left to right.

We note that the assumption of uniformly distributed shocks simplifies the analysis by generating the stark learning environment described above. However, the result that extreme (new) policies facilitate voter learning holds more generally, as it simply requires that the noise distribution satisfies the monotone likelihood ratio property.[12] Furthermore, our central results on when

[11]Notice that this result follows from our assumption that $d_{sq} = 0$. We use this normalization to simplify notation, but our qualitative results below simply require that if the policy remains at the status quo, then no new (direct) information is generated on that dimension. For example, we could assume that if $d_1 = d_{sq}$, then the voter does not observe a new realization of her utility on d.

[12]For example, Bils and Izzo (2022) show that this result holds under normally distributed shocks. There, every

the incumbent expands the scope of policymaking, and the strategic substitution effect described in the introduction, do not require that more extreme policies are more informative. These results simply require that moving closer to the status quo, or away from it, influences the amount of voter learning.

4.1.2 Indirect Learning

Lemma 2 tells us that the voter can directly learn her best policy on each dimension by observing how much she liked or disliked the policy implemented on that issue. Next, we show that such direct learning is not the only way the voter can gain new information. Due to the correlation between dimensions in our model, learning can spill over from one dimension to the other. Thus, the voter can learn something about her optimal policy on the primary dimension if she observes an informative outcome on the secondary dimension, and vice versa. In other words, the connections between dimensions allow for a form of *indirect* learning, where insights gained in one dimension can inform preferences in another.

The result below shows how the voter's posterior belief on X depends on the outcome of the secondary dimension Z (symmetric results apply to the posterior on Z as a function of the outcome on X). Recall that $\tilde{\mu}^x$ is the voter's *interim* beliefs, as a function only of her realized utility on X. Instead, we denote μ^x the voter's final posterior on x_v, as a function of her realized utility on *both* dimensions X and Z. Then, we have:

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

$$\mu^x = \tilde{\mu}^x. \tag{6}$$

Suppose instead that the voter learns that $z_v = \alpha$. Then:

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

outcome is somewhat informative, but never fully so. Nonetheless, extreme policies continue to facilitate voter learning and thus decrease the variance in the posterior distribution.

Finally, if the voter learns $z_v = -\alpha$, we have:

$$\mu^x(\tilde{\mu}^x, -\alpha, \rho) = \frac{\tilde{\mu}^x(1-\rho)}{\tilde{\mu}^x(1-\rho) + (1-\tilde{\mu}^x)\rho} \leq \tilde{\mu}^x. \tag{8}$$

The proof simply follows by applying Bayes rule, and is therefore omitted. 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 from Lemma 3 we have $\mu^x = \tilde{\mu}^x$. Instead, when no direct learning occurs on X (i.e., $\tilde{\mu}^x = \pi_x$), learning spillovers determine the voter's posterior. If the voter learns that her ideal point on Z is a right-wing (left-wing) one, she becomes more convinced that her optimal policy on X is right-wing (left-wing) as well. More substantively, 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. Moreover, the higher the correlation across dimensions ρ, the stronger these learning spillovers.

5 The Incumbent's Problem

The findings in the previous sections shed light on how the incumbent's decisions in our framework impact his expected payoff. There are two key effects at play here. 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 affect his retention chances and thus his expected second-period payoff. This, in turn, depends on voter learning. This information effect operates through the two previously discussed channels. 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. This trade-off clearly appears in the incumbent maximization problem, which we can express as follows:

$$\max_{x_1, z_1} \; -\lambda_I^x(x_1 - x_I)^2 - \lambda_I^z(z_1 - z_I)^2 - (1 - \mathbb{P}(x_1, z_1)) \left[\lambda_I^x(x_I - x_C)^2 + \lambda_I^z(z_I - z_C)^2\right], \tag{9}$$

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

Recall that, from Lemma 2, 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 (formally, the sign of $\frac{\partial \mathbb{P}(x_1,z_1)}{\partial x_1}$ and $\frac{\partial \mathbb{P}(x_1,z_1)}{\partial z_1}$), the incumbent will have incentives to distort his choice either to the extreme or towards the status quo $d_{sq} = 0$.[13] In the following analysis, we will see that the 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 assumptions:

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 $\lambda_I^x > \underline{\lambda}_I^x$, $z_I < z'$ and $x_I < x'$.*

