

Learning in a Complex World: How Multidimensionality Affects Policymaking

Jesse M. Crosson ^{*} Giovanna M. Invernizzi [†] Federica Izzo [‡]

September 26, 2023

Abstract

Governments are confronted with a wide range of issues, from international affairs, to social security, to the economy. We develop a model of accountability to study policymaking in this multidimensional world. Our aim is to tackle several questions: When do officeholders address all policy dimensions that are relevant for the voters, and when do they instead focus on a subset of them? What types of reforms do policymakers pursue when they have broader versus narrower policy agendas? How do interconnections across dimensions impact policymaking? We begin with the observation that policymaking is complex, due to uncertainty about policy consequences. Voters observe outcomes and adjust their beliefs about optimal policies. When different policy issues are connected in voters' minds, learning spills over from one issue to the others. In this context, trailing incumbents tend to adopt comprehensive policy programs covering interconnected issues, while leading incumbents prioritize fewer, more independent dimensions. Additionally, we identify a substitution effect between highly connected dimensions in voters' minds, where multidimensionality can lead to either moderation or extremism in the primary policy dimension, depending on the incumbent's electoral prospects. Our results challenge the notion that a unidimensional model adequately represents policymaking in a multidimensional world where preferences across dimensions are correlated.

^{*}Purdue University, *Email:* jcrosson@purdue.edu

[†]Corresponding author. Duke University, *Email:* giovanna.invernizzi@duke.edu

[‡]UCSD, *Email:* fizzo@ucsd.edu

1 Introduction

In today’s complex and multifaceted world, governments are confronted with a wide range of issues, from international affairs, to social security, to the economy. For elected politicians, a critical challenge is to determine the prioritization of these dimensions and whether any should be left unaddressed. In fact, according to classic accounts of power and influence, the decision regarding which policy dimensions to pursue and which to avoid (Bachrach and Baratz, 1963) is perhaps the most influential factor in determining political winners and losers. Beyond the selection of the relevant dimensions to be acted upon, officeholders must set the tone of policymaking on each issue. In a multidimensional world, the choice of which issue positions should “go together” is also a strategic consideration of tremendous political consequence, as Bawn et al. (2012) and others have shown.

In this paper, we propose a game-theoretic model that examines the decision-making of policymakers within this multidimensional landscape. Our aim is to tackle several key questions: When do officeholders choose to address all policy dimensions that are relevant for the voters, and when do they instead focus on a subset of them? If some dimensions are left unaddressed, which ones are they? What types of reforms do policymakers pursue when they have broader versus narrower policy agendas? Is there a tendency for more narrow programs to implement more extreme, or moderate, reforms? How do connections across dimensions influence the decision-making of policymakers?

To delve into these questions, we adopt an accountability framework that builds on a fundamental assumption: policymaking is intricate, and both voters and elected officials face uncertainties when evaluating the potential consequences of different policies. As such, policymakers must navigate through unknowns while considering the multidimensional nature of the issues they confront (Callander, 2011; Tavits, 2007).

In this uncertain world, the consequences of the reforms implemented by officeholders play a significant role in determining their chances of reelection. A substantial body of scholarship has in fact argued that voters respond to informational challenges by assessing their personal well-being (Fiorina, 1981; Stimson, 2018). Voters thus react to the results of policy choices—not simply

the substance of the policies themselves (e.g., Fiorina 1978, Alt, Bueno de Mesquita and Rose 2011). If the incumbent's past policy choices lead to favorable outcomes, voters' evaluation of the broader policy programs espoused by the incumbent improves, thus increasing the likelihood of their reelection. Importantly, the inferences voters draw when observing outcomes depend on the exact policies implemented by the officeholder (as in Izzo, Forthcoming).

Intuitively, when a policy remains at the status quo, it doesn't provide any new information about that particular dimension. On the other hand, when new policies are introduced, they offer opportunities for voters to learn—and more extreme policies yield more informative outcomes. For instance, consider a scenario where a voter experiences a favorable outcome from an extremely leftist policy. In such a case, it becomes evident that this policy aligns well with the voter's interests. However, the outcome of a moderate policy is significantly less informative. Even if the policy moves slightly in a direction that isn't ideal for the voter, random chance may still allow them to experience relatively high welfare, making it harder to discern their true preferences or alignment with the policy.

This learning process becomes more complex in a multidimensional world, where voters may perceive different issue areas as interconnected. 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, are beneficial. In other words, voters can learn about how well a candidate's program fits their preferences in one dimension by observing the policy outcome in another. Voters, therefore, revise their preferences by taking into account the policy changes implemented by the incumbent, the outcomes that arise from those changes, and their own beliefs about how different policy areas are interconnected or correlated.

Taken together, then, the incumbent can control the amount of voter learning on each dimension both directly (via the policy on that dimension) *and* indirectly (via his choices on the other correlated dimensions that generate learning spillovers). Consequently, incumbent policymakers face a complex web of incentives as they consider the impact of their policy choices on voters' well-being and, in turn, their own chances of reelection. These incentives then 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 analysing 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's ideological preferences may be multidimensional, 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 leading ex-ante will always face incentives to prevent voter learning. For this incumbent, in fact, generating no new information guarantees his initial advantage is preserved, and the voter will choose to reelect him. Recall that more extreme policies generate more information. As a consequence, the leading incumbent will always pursue an economic policy that is more moderate than his own ideological preferences would dictate. 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 incentives to expand or contract the scope of policymaking in a multidimensional world (i.e., a world where the voter cares about multiple dimension). Moreover, the extent of correlation between distinct policy issues assumes a crucial role in shaping this choice.

