Learning in a Complex World:
How Multidimensionality Affects Policymaking

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

March 31, 2025

Abstract

Policymaking inherently involves uncertainty about policy consequences. This uncertainty leads voters to rely on their past experiences when forming beliefs about policies, which ultimately shapes their electoral decisions. Officeholders, in turn, strategically weigh how their policy choices between bold initiatives and incremental options will affect voter learning and thus their electoral prospects. While conventional wisdom suggests leading incumbents always prefer safe policies and trailing ones pursue risky "gambles for resurrection," this paper develops a model of policy experimentation showing that this logic often breaks down in multidimensional policy contexts. Our central result identifies a strategic substitution effect, whereby trailing incumbents may moderate on certain issues, while leading ones might pursue riskier policies than they otherwise would. This finding yields novel empirical implications regarding issue salience, policy extremism, and the incumbent's electoral prospects, which we illustrate through a preliminary test using U.S. data.

*Vanderbilt University, *Email*: `federica.izzo@vanderbilt.edu`

†Bocconi University, *Email*: `giovanna.invernizzi@unibocconi.it`

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

1 Introduction

Policymaking is an exercise in navigating uncertainty. Elected officials and voters alike grapple with imperfect information about how different policy alternatives will ultimately shape outcomes (Callander, 2011). Voters, in turn, draw on their past experiences to form beliefs about policy consequences, and these beliefs inform their electoral choices (Fiorina, 1981; Stimson, 2018). As a result, officeholders must carefully consider how their policy decisions affect voter learning about the most effective policies—and thus influence their own prospects for reelection.

Conventional wisdom suggests that politicians who are electorally secure always have incentives to be more cautious than their ideological preferences would dictate—adopting moderate reforms to limit voter learning—while those at risk of losing office always pursue bolder policies to "gamble on resurrection" (Rose-Ackerman, 1980; Downs and Rocke, 1994). In this paper, we show that this conventional wisdom, rooted in a unidimensional logic, no longer holds once we acknowledge the multidimesional nature of policy making. When incumbents must navigate multiple policy dimensions, those who are trailing can be induced to moderate on certain issues, while those who are leading may be emboldened to pursue riskier policies than they would otherwise.

This paper develops a game-theoretic model of electoral accountability in which an incumbent makes policy choices across two correlated dimensions. We designate one of the two dimensions as the "primary" one, in that voters weigh it more heavily in their electoral decisions and they initially hold more information on the policy consequences on this dimension than on the secondary one. The incumbent faces two key decisions: whether to engage in policymaking on each dimension and, if so, what policy to implement.

When a policy remains at the status quo, voters gain no new information about the policy-relevant state of the world on that dimension. In contrast, new policies create opportunities for voter learning, with more radical reforms yielding more informative outcomes. This reflects that substantial changes reveal fundamental strengths and weaknesses of the policy approach that would remain hidden during incremental adjustments. For example, a major healthcare reform reveals more about underlying economic and social structures than minor adjustments to existing programs.

Importantly, in a multidimensional world learning about one dimension may also generate policy-relevant information about another. As such, when there are multiple interconnected dimensions, informational spillovers complicate the learning process. For example, if a voter's experience suggests that liberal economic reforms are optimal for her, then she may also become more inclined to believe that liberal reforms on other issues, such as healthcare, are beneficial. In other words, when issues are correlated, outcomes on one dimension influence voters' beliefs on the others. Thus, an incumbent can control the amount of voter learning both directly (via the policy on that dimension) and indirectly (via his choices on other correlated dimensions). Consequently, the incumbent faces a complex web of incentives as he considers the impact of policies on voter learning and his own ideological preferences.

Our central result identifies a strategic substitution effect that emerges under correlated dimensions. In particular, we show that, when incumbents engage in policymaking across multiple dimensions, they adopt radical reforms in one dimension while pursuing more conservative policies in others. The way that this substitution effect operates is determined by the incumbent's initial electoral prospects and the electoral salience of each issue: While a trailing incumbent pushes bolder policy experiments on the secondary policy dimension and adopts more moderate reforms on the primary one, a leading incumbent takes the opposite approach. This result generates new theoretical insights into the effects of electoral accountability on policy experimentation and produces novel empirical predictions, which we further discuss below. While this substitution effects counters the conventional wisdom on "gambling on resurrection", as we discuss below, it is still rooted in the incumbent's incentives to control information.