Recall that d' is smallest policy that guarantees direct learning on dimension d. These assump-

[13] This is clear from the first-order necessary conditions for an interior maximum, which are respectively:

$$(x_1) \quad -2\lambda_I^x(x_1 - x_I) + \frac{\partial \mathbb{P}(x_1, z_1)}{\partial x_1}\left(\lambda_I^x(x_I - x_C)^2 + \lambda_I^z(z_I - z_C)^2\right) = 0, \tag{10}$$

$$(z_1) \quad -2\lambda_I^z(z_1 - z_I) + \frac{\partial \mathbb{P}(x_1, z_1)}{\partial z_1}\left(\lambda_I^x(x_I - x_C)^2 + \lambda_I^z(z_I - z_C)^2\right) = 0. \tag{11}$$

tions then guarantee that in a unidimensional world, that is, a world where the voter only cares about the primary dimension, the incumbent always chooses to act on this dimension. Furthermore, the incumbent never finds it optimal to adopt his statically preferred policy on dimension d if that is the only electorally relevant one.

5.1 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. In our setting, this also implies that there are no learning spillovers across dimensions.[14] Consequently, the incumbent's retention probability is only a function of his policy on the primary dimension. Recall that we define an outcome on dimension D as informative if it induces a (interim) posterior $\tilde{\mu}^d \neq \pi_d$. Further, an informative outcome is favorable to the right-wing incumbent if it induces $\tilde{\mu}^d = 1$, and unfavorable if $\tilde{\mu}^d = 0$. We have:

Remark 1. *Suppose that $\lambda_v^z = 0$. Then:*

(i) a leading incumbent is reelected unless the outcome on X is informative and unfavorable, i.e., the voter learns that $x_v = -\alpha$;

(ii) a trailing incumbent is reelected if and only if the outcome on X is informative and favorable, i.e., the voter learns that $x_v = \alpha$.

Suppose the incumbent is ex-ante *trailing*. If the voter receives no new information, she will choose to oust him. This right-wing incumbent is then reelected if and only if the voter observes an informative policy outcome and learns that right-wing policies are optimal for her. In contrast, a *leading* incumbent can only be damaged by information. If the voter learns nothing new, this incumbent will be reelected for sure.

Denote by d_u the incumbent's optimal policy on dimension d when the incumbent's strategic problem is unidimensional. Applying Remark 1 to the incumbent's maximization problem, we obtain:

[14] This follows from the fact that, if $\lambda_v^z = 0$, the voter's observed payoff on dimension Z is pure noise.

Proposition 1. *Suppose $\lambda_v^z = 0$. In equilibrium, the incumbent always implement his bliss point on the secondary dimension Z ($z_u = z_I$). 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, $x_u > x_I$.

When the voter cares solely about dimension X, the incumbent's policy choice on Z is inconsequential for his retention chances. As such, the incumbent simply consider his static payoff and implements his ideologically preferred policy on this dimension. 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. Recall that more extreme policies facilitate voter learning. Then, a trailing 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 wants to avoid risks, and distorts policy towards 0 so as to minimize information. Notice that, since any pair of policies x and $-x$ induces the same amount of learning (Lemma 2), the right-wing incumbent never implements a policy to the left of 0.

Having characterized equilibrium policy in this unidimensional benchmark, we now move to analyzing the incumbent's policy choices in the multidimensional case, i.e., when $\lambda_v^z > 0$. Our objective is 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.

5.2 Multidimensional World

When the voter cares about multiple dimensions ($\lambda_v^z > 0$), the incumbent's strategic problem becomes multidimensional as well. To begin, we characterize the incumbent's probability of winning in this multidimensional world. For this purpose, it is useful to introduce the following:

Definition 2. *Let*

$$\widehat{\rho} \equiv \begin{cases} \frac{\left(1-\widehat{\mu}_v^x(0)\right)\pi_x}{\pi_x\left(1-2\widehat{\mu}_v^x(0)\right)+\widehat{\mu}_v^x(0)} & \text{if } \pi > \frac{1}{2} \\ \frac{\left(1-\pi_x\right)\widehat{\mu}_v^x(1)}{\pi_x\left(1-2\widehat{\mu}_v^x(1)\right)+\widehat{\mu}_v^x(1)} & \text{if } \pi < \frac{1}{2} \end{cases}$$

where $\widehat{\mu}^x(0)$ *and* $\widehat{\mu}^x(1)$ *are as defined in Lemma 1.*

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

Recall that, by Assumption 2, if the voter knew the location of the optimal policy on the primary policy X, the secondary policy Z would be electorally irrelevant. Our next result instead highlights that, under uncertainty, when the correlation across dimensions is high even the secondary issue that the voter cares relatively little about may have a crucial impact on her electoral choice.