In our model, trailing incumbents are motivated to promote voter learning. As a result, they have an incentive to expand the scope of their policy agenda, even including issues where they may lack ideological preferences. This expansion in fact creates additional opportunities for generating information, maximizing a trailing incumbent's electoral chances. This strategic behavior even extends to issues that the voter may not prioritize, if she perceives these issues to be correlated with the primary policy dimension she has stronger preferences over. In the absence of such correlations, these secondary issues would not significantly influence the voter's electoral decision. In this scenario, the incumbent would only take action on these issues if he held ideological preferences over them. In contrast, strong correlations generate learning spillovers. This means that policy

outcomes on secondary issues can provide valuable information about the primary dimension that the voter cares about. Consequently, secondary policy choices become strategically significant for an incumbent who seeks reelection.

Electorally leading incumbents face opposite incentives. As described above, these incumbents wish to avoid generating fresh policy information, as it could potentially undermine their existing electoral advantage. This concern becomes particularly pronounced when dealing with issues that exhibit a strong correlation with the primary dimension that the voter prioritizes, due to the possibility of learning spillovers. Consequently, leading incumbents may opt to leave these policy matters unattended, even if their own ideological inclinations would otherwise compel them to take action. By contrast, low correlation allows leading incumbents to legislate on the secondary dimension without risking their electoral advantage with voters.

Overall, the model implies that incumbent politicians who are lagging in electoral support tend to adopt more comprehensive policy programs that encompass issues that are interconnected in voters' mind, while those in the lead tend to concentrate on a narrower set of issues, prioritizing dimensions that are relatively independent or orthogonal to each other. Notice that these results are driven by the voters' *perception* of correlations across issues, and do not require that such correlations are genuine.

Next, we study how multidimensionality influences the policies that are implemented in equilibrium on each of the issues. We find that, when policy dimensions exhibit strong correlations, the presence of secondary policy issues fundamentally distorts the programs pursued by the incumbent on the primary dimension. Specifically, we reveal the emergence of a strategic substitution effect between dimensions that emerges as a consequence of learning spillovers. Instead, when issues are relatively orthogonal to each other, this distortion does not emerge and the policies pursued by the incumbent on the primary issue align with those in the unidimensional world.

Let's begin by examining the scenario of a trailing incumbent. In the context of a unidimensional environment, this incumbent consistently adopts extreme policies to encourage voter learning. Moving to a multidimensional world, naive intuition would suggest that the desire to facilitate voter

learning would drive the trailing incumbent to continue pursuing extreme policies across all available issue domains. Instead, we show that in equilibrium, the possibility to expand the policy agenda actually leads to moderation on the primary dimension. To understand this result, notice that policy outcomes on secondary issues are electorally relevant only when the voter doesn't acquire new direct information about the primary dimension, which she cares more about. By implementing moderate, and thus less informative, policies on the primary issue the incumbent can then amplify the impact of learning spillovers, ultimately maximizing his electoral chances. Trailing incumbents who, in a unidimensional context, show a propensity for risk-taking may display risk aversion when confronted with multiple dimensions.

Symmetric results hold when the incumbent is electorally advantaged. A leading incumbent always implements moderate policies in the unidimensional case. Instead, when pursuing a multidimensional policy program he chooses to adopt more extreme stances on the primary dimension. He does so in order to counteract damaging learning spillovers, i.e., in hopes of generating enough positive information on the primary economic dimension to counteract any negative electoral consequences from the secondary one. In other words, policymakers who are risk-averse in a unidimensional world may instead become risk-loving on the main policy issue when they choose to expand the scope of policymaking.

Finally, when we compare equilibrium policy choices across correlated issues, our theory suggests that leading incumbents are inclined to adopt more extreme policies on issues of significant importance to voters, while they tend to opt for more moderate policies on other dimensions. Conversely, this trend reverses for incumbents facing an electoral disadvantage.

Taken together, 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 observed correlation of preferences across various issues (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

doesn't distort the fundamental nature of policymakers' strategic challenges when addressing each issue individually. Consequently, a unidimensional model serves as a suitable approximation of a multidimensional world. However, when issues exhibit high levels of correlation in the minds of voters, policymaking in a multidimensional world takes on a significantly distinct character from the unidimensional case. In such scenarios, the multidimensional problem isn't merely the sum of multiple unidimensional problems. To truly understand policymaking, one must account for its inherent multidimensionality.

1.1 Related Literature

Our focus on voter learning contrasts with much of the existing literature on multidimensional policymaking, which instead emphasizes how the presence of additional dimensions may create issues of equilibrium existence in models of spatial competition (see Duggan 2005), or facilitate logrolling and influence coalition building within legislatures (e.g., Banks and Duggan 2000; 2006). Further, this literature generally assumes that all available dimensions must be legislated on. Thus, we offer among the first formal examinations of policymakers' decision of whether and when to expand policymaking activities to new dimensions.

In this perspective, 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* by officeholders, in a world where voters themselves are uncertain about the optimal policy.

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

by experience. In addition, our paper differs from most of these works as it considers multiple *ideological* policy dimensions.

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). Most of this literature considers a binary policy space, with one risky option and one safe option.¹ 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.