When choosing which policy to implement on each dimension, the incumbent in our model considers his own ideological preferences as well as the electoral benefits of either facilitating or hampering voter learning. A trailing incumbent needs to generate information in order to get reelected, while a leading one is is guaranteed reelection if the voter learns nothing new. Because more radical reforms facilitate voter learning, in a unidimensional world a trailing incumbent always pursues a reform that is farther from the status quo compared to what he would choose absent reelection incentives. In contrast, a leading incumbent is induced to moderate, moving closer to the status quo. Thus, in our model, policymaking in a unidimensional world aligns with the conventional intuition.

A naïve extension of this logic might suggest that, when pursuing a multidimensional policy agenda, a leading incumbent would moderate across all dimensions, while a trailing incumbent would become extreme everywhere. But this reasoning can fall apart once we account for learning spillovers—the fact that policy outcomes on one issue shape voter beliefs about others. Even a low-salience dimension can swing an election if the resulting outcome shifts voter perceptions on more politically central issues. However, this spillover effect weakens when voters can directly learn from high-salience issues, reducing the indirect influence of secondary dimensions.

With this in mind, consider a leading incumbent who engages in policymaking across multiple correlated dimensions. Negative outcomes on secondary issues could induce learning spillovers onto the primary dimension that ultimately cost him the election. The incumbent's leading status implies that on high-salience issues, where initial uncertainty is lower, positive outcomes are more likely. Therefore, in order to mitigate the risk of damaging learning spillovers, the incumbent implements more moderate reforms on secondary issues while pursuing more radical, and therefore more informative, policies on electorally salient dimensions. A trailing incumbent faces opposite incentives: to capitalize on positive learning spillovers, he pursues more moderate reforms on more salient dimensions and more radical ones on secondary issues. By doing so, he maximizes the chance that voter learning works in his favor, even if direct information about the primary dimension remains limited. Under certain conditions, this substitution effect becomes so pronounced that a leading incumbent adopts a more radical primary-dimension reform than he would without electoral incentives, while a trailing incumbent pursues a more moderate one—in sharp contrast with the conventional intuition.

Our theoretical results generate clear empirical implications: the correlation between electoral salience and the radicalism of policy reforms depends on the incumbent's electoral position. A leading incumbent should pursue more extreme policies (relative to the status quo) on high-salience issues, while a trailing incumbent should do the opposite, shifting extremism to secondary issues. Using data from U.S. policymaking, we find preliminary evidence consistent with our theoretical expectations. While further empirical work is needed, these results suggest a new way of understanding the strategic calculus of policymaking in a multidimensional world.

Ultimately, this paper has wide implications for our approach to studying policymaking. Both theoretically and empirically, it is common for scholars to assume that a one-dimensional world closely approximates the multidimensional reality. Scholars often justify this simplification by pointing to correlations in voter preferences across issues (Converse, 1964; McMurray, 2014). Yet, our framework highlights that these very correlations introduce strategic interactions that fundamentally alter policymaking. When issues are truly independent, treating policymaking as a series of separate, unidimensional decisions works in our framework. But when issues are correlated—whether structurally or in voters' perceptions—this logic breaks down. The substitution effect reshapes incentives, meaning that policy choices across dimensions interact in ways a unidimensional model cannot capture. To truly understand policymaking, one must account for its inherent multidimensionality.

2 Related Literature

Our paper contributes to the literature on multidimensional policymaking. Most of the works in this literature 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 the optimal direction of policy programs and learns by experience.

In this sense, our paper connects to 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). Particularly noteworthy in this literature is Blumenthal (2025), who also examines policy experimentation, but assumes voters face uncertainty about politicians' types, which can be ideological, benevolent, or captured by special interests. We consider politicians with well-defined ideologies, and therefore assume their ideal policies are common knowledge.[1] Furthermore, this paper, like most of the extant literature, considers a single policy dimension and only allows the incumbent to choose between one of two policies: a safe one and a risky one.

In contrast, we innovate by considering a continuum of policies on each of multiple dimensions, which allows us to analyze the intensity of experimentation and reveal nuances in policymaking strategies that binary models cannot capture. While Callander (2011) and Callander and Hummel (2014) also study experimentation with continuous policies, they abstract from the dynamic electoral considerations central to our analysis. Moreover, our learning framework differs fundamentally: where Callander's voters experiment to learn precise policy consequences within a known ideological direction, our voters seek to determine which ideological direction is optimal in the first place.

The learning technology we use relates more closely to the models introduced in Ashworth, Bueno de Mesquita and Friedenberg (2017) and Izzo (2023). However, these works focus on a single-dimensional choice for the incumbent. Furthermore, Ashworth, Bueno de Mesquita and Friedenberg (2017) considers 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 experiment.