Lemma 4. *Let* $\lambda_v^z > 0$*.*

1. *If the correlation across dimensions is low, then the incumbent's choice on the secondary dimension* Z *is electorally irrelevant and his probability of winning is the same as in the unidimensional benchmark.*

2. *Suppose instead the correlation across dimensions is high. The incumbent's retention probability is a function of his policy choice on both dimensions. Specifically:*

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

To understand the results from Lemma 4, suppose first that π_x is relatively high and $\rho = \frac{1}{2}$. Recall that the voter cares more about the primary dimension X compared to the secondary one, Z. For instance, the voter's concerns about the economy outweigh those about healthcare. Then, the absence of correlation between these dimensions implies that policy outcomes on the secondary dimension hold little electoral significance. Under $\rho = \frac{1}{2}$, learning spillovers can never emerge. Further, even if the voter observes an uninformative outcome on the economy but updates against the incumbent on healthcare, her prior on the primary dimension is sufficiently high that she still prefers reelecting the incumbent. Similarly, a positive outcome on the secondary issue alone is insufficient to resurrect a trailing incumbent.

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.[15] 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.

5.2.1 Policymaking in a Multidimensional World

Building on Lemma 4, we now characterize the incumbent's policy choice in this multidimensional world. To more clearly illustrate the policymaker's strategic incentives, we begin by analyzing a special case where the candidates only care about the primary dimension, $\lambda_I^z = \lambda_C^z = 0$. Under this assumption, even though she cares about both dimensions, the voter's retention decision does

[15]Notice that, when π_x is very close to $\frac{1}{2}$, the critical value $\widehat{\rho}$ may actually be lower than $\frac{1}{2}$.

not *directly* depend on her beliefs over Z, since she anticipates that neither I nor C will act on Z in the second period.[16] More specifically, the voter's optimal retention rule is exactly the same as in the unidimensional case: she retains the right-wing incumbent if and only if $\mu^x > \frac{1}{2}$. However, by Lemma 3, the voter's posterior on X is a function of her realized utility on Z. Therefore, even though the incumbent's ideological preferences are unidimensional, his strategic problem is *multidimensional*. This assumption thus allow us to isolate the strategic incentives emerging solely due to the learning spillovers across dimensions. In section 5.2.4, we complete the analysis allowing for $\lambda_I^z, \lambda_C^z > 0$.

5.2.2 The Dimensionality of Policymaking

First, we note that, as in the benchmark model, the incumbent always acts on the primary dimension X:

Remark 2. *The incumbent always acts on X in equilibrium, $x_1^* \neq x_{sq}$.*

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 4, characterizes the conditions under which the incumbent chooses to open this secondary policy dimension, despite not having ideological preferences to do so.

Corollary 1. *Suppose $\lambda_v^z > 0$ and $\lambda_I^z = \lambda_C^z = 0$. Then, the incumbent chooses to act on Z (i.e., $z_1 \neq z_{sq}$) 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 small (according to Definition 2). In this case, a trailing incumbent is indifferent between acting on the secondary dimension and keeping

[16] Recall that we assume that when indifferent the incumbent chooses not to act on Z.

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.

5.2.3 Multidimensionality and Extremism: the Substitution Effect

Next, we study how the multidimensionality of voter's preferences influences the nature of the policies pursued by the incumbent. The first result follows straightforwardly from the above discussion and Lemma 4:

Corollary 2. *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, he always implements a fully informative policy on this dimension, i.e., $z_1^* \geq z'$.*

Even though the incumbent does not have ideological preferences over dimension Z, his strategic incentives to facilitate voter learning induce policy extremism on this secondary dimension. Thus, in equilibrium, we either observe inaction on the secondary dimension (when ρ is low or the incumbent is leading), or we observe the incumbent pursuing extreme policies on this dimension (when ρ is high and the incumbent is trailing). In our setting, the incumbent's incentives to control information imply that extreme policymaking need not follow from extreme ideological preferences; in the multidimensional case, extreme policymaking can emerge on issues over which policymakers have no ideological preferences at all.