The learning technology we use relates to the models introduced in Ashworth, Bueno de Mesquita and Friedenber (2017) and Izzo (2022). In both models, as in ours, the policymaker chooses between a continuum of actions, and his choice determines how informative the resulting outcome is going to be for the voter. Ashworth, Bueno de Mesquita and Friedenber (2017), however, consider a continuous choice of effort which is unobserved by the voter, whereas we study an ideological policy choice which is observed by the voter. As a consequence, the voter in our model updates her beliefs (and thus ideological preferences) based on the implemented policy as well as the outcome of the reform. Izzo (2022) is closest to our setup, as it also focuses on policymaking along an ideological dimension and assumes voters observe the location of the policy choice. However, Izzo (2022), presumes policymaking is unidimensional. Our contribution is then to adapt the learning technology to study policymaking spanning multiple correlated dimensions, analyzing the officeholder’s choice over which dimensions to legislate on, as well as the interactions between the different available dimensions.

¹Strumpf (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.

2 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, $D \in \{X, Z\}$. If he chooses to act on dimension D , then he selects a policy $d_t \in \mathbb{R}$ to be implemented. If he chooses not to act on dimension D , then the status quo d_{sq} remains in place in that period. For simplicity, we normalize the status quo on each dimension to 0, $x_{sq} = z_{sq} = 0$. To avoid trivialities, we assume that if indifferent the officeholder chooses not to act on dimension d .

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

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

$$\Pr(x_v = \alpha) = \pi_x,$$

and

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

Thus, players believe the dimensions are positively correlated in a symmetric way. As the voter's initial beliefs on dimension X shift to the right, so does her prior on dimension Z . The ex-ante probability that $z_v = \alpha$, which we denote as π_z , is then given by $\rho \pi_x + (1 - \rho)(1 - \pi_x)$.

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 per-period total utility is

$$u_t^i = - \sum_d \lambda_i^d (d_t - d_i)^2,$$

where λ_i^d , for $d \in \{X, Z\}$, is the weight that player i places on dimension d . For simplicity, let $\lambda_i^x = \lambda_i = 1 - \lambda_i^z$.

On each dimension, the voter observes her realized utility plus a random shock, $-\lambda_i^d (d_t - d_i)^2 + \varepsilon_{d,t}$. The random shock is drawn *i.i.d.* in each period and for each dimension, with $\varepsilon_{d,t} \sim U \in \left[-\frac{1}{2\psi_d}, \frac{1}{2\psi_d}\right]$. 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 3.1 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 d_{sq} .
2. V observes I 's choice and her realized utility on each dimension.
3. V chooses whether to re-elect I or replace her with C .
4. The winner of the election takes office, then chooses whether to act on each dimension or instead keep the status quo.

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 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, to focus on what the voter learns from her lived experiences (i.e., the inferences she draws upon observing realized outcomes).

3 Equilibrium Analysis

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

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

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

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

Proof. All Proofs are collected in the Appendix.² □

When the voter only cares about the primary dimension ($\lambda_v = 1$), it follows from (1) that the right-wing incumbent is reelected as long as the voter believes her optimal policy on dimension X is more likely to be a right-wing one ($\widehat{\mu}^x = \frac{1}{2}$). Instead, when the voter cares about both dimensions ($\lambda_v < 1$) she becomes more lenient with the incumbent on dimension 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.³

To streamline the presentation of the results, we will assume that the voter cares sufficiently about the primary dimension X . Specifically, λ_v is sufficiently large that if the voter believes her

²More generally, since when indifferent a policymaker chooses to leave the status quo unaddressed, the voter re-elects the right-wing incumbent if and only if $\mu^x > \frac{1}{2} - \mathbb{I}_z \frac{(1-\lambda_v)z_I}{\lambda_v x_I} \frac{2\mu^z - 1}{2} \equiv \widehat{\mu}^x(\mu^z)$, where $\mathbb{I}_z = 0$ if $\lambda_I = \lambda_C = 1$ and $\mathbb{I}_z = 1$ if $\lambda_I < 1$ and $\lambda_C < 1$.

³Given our symmetry assumption $d_I = -d_c$, polarization on dimension D here is given simply by $2d_I$.

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

Before continuing with the analysis, let us introduce some useful definitions. By plugging in $\mu^x = \pi_x$ and $\mu^z = \pi_z = \pi_x \rho + (1 - \pi_x)(1 - \rho)$ into Equation 1, 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

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

3.1 Voter Learning

In this scenario, the voter gains insights into her preferred policies for each dimension through her real-life experiences. However, these experiences only represent a noisy signal of the genuine alignment between the voter’s interests and the implemented policy, and this complicates the voter’s inference problem (as in Izzo Forthcoming). 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 both her optimal platform on that dimension (*direct learning*) and on the policy-relevant state of the world on the others (*indirect learning*). For instance, suppose that the economy and healthcare 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 their healthcare policies as well. In what follows, we fully characterize these processes of direct and indirect learning in our framework.

3.1.1 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 $\tilde{\mu}^d$. The statements below refer to dimension X , though expressions for dimension Z are analogous. The key feature of the voter learning in this setup is that new and more extreme policies generate more information:

Lemma 2 (Direct Learning). *The voter's learning satisfies the following properties:*

- (i) *Her interim posterior $\tilde{\mu}^x$ takes one of three values, $\tilde{\mu}^x \in \{0, \pi_x, 1\}$;*
- (ii) *If the incumbent does not act on dimension X , then $\tilde{\mu}^x = \pi_x$;*
- (iii) *If the incumbent acts on dimension X , the amount of learning is a function of the implemented policy x_1 , as more extreme policies increase the probability that $\tilde{\mu}^x \neq \pi_x$. Furthermore, there exists a unique policy x such that if $|x_1| \geq |x|$, then $\tilde{\mu}^x \neq \pi_x$ 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 quite intuitive. Suppose the incumbent acts on the primary dimension, perhaps by lowering taxes. Lowering taxes may be good for the voter, because it may spur economic growth, or bad, because it reduces redistribution and welfare spending. Thus, 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. 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*.