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

Proposition 2 (**Substitution Effect**). *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, his policy choice on the primary dimension x_1^* satisfies $x_1^* < x_I < x_u$.*

When a trailing incumbent cannot exploit the secondary dimension (i.e., $\lambda_v^z = 0$ or $\rho < \widehat{\rho}$) he always has strategic incentives to gamble on the primary one. Thus, this incumbent implements a policy more extreme than his ideological preference, $x_u > x_I$. Naïve intuition would suggest that in a multidimensional world, this incumbent should have incentives to pursue extreme policies on all issues. Instead, Proposition 2 highlights that a high correlation generates a *strategic substitution effect* between policy dimensions. This correlation allows the trailing incumbent to exploit learning spillovers, inducing *moderation* on the primary dimension: $x_1^* < x_I$. Notice that 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 better understand why, 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—pushing the voter to learn on the secondary one.

Notice that the above discussion implies that *if* the incumbent chooses to gamble on Z, then he prefers avoiding risks on X to magnify the impact of the learning spillovers. However, in principle, the incumbent may sometimes 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, the 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 policies on immigration. Recall 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, when $\pi_x < \frac{1}{2}$ we have that $\pi_z > \pi_x$. 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, the trailing right-wing incumbent prefers to gamble on Z, hoping to exploit a false positive—that is, generate a favorable outcome on Z and thus induce the voter to positively update on x_v as well, even when the optimal policy on the primary dimension is actually a left-wing one.

Thus, given these dynamics, in equilibrium the incumbent will never gamble on both dimensions. Rather, if the correlation is too low to exploit the learning spillovers, he will have no strategic incentives to act on Z and will continue gambling on X. If instead the correlation is high, he will gamble on Z but prefer avoiding risky choices on X.

5.2.4 General Model

The results of the previous section are useful to isolate the strategic incentives generated by the learning spillovers. Here, we complete the analysis by studying the incumbent's policy choice on each dimension in the general model, where both the voter and the politicians care about all issue areas (i.e., $\lambda_i^z > 0$ for $i \in \{I, V, C\}$). In contrast to the analysis presented above, the incumbent now has both strategic *and ideological* preferences over the secondary dimension Z. As above, our goal is to characterize the conditions under which the incumbent chooses to open the secondary dimension, and to examine how this influences his policy on the primary one.

First, we establish that, as in the benchmark cases analyzed above, the incumbent always acts on the primary dimension in equilibrium:

Remark 3. *The incumbent always acts on X in equilibrium: $x_1^* \neq x_{sq}$.*

Next, we characterize conditions under which the incumbent chooses to act on the secondary dimension as well:

Proposition 3.

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.*

A trailing incumbent has *both* ideological and strategic reasons to act on the secondary dimension. In equilibrium, he will therefore always choose to do so. By contrast, a leading incumbent faces a trade-off. On the one hand, he would like to implement his preferred policy on the secondary dimension. On the other hand, as the results of the previous section demonstrate, acting on the secondary dimension may hurt his reelection chances—and thus his expected future payoff.

If ρ is low, this tradeoff does not bite: the leading incumbent can survive reelection even if the outcome on the secondary dimension reveals damaging information. He can therefore implement his preferred policy on the secondary dimension while avoiding the negative electoral consequences.

Instead, if ρ is high, then a leading incumbent must balance dynamic electoral considerations and static ideological preferences. Proposition 3 identifies sufficient conditions under which the incumbent does and does not act on the secondary dimension. If ψ_z is sufficiently low, then the informativeness of policies on Z (i.e., the probability that a policy $z_1 \neq z_{sq}$ generates an informative outcome) is low. The incumbent can then enact a policy close to his static ideal without suffering damaging electoral consequences. If instead ψ_z is sufficiently high, and the incumbent has sufficiently strong preferences on the primary dimension, then strategic electoral considerations dominate and the incumbent prefers to avoid risks keeping the policy at the status quo $z_1^* = z_{sq}$.

In short, a trailing incumbent has incentives to expand the scope of policymaking, even incorporating correlated dimensions he has no ideological reasons to act on. In contrast, a strong between-dimension correlation can encourage a leading incumbent to contract the scope of policy-making, even inducing him to abandon issues he cares about.

Generalizing our central result from the previous section, Proposition 4 then highlights that the substitution effect described in Proposition 2 continues to emerge when the incumbent cares about both policy dimensions. As above, a trailing incumbent becomes more moderate on the primary dimension when he can act on a secondary one. The result is *reversed* for a leading incumbent: here, the strategic importance of dimension Z generates more *extremism* on X. Once more, this result shows how extending the unidimensional baseline to multiple policy issues fundamentally alters the nature of policymaking.

Proposition 4. *Suppose that the incumbent chooses to act on the secondary dimension in equilibrium. When the correlation with the primary dimension is low, we have $z_1^* = z_u = z_I$ and $x_1^* = x_u$. Suppose instead the correlation is high. Then:*

- *when the incumbent is trailing, we have $z_1^* > z_u = z_I$ and $x_1^* < x_u$;*
- *when the incumbent is leading, we have $z_1^* < z_u = z_I$ 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 equilibrium policy on *both* issues 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 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 tastes induce this incumbent to act on the secondary dimension, he undertakes more electoral risk. He may in fact generate an informative and unfavorable outcome on Z and 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. Finally, recall that given Assumption 2, direct learning on X renders outcomes on Z electorally irrelevant. Together, these observations imply that in order to counteract the detrimental effects of learning on the secondary dimension, the leading incumbent has incentives to *facilitate* direct learning on the primary one. In other words, a leading incumbent—who always operates in a risk-averse (moderate) fashion in a unidimensional world—becomes more risk-loving on the primary dimension when he chooses to expand the scope of policymaking.

Before concluding, we study how the incumbent's ideological preferences influence the kinds of dimensions he chooses to pursue in equilibrium. Suppose that multiple secondary dimensions are available, but the incumbent is resource constrained, such that he cannot act on all available dimensions. Then,

Proposition 5. *All else equal:*

- *a leading incumbent is more likely (in the sense of set inclusion) to open dimensions for which he is more moderate;*
- *a trailing incumbent is more likely to open dimensions for which he is more extreme.*

The intuition is as follows. Consider first a leading incumbent. If he chooses to act on a secondary policy dimension, he knows that moderate policies are less electorally risky, as they are less apt to generate negative information. As such, if he selects a dimension for which his ideological preferences are moderate, ideological costs are also kept at a minimum. The opposite holds for a trailing incumbent, who instead tends to pursue dimensions he is more extreme over.

6 Conclusions

As our analysis underscores, the introduction of multidimensionality within our accountability setting dramatically influences the incentives that incumbents face, as they make decisions about whether and how to change policy.

For a trailing incumbent, the possibility of policymaking in multiple *correlated* 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. Further, we show that if given the opportunity to select a dimension for expansion, the trailing incumbent prefers dimensions for which his preferences are ideologically extreme. For a leading incumbent the opposite holds: While he pursues moderate policy in a unidimensional world, in a multidimensional one he chooses 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 is fundamentally different from the unidimensional case. As we have emphasized throughout the paper, when the correlation across dimensions is sufficiently strong, the multidimensional problem is more than just the 'sum' of multiple unidimensional problems. If the correlation between different issues is high, voter learning spills over across dimensions, altering the politician's willingness to pursue policy change. Then, a model where policy making is unidimensional does not do a good job in capturing the incentives and nature of policymaking in a multidimensional world.

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.

Bils, Peter and Federica Izzo. 2022. "Policymaking in Times of Crisis.".

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.

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.

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

Izzo, Federica. 2022. "With friends like these, who needs enemies?" *London School of Economics working paper. Available at https://www. federicaizzo. com/pdf/WFLT_ 2018_08_09. pdf* .

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

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.

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.

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.