When the policy is sufficiently moderate ($x_1 \in (-\frac{1}{4\alpha\psi_d\lambda_v}, \frac{1}{4\alpha\psi_d\lambda_v})$), 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. 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 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.⁴

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. 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 x_v .

We note that the assumption of uniformly distributed shocks simplifies the analysis by generating the stark learning environment described above. However, the crucial result that extreme (new) policies facilitate voter learning holds more generally, as it simply requires that the noise distribution

⁴Notice 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 .




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

satisfies the monotone likelihood ratio property.⁵

3.1.2 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 of the policy implemented on that issue. Our next result shows that such direct learning is not the only way the voter can gain new information in our multidimensional scenario. Due to the correlations between dimensions in our model, the voter can actually learn 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.

Formally, the result below shows how the voter’s posterior belief on X depends on the outcome of the secondary dimension Z . 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:*

$$\mu^x = \tilde{\mu}^x. \tag{2}$$

⁵For example, [Bils and Izzo \(2022\)](#) shows that this result holds under normally distributed shocks. There, every outcome is somewhat informative but never fully so. Nonetheless, extreme policies continue to facilitate voter learning and thus increase the variance in the posterior distribution.

Suppose instead that she learns that $z_v = \alpha$, we have:

$$\mu^x(\tilde{\mu}^x, \alpha, \rho) = \frac{\tilde{\mu}^x \rho}{\tilde{\mu}^x \rho + (1 - \tilde{\mu}^x)(1 - \rho)} \geq \tilde{\mu}^x. \quad (3)$$

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

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

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 $\mu^x = \tilde{\mu}^x$. 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 liberal economic policy will predispose the voter toward liberal policy on healthcare as well. Moreover, the higher the correlation across dimensions ρ , the stronger these learning spillovers.

4 The Incumbent's Problem

The findings in the previous sections shed light on how the incumbent's decisions in our scenario 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 expected second-period payoff. This, in turn, depends on voter learning. The information effect operates through two distinct 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 voters 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. This trade-off clearly appears in the incumbent maximization problem, which we can express as follows:

$$\max_{x_1, z_1} -\lambda_I(x_1 - x_I)^2 - (1 - \lambda_I)(z_1 - z_I)^2 - (1 - \mathbb{P}(x_1, z_1))\left(\lambda_I(x_I - x_C)^2 + (1 - \lambda_I)(z_I - z_C)^2\right), \quad (5)$$

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

Recall that, from Lemma 2, more extreme policies that move farther from the status quo are more likely to generate informative outcomes. Thus, depending on whether information is electorally beneficial or not (i.e., focusing on right-wing policies, 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 $d_{sq} = 0$.⁶ In what follows, we will see that whether one or the other distortion emerges in equilibrium depends on the incumbent's ex-ante prevailing electoral chances *and* whether he chooses to act only on a single dimension or expand the scope of policymaking.

4.1 Unidimensional Benchmark

To better understand the strategic incentives within our environment, it is useful to begin by analyzing a unidimensional benchmark. For this purpose, suppose that $\lambda_v = 1$, so that the voter

⁶This is clear from the first-order necessary conditions for an interior maximum, which are respectively:

$$(x_1) \quad -2\lambda_I(x_1 - x_I) + \frac{\partial \mathbb{P}(x_1, z_1)}{\partial x_1} \left(\lambda_I(x_I - x_C)^2 + (1 - \lambda_I)(z_I - z_C)^2\right) = 0, \quad (6)$$

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

only cares about the primary dimension X . In our setting, this also implies that there are no learning spillovers across dimensions. Consequently, the incumbent's likelihood of being reelected, denoted as \mathbb{P} , is only a function of his policy on the primary dimension:

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

- *a trailing incumbent is reelected **if and only if** the outcome on X is informative and favorable, i.e., the voter learns that $x_v = \alpha$;*
- *a leading incumbent is reelected **unless** the outcome on X is informative and unfavorable, 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 in fact be reelected for sure.

Applying Remark 1 to the incumbent's maximization problem, and denoting d_u the incumbent's optimal policy on dimension d under $\lambda_v = 1$,⁷ we obtain:

Proposition 1. *Suppose $\lambda_v = 1$. In equilibrium:*

- *The incumbent always implement his bliss point on the secondary dimension Z , $z_u = z_I$;*
- *On the primary dimension X*
 - *A leading incumbent implements a policy more moderate than his bliss point, $x_u \leq x_I$;*
 - *A trailing incumbent implements a policy more extreme than his bliss point, $x_u \geq x_I$*

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

⁷The subscript indicates that the incumbent's strategic problem is unidimensional.

implements his ideologically preferred policy on this dimension. In contrast, the implemented X -dimension policy 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 is risk-averse, and distorts policy towards 0 so as to minimize information. Notice that, since any pair of policies x and $-x$ induces the same amount of learning (Lemma 2), the right-wing incumbent never implements a policy to the left of 0.