Appendix

Main Results - Proofs

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

$$\begin{aligned} &-\lambda_v^x[\mu^x(x_I-\alpha)^2+(1-\mu^x)(x_I+\alpha)^2]-\mathbb{I}_z\Big\{\lambda_v^z[\mu^z(z_I-\alpha)^2+(1-\mu^z)(z_I+\alpha)^2]\Big\}> \\ &-\lambda_v^x[\mu^x(x_C-\alpha)^2+(1-\mu^x)(x_C+\alpha)^2]-\mathbb{I}_z\Big\{\lambda_v^z[\mu^z(z_C-\alpha)^2+(1-\mu^z)(z_C+\alpha)^2]\Big\}. \end{aligned} \tag{12}$$

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

$$2\lambda_v^x\mu^x x_I\alpha - \lambda_v^x x_I\alpha + \mathbb{I}_z\Big(2\lambda_v^z\mu^z z_I\alpha - \lambda_v^z z_I\alpha\Big) > 0$$

which rearranged yields:

$$\mu^x > \frac{1}{2} - \mathbb{I}_z\frac{\lambda_v^z z_I}{\lambda_v^x x_I}\frac{(2\mu_v^z-1)}{2} \equiv \widehat{\mu}^x(\mu^z). \tag{13}$$

□

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

(i) A possible payoff realization for V given the incumbent's choice (x_t), and conditioning on the true state x_v has to fall within:

$$\left[-\lambda_v^x(x_t-x_v)^2-\frac{1}{2\psi_x}, -\lambda_v^x(x_t-x_v)^2+\frac{1}{2\psi_x}\right]. \tag{14}$$

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 (14) is empty. This requires:

$$-\lambda_v^x(x_t+\alpha)^2 + \frac{1}{2\psi_x} + \lambda_v^x(x_t-\alpha)^2 + \frac{1}{2\psi_x} \leq 0 \tag{15}$$

which rearranged yields:

$$x_t \geq \frac{1}{4\alpha\lambda_v^x\psi_x}. \tag{16}$$

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

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

Since the two probabilities are symmetric, we have

$$\begin{aligned}\Pr(L_x = 1 | \pi_x, 0 < x_t < x') =& \Pr\left(-\lambda_v^x(x_t-\alpha)^2 + \varepsilon_{x,t} > -\lambda_v^x(x_t+\alpha)^2 + \frac{1}{2\psi_x}\right) \\ =& \Pr\left(\varepsilon_{x,t} < 4\lambda_v^x\alpha x_t - \frac{1}{2\psi_x}\right) \\ =& 4\alpha x_t\lambda_v^x\psi_x. \end{aligned} \tag{17}$$

The proof for dimension Z is analogous therefore omitted. □

Proof of Remark 1. By Lemma 1, if $\lambda_v^z = 0$ 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 ($\pi_x < 1/2$). 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. □

Proof of Proposition 1. First, notice that given the assumption that $x_I < x'$ 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 ($\pi_x \geq \frac{1}{2}$). Using Remark 1 we then have that $\mathbb{P}(x_1, z_1) = 1 - (1 - \pi_x) \Pr(L_x = 1 | \pi_x, 0 < x_t \leq x')$. Using the proof of Lemma 2, we can substitute the value of this probability into the incumbent's maximization problem, which becomes:

$$- \lambda_I^x (x_1 - x_I)^2 - 4\alpha\psi_x\lambda_v^x x_1(1 - \pi_x)\Big(\lambda_I^x(x_I - x_C)^2 + \lambda_I^z(z_I - z_C)^2\Big). \tag{18}$$

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:

$$- 2\lambda_I^x(x_1 - x_I) - 4\alpha\psi_x\lambda_v^x(1 - \pi_x)K = 0. \tag{19}$$

Rearranging (19) 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 3 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

$$- \lambda_I^x(x_1 - x_I)^2 - (1 - 4\alpha\psi_x\lambda_v^x x_1\pi_x)K, \tag{20}$$

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

$$x_1 = \min\left\{x', x_I + \frac{2\alpha\psi_x\pi_x\lambda_v^x}{\lambda_I^x}K\right\} > x_I. \tag{21}$$

$\square$

Proof of Lemma 4. Suppose that the incumbent is leading ($\pi_x > \frac{1}{2}$). Then, $\widehat{\rho}$ solves:

$$\mu^x(\varnothing, -\alpha, \rho) = \widehat{\mu}_v^x(0), \tag{22}$$

where

$$\mu^x(\varnothing, -\alpha, \rho) = \frac{(1-\rho)\pi_x}{(1-\rho)\pi_x + \rho(1-\pi_x)}, \tag{23}$$

which yields:

$$\widehat{\rho} = \frac{\big(1 - \widehat{\mu}_v^x(0)\big)\pi_x}{\pi_x\big(1 - 2\widehat{\mu}_v^x(0)\big) + \widehat{\mu}_v^x(0)}. \tag{24}$$