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

4.2 Multidimensional World

When the voter's preferences span multiple dimension, $\lambda_v < 1$, the incumbent's strategic problem becomes multidimensional as well. To start studying this problem, we characterize the incumbent's probability of winning in this multidimensional world. For this purpose, it is useful to introduce the following definition:

Definition 2. *Let*

$$\hat{\rho} = \begin{cases} \frac{\pi_x(1-\Lambda^+)}{\Lambda^+(1-2\pi_x)+\pi_x} & \text{if } \pi > \frac{1}{2} \\ \frac{K^-(1-2\pi_x)}{\Lambda^-(1-\pi_x)+\pi_x} & \text{if } \pi < \frac{1}{2} \end{cases}$$

where $\Lambda^- = \frac{1}{2} \left(1 - \frac{1-\lambda_v}{\lambda_v} \frac{z_I}{x_I} \right)$ and $\Lambda^+ = \frac{1}{2} \left(1 + \frac{1-\lambda_v}{\lambda_v} \frac{z_I}{x_I} \right)$.

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

Our results highlight that when perceived correlations across dimensions are high, even issues the voter cares relatively little about may have a crucial impact on her electoral choice.

Lemma 4. *Let $\lambda_v < 1$. If the correlation between policy dimensions is low, then the incumbent's probability of being reelected is the same as in the unidimensional benchmark.*

If instead the correlation is high, then the incumbent's probability of being reelected 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***
 - *if the outcome on X is uninformative but 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 but 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. 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.

When $\rho > \frac{1}{2}$, i.e., the two dimensions are correlated in the voter's mind, indirect learning contributes to the evolution of the voter's preferences. In this context, observing a negative (positive) outcome on healthcare leads the voter to adjust her beliefs on the economy against (in favor of) the incumbent as well. However, 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.⁸

⁸Notice that, when π_x is very close to $\frac{1}{2}$, the critical value $\hat{\rho}$ may actually be lower than $\frac{1}{2}$.

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.

4.2.1 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, $\lambda_I = \lambda_C = 1$. Here, even though she cares about both dimensions, the voter’s retention decision does not *directly* depend on her beliefs over Z , since she anticipates that neither I nor C will act on Z in the second period.⁹ More specifically, the voter’s optimal retention rule is exactly the same as in the unidimensional case: she retains the right-wing incumbent if and only if $\mu^x > \frac{1}{2}$. However, by Lemma 3, the voter’s posterior on X is a function of her realized utility on Z . Therefore, even though the incumbent’s ideological preferences are unidimensional, his strategic problem is *multidimensional*.

These assumptions allow us to isolate the strategic incentives emerging solely due to the learning spillovers across dimensions. In section 4.2.4, we complete the analysis allowing for $\lambda_I, \lambda_C < 1$.

4.2.2 The Dimensionality of Policymaking

First, we characterize the conditions under which the incumbent chooses to open the secondary policy dimension:

Proposition 2. *Suppose $\lambda_v < 1$ and $\lambda_I = \lambda_C = 1$. Then, the incumbent chooses to act on Z (i.e., $z_1 \neq z_{sq}$) if and only if he is trailing and the correlation with the primary dimension is high.*

⁹Recall that we assume that when indifferent the incumbent chooses not to act on Z .

It follows straightforwardly from Lemma 4 that 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 too small (i.e., the secondary dimension is insufficiently correlated with the primary). In this case, a trailing incumbent is indifferent between acting on the secondary dimension and keeping the status quo and (by assumption) chooses not to act. If instead ρ is sufficiently large, the trailing incumbent can exploit learning spillovers to increase his probability of resurrecting himself. In equilibrium he will therefore always choose to expand the scope of policymaking to the secondary dimension, even if he has no ideological taste for it.

4.2.3 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 2:

Proposition 3. *Suppose that the correlation between the two dimensions is high, so that in equilibrium the incumbent 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, extreme policymaking need not follow from extreme ideological preferences.

We next characterize 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 4. *Suppose that the correlation between the two dimensions is high, so that in equilibrium the incumbent chooses to act on the secondary dimension Z . Then, his policy choice on the primary one x_1^* satisfies $x_1^* < x_I \leq x_u$.*

When a trailing incumbent cannot exploit the secondary dimension (i.e., $\lambda_v = 1$ 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, $x_u > x_I$. Naive intuition would suggest that in a multidimensional world, this incumbent should have incentives to continue pursuing extreme policies on all issues. Instead, Proposition 4 highlights that a high correlation generates a strategic substitution effect between policy dimensions. This correlation allows the trailing incumbent to exploit learning spillovers (i.e., $\lambda_v < 1$ and $\rho > \widehat{\rho}$), inducing *moderation* on the primary dimension: $x_1^* < x_I$.

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, immigration affects the voter’s selection only if she fails to learn directly about 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 voter to learn on the secondary.

Notice that the above discussion implies that if the incumbent chooses to gamble on Z , then he becomes risk-averse on X to magnify the impact of the learning spillovers. However, in principle, the incumbent may 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. Proposition 4 shows that this *never* occurs in equilibrium.¹⁰

¹⁰Indeed, if ρ is sufficiently large such that outcomes on the secondary dimension can be electorally relevant, the trailing incumbent always chooses to gamble on Z and pursue moderate policies on X .

The reason lies in the fact 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; However, when $\pi_x < \frac{1}{2}$ we have that $\pi_z > \pi_x$, where π_z is the prior probability that the voter’s optimum on Z is a right-wing one. In other words, under when the incumbent is trailing the probability that his preference are aligned with the voter’s 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 if the optimal policy on the primary dimension is actually a left-wing one.

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

4.2.4 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., $\lambda_i < 1$ for $i \in \{I, V, C\}$). In contrast with the analysis presented above, the incumbent now has both strategic *and ideological* preferences over the secondary dimension Z . As above, our goal is to characterize the conditions under which the incumbent chooses to open the secondary dimension, and to examine how this influences his policy on the primary one.

Proposition 5. *A trailing incumbent always acts on the secondary dimension Z . Suppose instead the incumbent is leading. Then he acts on Z if and only if*

- *He cares enough about this dimension, i.e., λ_I is sufficiently small, or*
- *The correlation with the primary dimension is low, $\rho < \hat{\rho}$*

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 one hand, he has ideological preferences over the secondary dimension and statically prefers to act on it. 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 sufficiently 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 and make *ideological/policy* gains, while nevertheless avoiding the negative *electoral* consequences. However, if the correlation ρ is high, then dynamic considerations emerge and the incumbent only acts on Z if his ideological preferences on this dimension are sufficiently strong (i.e., λ_I is low).

In short, a trailing incumbent always has incentives to expand the scope of policymaking, even incorporating 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 policymaking, even inducing him to abandon issues he cares about.

The next result highlights that the substitution effect described in the previous section 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 6. *Suppose that the incumbent chooses to act on the secondary dimension. 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 $z_1^* \geq z_u = z_I$ and $x_1^* \leq x_u$;
- when the incumbent is leading, we have $z_1^* \leq z_u = z_I$ and $x_1^* \geq x_u$.

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

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.

Suppose 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 gamble 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 1, 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, we have:

Proposition 7. *All else equal, a leading incumbent prefers to open dimensions for which he is more moderate. Suppose instead the incumbent is trailing. Then, all else equal, he prefers 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.

5 Implications and 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 always expand the scope of policymaking and, at the same time, *moderate* on the primary dimension. This substitution effect emerges because the different dimensions are connected in the voter's mind, 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, multidimensionality generates the opposite result. Here, while the incumbent tends to pursue moderate policy in a unidimensional world, multidimensionality generates incentives for extremism on the primary policy dimension in order to mitigate the negative consequences of learning spillovers.

These results have important implications for how we think about the mapping from a unidimensional world to a multidimensional one. Both theoretically and empirically, efforts to understand multidimensional politics have repeatedly been stymied by several challenges, ranging from equilib-

rium existence in theoretical models, to the inability to fit too many parameters in empirical ones. A common response in the literature is sidestep the issue, and assume that a unidimensional world is a close enough analogue of the multidimensional one. In particular, this assumption is often considerable reasonable because of the observation that preferences across issues are surprisingly correlated. In the words of Converse (1964) (p. 207) “if a person is opposed to the expansion of social security, he is probably a conservative and is probably opposed as well to any nationalization of private industries, federal aid to education, sharply progressive income taxation, and so forth”. When the different issue areas are connected in such a way in political actors’ mind, this 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 in the voters’ mind 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 the issues is high, voter learning will spill over between the 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

- Alt, James, Ethan Bueno de Mesquita and Shanna Rose. 2011. “Disentangling accountability and competence in elections: evidence from US term limits.” *The Journal of Politics* 73(1):171–186.
- 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 Friedenber. 2017. “Accountability and information in elections.” *American Economic Journal: Microeconomics* 9(2):95–138.
- Bachrach, Peter and Morton S Baratz. 1963. “Decisions and nondecisions: An analytical framework.” *American political science review* 57(3):632–642.
- Banks, Jeffrey s. and John Duggan. 2000. “A Bargaining Model of Collective Choice.” *American Political Science Review* 94(1):73–88.
- Banks, Jeffrey S and John Duggan. 2006. “A general bargaining model of legislative policy-making.” *Quarterly Journal of Political Science* 1(1):49–85.
- 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.
- 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.
- Duggan, John. 2005. “A survey of equilibrium analysis in spatial models of elections.” *Unpublished manuscript* .
- Fiorina, Morris P. 1978. “Economic retrospective voting in American national elections: A micro-analysis.” *American Journal of political science* pp. 426–443.
- 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* .
- 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* .
- McMurray, Joseph. 2014. Why the political world is flat: an endogenous left-right spectrum in multidimensional political conflict. In *APSA 2014 Annual Meeting Paper*.
- 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.
- 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.

Appendix

Main Results - Proofs

To reduce the number of cases under consideration, we will assume that $x_I < x'$ and $z_I < z'$.

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:

$$\begin{aligned} & -\lambda_v[\mu^x(x_I - \alpha)^2 + (1 - \mu^x)(x_I + \alpha)^2] - (1 - \lambda_v)[\mu^z(z_I - \alpha)^2 + (1 - \mu^z)(z_I + \alpha)^2] > \quad (8) \\ & -\lambda_v[\mu^x(x_C - \alpha)^2 + (1 - \mu^x)(x_C + \alpha)^2] - (1 - \lambda_v)[\mu^z(z_C - \alpha)^2 + (1 - \mu^z)(z_C + \alpha)^2]. \end{aligned}$$

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

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

which rearranged yields:

$$\mu_v^x > \frac{1}{2} + \frac{(1 - \lambda_v)z_I (1 - 2\mu_v^z)}{\lambda_v x_I} \equiv \hat{\mu}_v^x(\mu^z). \quad (9)$$

□

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:

$$\left[-\lambda_v(x_t - x_v)^2 - \frac{1}{2\psi_x}, -\lambda_v(x_t - x_v)^2 + \frac{1}{2\psi_x} \right]. \quad (10)$$

We can immediately see that if V observes $u_v^t > -\lambda_v(x_t + \alpha)^2 + \frac{1}{2\psi_x}$, she knows for sure that she