Since the RHS in 23 is decreasing in ρ, under $\rho > \widehat{\rho}$ 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 $\pi_x < \frac{1}{2}$, $\widehat{\rho}$ satisfies:

$$\mu^x(\varnothing, \alpha, \rho) = \widehat{\mu}_v^x(1), \tag{25}$$

where

$$\mu^x(\varnothing, \alpha, \rho) = \frac{\pi_x\rho}{\pi_x\rho + (1-\pi_x)(1-\rho)}. \tag{26}$$

Combining the above, we have

$$\widehat{\rho} = \frac{(1-\pi_x)\widehat{\mu}_v^x(1)}{\pi_x(1-2\widehat{\mu}_v^x(1))+\widehat{\mu}_v^x(1)}. \tag{27}$$

The RHS of 26 is increasing in ρ, therefore under $\rho > \widehat{\rho}$ 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 $z_v = \alpha$.

Thus, we have:

- $\mathbb{P}(x_1, z_1) = 1 - \Pr(L_x(x_1) = 1)(1-\pi_x)$ when the incumbent is leading and $\rho < \widehat{\rho}$;
- $\mathbb{P}(x_1, z_1) = 1 - \Pr(L_x(x_1) = 1)(1-\pi_x) - \Big(1 - \Pr(L_x(x_1) = 1)\Big)\Pr(L_z(z_1) = 1)(1-\pi_z)$ 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 4 and the assumption that when indifferent the incumbent prefers not to act on Z. □

Proof of Corollary 2. 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 $z_1^* \geq z'$. □

Proof of Proposition 2. 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

$$\mathbb{P} = 4\alpha\psi_x\lambda_v^x x_1\pi_x + (1 - 4\alpha\psi_x\lambda_v^x x_1)p^\dagger\pi_z, \tag{28}$$

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

$$\mathbb{P} = 4\alpha\psi_x\lambda_v^x\pi_x x_1 + (1 - 4\alpha\psi_x\lambda_v^x x_1)\pi_z. \tag{29}$$

Plugging (29) into the incumbent's problem yields the following first-order condition:

$$(x_1) \quad -2\lambda_I^x(x_1 - x_I) + 4\alpha\psi_x\lambda_v^x(\pi_x - \pi_z)K = 0. \tag{30}$$

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 (30) is negative at $x_1 \geq x_I$. Since the utility is concave, it follows that $x_1^* < x_I$. □

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 30 we can verify that Assumption 3 guarantees that $x_1^* > 0$, since the incumbent's marginal utility is positive at $x_1 = 0$. □

General Model

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:

$$-2\lambda_I^x(x_1 - x_I) + K\frac{\partial \mathbb{P}}{\partial x_1} + \chi_1 - \chi_2 = 0 \tag{31}$$

$$-2\lambda_I^z(z_1 - z_I) + K\frac{\partial \mathbb{P}}{\partial z_1} + \chi_3 - \chi_4 = 0 \tag{32}$$

$$x_1 \geq 0 \quad \wedge \quad \chi_1 x_1 = 0 \tag{33}$$

$$x_1 - x' \leq 0 \quad \wedge \quad \chi_2(x' - x_1) = 0 \tag{34}$$

$$z_1 \geq 0 \quad \wedge \quad \chi_3 z_1 = 0 \tag{35}$$

$$z_1 - z' \leq 0 \quad \wedge \quad \chi_4(z' - z_1) = 0 \tag{36}$$

$$\chi_1, \chi_2, \chi_3, \chi_4 \geq 0. \tag{37}$$

Proof of Remark 3. 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. 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. Then, $\frac{\partial \mathbb{P}}{\partial x_1}$ is a function of z_1. For a leading incumbent, we have:

$$\frac{\partial \mathbb{P}}{\partial x_1} = 4\alpha\psi_x\lambda_v^x\Big[\pi_x - \Big(1 - (1 - \pi_z)4\alpha\psi_z z_1\lambda_v^z\Big)\Big] \tag{38}$$

Thus, our necessary condition becomes:

$$2\lambda_I^x x_I + K4\alpha\psi_x\lambda_v^x\Big[\pi_x - \Big(1 - (1 - \pi_z)4\alpha\psi_z z_1\lambda_v^z\Big)\Big] < 0. \tag{39}$$

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:

$$\frac{\partial \mathbb{P}}{\partial x_1} = 4\alpha\psi_x\lambda_v^x\Big(\pi_x - \pi_z 4\alpha\psi_z\lambda_v^z z_1\Big), \tag{40}$$

and our necessary condition becomes:

$$2\lambda_I^x x_I + K4\alpha\psi_x\lambda_v^x\Big(\pi_x - \pi_z 4\alpha\psi_z\lambda_v^z z_1\Big) < 0. \tag{41}$$

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. □

Proof of Proposition 3. 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 $\rho > \widehat{\rho}$. 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:

$$2\lambda_I^z z_I + K(1 - 4\alpha\psi_x\lambda_v^x x_1\pi_x)4\alpha\psi_z\lambda_v^z\pi_z < 0, \tag{42}$$

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

$$2\lambda_I^z z_I - (1 - \pi_z)(1 - 4\alpha\widehat{x}\psi_x\lambda_v^x)4\alpha\psi_z\lambda_v^z K < 0, \tag{43}$$

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

$$\widehat{x} = x_I - \frac{2\alpha\psi_x\lambda_v^x(1 - \pi_x)K}{\lambda_I^x}. \tag{44}$$

We can immediately see that the LHS of (43) 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 (43) 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 (43) 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 $\Big(\widehat{x}\big(\widehat{z}\big), \widehat{z}\big(\widehat{x}\big)\Big)$.

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

$$-2\lambda_I^x(x' - x_I) + K4\alpha\psi_x\lambda_v^x\Big[\pi_x - \Big(1 - (1-\pi_z)4\alpha\psi_z\lambda_v^z z_I\Big)\Big] > 0. \tag{45}$$

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

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

$$\widehat{z} = z_I - \frac{2\alpha\psi_Z\lambda_v^z(1-\pi_z)K(1 - 4\alpha\psi_x\lambda_v^x\hat{x})}{\lambda_I^z} > 0. \tag{46}$$

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. $\square$

Proof of Proposition 4. If $\rho < \widehat{\rho}$ then the incumbent's strategic problem is identical to the unidimensional case, therefore the equilibrium policy choice on both dimensions is the same as in the baseline.

Suppose instead, $\rho > \widehat{\rho}$. 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), $\Big(\widehat{x}(z'), z'\Big)$, $\Big(\widehat{x}(\widehat{z}), \widehat{z}(\widehat{x})\Big)$ 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 $\lambda_v^z > 0$:

$$-2\lambda_I^x(x' - x_I) + 4\alpha\psi_x\pi_x\lambda_v^x K \geq -2\lambda_I^x(x' - x_I) + 4\alpha\psi_x\lambda_v^x\Big[\pi_x - 4\alpha\psi_z\lambda_v^z z_1\pi_z\Big]K \tag{47}$$

which reduces to

$$\pi_x \geq \pi_x - 4\alpha\psi_z\lambda_v^z z_I\pi_z, \tag{48}$$

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). $\square$

Proof of Proposition 5. Suppose that the incumbent has multiple secondary dimensions $\tilde{D}$ available to open, but can only choose one. Applying the envelope theorem, we can characterize how the incumbent's equilibrium utility changes if he chooses to open dimensions with different features in the first period. For simplicity, we will assume that in the second period the officeholder implements his ideologically preferred policy on all dimensions, and denote $\tilde{K}$ the cost of losing the election in this augmented multidimensional world. Further, we denote $\tilde{d}_I$ the incumbent's ideal point on dimension $\tilde{d}$, $\rho_{\tilde{d}}$ the correlation between X and $\tilde{D}$, and $\psi_{\tilde{d}}$ the precision of the shock term on dimension $\tilde{D}$. Then, we have

$$\frac{\partial U_I^*}{\partial \tilde{d}_I} = 2(d_1 - \tilde{d}_I). \tag{49}$$

From Proposition 4 we know that $d_1 \geq \tilde{d}_I$ iff $\pi_x < \frac{1}{2}$. Therefore $\frac{\partial U_I^*}{\partial \tilde{d}_I} \geq 0$ iff $\pi_x > \frac{1}{2}$. As an aside, note that here we are not treating $\tilde{K}$ as a function of $\tilde{d}_I$, since we are comparing utility across dimensions and the cost of losing does not depend on which dimension the incumbent chooses to open in the first period. $\square$