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

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

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

Since $\varepsilon_{x,t}$ is uniformly distributed, we have $f(u_v^t + \lambda_v(x_t + \alpha)^2) = f(u_v^t + \lambda_v(x_t - \alpha)^2)$, 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 (10) is empty. This requires:

$$-\lambda_v(x_t + \alpha)^2 + \frac{1}{2\psi_x} + \lambda_v(x_t - \alpha)^2 + \frac{1}{2\psi_x} \leq 0 \quad (11)$$

which rearranged yields:

$$x_t \geq \frac{1}{4\alpha\lambda_v\psi_x}. \quad (12)$$

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

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

Since the two probabilities are symmetric, we have

$$\begin{aligned} \Pr(L_x = 1 | \pi_x, 0 < x_t < x') &= \Pr\left(-\lambda_v(x_t - \alpha)^2 + \varepsilon_{x,t} > -\lambda_v(x_t + \alpha)^2 + \frac{1}{2\psi_x}\right) \\ &= \Pr\left(\varepsilon_{x,t} < 4\lambda_v\alpha x_t - \frac{1}{2\psi_x}\right) \\ &= 4\alpha x_t \lambda_v \psi_x, \end{aligned} \quad (13)$$

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

The proof for dimension Z is analogous therefore omitted. \square

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

$$-\lambda_I(x_1 - x_I)^2 - 4\alpha\psi_x x_1(1 - \pi_x) \left(\lambda_I(x_I - x_C)^2 + (1 - \lambda_I)(z_I - z_C)^2 \right), \quad (14)$$

which yields the following first-order necessary condition (which is also sufficient since the problem is concave):

$$-2\lambda_I(x_1 - x_I) - 4\alpha\psi_x(1 - \pi_x) \left(\lambda_I(x_I - x_C)^2 + (1 - \lambda_I)(z_I - z_C)^2 \right) = 0. \quad (15)$$

Rearranging (15) yields:

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

It follows that

$$x_1 = \max \left\{ 0, x_I - \frac{2\alpha\psi_x(1 - \pi_x)}{\lambda_I} \left(\lambda_I(x_I - x_C)^2 + (1 - \lambda_I)(z_I - z_C)^2 \right) \right\}. \quad (16)$$

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

$$-\lambda_I(x_1 - x_I)^2 - (1 - 4\alpha\psi_x x_1 \pi_x) \left(\lambda_I(x_I - x_C)^2 + (1 - \lambda_I)(z_I - z_C)^2 \right), \quad (17)$$

which yields the following first-order necessary condition (which is also sufficient):

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

which rearranged yields:

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

It follows that

$$x_1 = \min \left\{ x_I + \frac{2\alpha\psi_x\pi_x}{\lambda_I} (\lambda_I(x_I - x_C)^2 + (1 - \lambda_I)(z_I - z_C)^2), \frac{1}{4\alpha\lambda_v\psi_x} \right\}. \quad (18)$$

□

Proof of Lemma 4. Suppose that $\pi_x > \frac{1}{2}$. Then $\hat{\rho}$ solves:

$$\mu^x(\emptyset, -\alpha, \rho) = \hat{\mu}_v^x(0), \quad (19)$$

where

$$\mu^x(\emptyset, -\alpha, \rho) = \frac{(1 - \rho)\pi_x}{(1 - \rho)\pi_x + \rho(1 - \pi_x)}, \quad (20)$$

which yields:

$$\hat{\rho} = \frac{1 - \hat{\mu}_v^x(0)\pi_x}{1 - \pi_x\hat{\mu}_v^x(0)}. \quad (21)$$

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

$$\mu^x(\emptyset, \alpha, \rho) = \hat{\mu}_v^x(1), \quad (22)$$

where

$$\mu^x(\emptyset, \alpha, \rho) = \frac{\pi_x\rho}{\pi_x\rho + (1 - \pi_x)(1 - \rho)}. \quad (23)$$

Combining the above, we have

$$\hat{\rho} = \frac{(1 - \pi_x)\hat{\mu}_v^x(1)}{\pi_x(1 - 2\hat{\mu}_v^x(1)) + \hat{\mu}_v^x(1)}. \quad (24)$$

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

□

Proof of Proposition 2. Follow from Lemma 4 and the assumption that when indifferent the incumbent prefers not to act on Z . □

Proof of Proposition 3. Recall that under $\lambda_I = 1$ the incumbent's utility depends on z_1 only via the voter learning. Further, if the incumbent chooses to act on Z in equilibrium it must be the case that his probability of winning is increasing in the probability of generating an informative outcome on Z . This yields that in equilibrium the incumbent will always choose to implement a fully informative policy $z_1^* > z'$. □

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

$$4\alpha\psi_x \pi_x x_1 + (1 - 4\alpha\psi_x x_1)\pi_z. \quad (25)$$

Note that, given $\pi_x < \frac{1}{2}$, $\pi_z = \pi_x \rho + (1 - \pi_x)(1 - \rho) > \pi_x$, therefore the incumbent's probability of winning is decreasing in x_1 . It follows from Equation 6 and 7 that $x_1^* < x_I$. □

Proof of Proposition 5. If $\rho < \hat{\rho}$ then the incumbent's retention chances are not a function of his choice on the Z dimension, and $z^* = z_I$.

Suppose instead, $\rho > \hat{\rho}$. A trailing incumbent's probability of being reelected is only a function of his policy choice both dimensions. However, we can see that the utility is always increasing in z_1 at $z_1 = 0$, for all values of x_1 . Thus, in equilibrium the trailing incumbent always acts on the secondary policy dimension. Indeed, under high correlation this holds true even if $\lambda_I = 1$, i.e., the incumbent has no ideological preferences to act on the secondary dimension.

Similarly, a leading incumbent chooses to open the secondary dimension if and only if his utility is increasing in z_1 at $z_1 = 0$. Under the assumption on ρ , I 's retention probability is given by:

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

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

$$2(1 - \lambda_I)z_I - (1 - \pi_z)(1 - 4\alpha \hat{x} \lambda_v \psi_x)4\alpha \psi_z(1 - \lambda_v)K > 0, \quad (26)$$

where \hat{x} solves

$$-2\lambda_I(x_1 - x_I) - 4\alpha \psi_x \lambda_v \left[1 - \pi_x - 4\alpha \psi_z(1 - \lambda_v)z_1(1 - \pi_z) \right] K = 0 \quad (27)$$

and is equal to:

$$\hat{x} = x_I - \frac{4\alpha \psi_x \lambda_v (1 - \pi_x) [\lambda_I 4x_I^2 + (1 - \lambda_I)4z_I^2]}{\lambda_I}. \quad (28)$$

Condition 26 is satisfied for $\lambda_I < \hat{\lambda}_I$. The expression for $\hat{\lambda}_I$ is lengthy therefore omitted. Intuitively, the incumbent opens the secondary dimension when he sufficiently cares about it.

□

Proof of Proposition 6. From the above, if $\rho < \hat{\rho}$ then the equilibrium policy choice on X is the same as in the unidimensional baseline.

Suppose instead, $\rho > \hat{\rho}$. Denote $\hat{d}(-d)$ the policy that satisfies the first order condition on dimension d , given the policy choice on dimension $(-d)$. Further, recall that d_u is the optimal policy on dimension d in the unidimensional benchmark.

Consider first a trailing incumbent. From inspection of the utility we can verify that there are only three possible equilibrium candidates: (x', z_I) , $(\hat{x}(z'), z')$ or $(\hat{x}(\hat{z}), \hat{z}(\hat{x}))$. This follows from two observations: first, the utility is always increasing in x_1 at $x_1 = 0$, for all values of z_1 , therefore it must be the case that $x_1^* > 0$; second, if $x_1 = x'$ the probability of winning is not a function of z_1 .

An inspection of the first order conditions gives us that $\hat{z}(\hat{x}) > z_I$ and $\hat{x}(z > 0) < x_u$. Thus, sufficient condition to ensure that $x_1^* \leq x_u$ is that $x_1^* = x'$ implies $x_u = x'$. First, suppose that the utility is concave. Then, the result follows from the fact that that if the incumbent's utility is increasing in x_1 at $x_1 = x'$ under $\lambda_v < 1$, then it must also be increasing under $\lambda_v = 1$:

$$-2\lambda_I(x' - x_I) + 4\alpha\psi_x\pi_x K \geq -2\lambda_I(x' - x_I) + 4\alpha\psi_x\lambda_v \left[\pi_x - 4\alpha\psi_z(1 - \lambda_v)z_1\pi_z \right] K \quad (29)$$

which reduces to

$$\pi_x \geq \lambda_v \left[\pi_x - 4\alpha\psi_z(1 - \lambda_v)z_1\pi_z \right], \quad (30)$$

which is always satisfied.

Next, suppose the utility is convex. Denote $u^m(x, z)$ the incumbent's utility in the multidimensional world and $u^u(x, z)$ his utility in the unidimensional baseline. Then, the result follows from the fact that, $u^m(x', z_I) = u^u(x', z_I)$ but $u^m(0, \hat{z}(0)) > u^m(0, z^I) > u^u(0, z^I)$.

Finally, consider a leading incumbent. from inspection of the utility we can verify that there are only four possible equilibrium candidates: $(0, \hat{z}(0))$, (z', z_I) , $(\hat{x}(0), 0)$ or $(\hat{x}(\hat{z}), \hat{z}(\hat{x}))$.

An inspection of the first order conditions gives us that $\hat{z}(\hat{x}) < z_I$ and $\hat{x}(z > 0) > x_u$. Thus, sufficient condition to ensure that $x_1^* \geq x_u$ is that $x_1^* = 0$ implies $x_u = 0$. First, suppose that the utility is concave. The result follows from the fact that if the incumbent's utility is decreasing in

x_1 at $x_1 = 0$ under $\lambda_v < 1$, then it must also be decreasing under $\lambda_v = 1$:

$$2\lambda_I x_I - 4\alpha\psi_x(1 - \pi_x)K \leq 2\lambda_I x_I - 4\alpha\psi_x\lambda_v \left[1 - \pi_x - 4\alpha\psi_z(1 - \lambda_v)z_1(1 - \pi_z)\right]K, \quad (31)$$

which reduces to

$$1 - \pi_x \geq \lambda_v \left[1 - \pi_x - 4\alpha\psi_z(1 - \lambda_v)z_1(1 - \pi_z)\right], \quad (32)$$

which is always satisfied.

Next, suppose the utility is convex. Denote $u^m(x, z)$ the incumbent's utility in the multidimensional world and $u^u(x, z)$ his utility in the unidimensional baseline. Then, the result follows from the fact that, $u^m(x', z_I) = u^u(x', z_I)$ but $u^m(0, \hat{z}(0)) < u^u(0, \hat{z}(0)) < u^u(0, z^I)$.

□

Proof of Proposition 7. 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

$$\frac{\partial U_I^*}{\partial \tilde{d}_I} = 2(d_1 - \tilde{d}_I). \quad (33)$$

From Proposition 6 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.

□