3 The Model

Players and actions. Our model has three players: two policy-motivated candidates—the incumbent, I, and a challenger, C—and a representative voter, V. In each of the two periods in the model,

[1]We note that, because of our assumption that the voter cares exclusively about outcomes (and thus, her utility is monotonic in policy), assuming uncertainty about the politicians' ideal points would have no effects on our results, as long as the voter knows which politician is right-wing and which one is left-wing. We briefly return to this point following the presentation of the model setup.

the officeholder chooses whether to act on one or both of two policy dimensions, $D \in \{X, Z\}$.[2] 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. At the end of the first period, the outcome of the first-period policy choice on each dimension publicly realizes, and the voter makes her electoral choice.

Policy Outcomes and Voter Payoff In this model, the voter has no intrinsic preferences over policies. Instead, she cares exclusively about outcomes. Specifically, her utility in period t is given by

$$u_{v,t} = \lambda_v^x o_t^x + \lambda_v^z o_t^z, \tag{1}$$

where o_t^d is the realized outcome on dimension d in period t, and $\lambda_v^d \geq 0$ is the weight dimension d has on the voter's welfare. The voter observes

In this setting, the voter's preferences over policies, and candidates, will be fully determined her beliefs about the relationship of policies to outcomes. We model this relationship in a simple way: on each dimension, outcomes are a linear function of the chosen policy, d_t, an unknown state of the world, ω_d, and an idiosyncratic shock, $\varepsilon_{d,t}$:

$$o_t^d = \omega_d d_t + \varepsilon_{d,t}. \tag{2}$$

The random shock $\varepsilon_{d,t}$ is drawn in each period and for each dimension from a uniform distribution with support $\left[-\frac{1}{2\psi_d}, \frac{1}{2\psi_d}\right]$.[3] The state ω_d can take one of two values: $\omega_d \in \{-\alpha, \alpha\}$. Straightfor-

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

[3]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 6.3 below.

wardly, the sign of ω_d indicates the optimal direction of policymaking for the voter: she benefits from right-wing policies if $\omega_d = \alpha$, and from left-wing policies otherwise. Notice that, since our initial status quo is normalized to 0 on each dimension, we can also interpret ω_d as capturing the optimal direction of reform for the voter. If $\omega_d = \alpha$, reforms that move the status quo to the right are optimal; If $\omega_d = -\alpha$, instead, liberal reforms are preferred by the voter.[4]

The state ω_d is initially unknown to all, and players share common prior beliefs:

$$\Pr(\omega_x = \alpha) = \pi_x \text{ and } \Pr(\omega_x = -\alpha) = 1 - \pi_x, \tag{3}$$

and

$$\Pr(\omega_z = \alpha | \omega_x = \alpha) = \Pr(\omega_z = -\alpha | \omega_x = -\alpha) = \rho \geq \frac{1}{2}. \tag{4}$$

Thus, the players believe the dimensions to be (weakly) positively correlated. As the belief about the state on dimension X shifts to the right, so does the prior about the state on dimension Z. It follows that the ex-ante probability that $\omega_z = \alpha$ is $\Pr(\omega_z = \alpha) = \pi_z = \rho\pi_x + (1-\rho)(1-\pi_x)$.

Notice that in our setting players initially have more information about the voter's optimal direction of reform 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.

Politicians. Because politicians represent constituencies with more well-defined policy interests, their optimal policies are state-independent. Alternatively, we may think about parties who have well-defined ideologies, and thus have preferences over policies rather than outcomes. For $i \in \{C, I\}$, payoff in period t is:

$$u_{i,t} = \lambda_i^x f_i^x(x_t) + \lambda_i^z f_i^z(z_t), \tag{5}$$

where f is twice differentiable, symmetric and concave, maximized at $d_t = d_i$. Thus, d_i is i's ideal point on dimension d. Without loss of generality, let $x_I > 0 > x_C$ and $z_I > 0 > z_C$. We note

[4]Because of our functional form assumption, the voter would optimally set the policy to either + or $-\infty$ under full information. It is important to note that this has no bearing on our results.

that we could allow the parties' preferences to have a state-dependent component (along with the ideological bias). This would complicate the analysis without changing the qualitative results.

It is also important to emphasize that allowing the voter to face uncertainty over the politicians' ideal points would have no effect on our results, as long as the voter knows that the incumbent's bliss point is to the right of the challenger's on each dimension (or vice versa). As our analysis below will make clear, because the voter's utility is monotonic in the implemented policy, her optimal decision in our model depends solely on whether she is sufficiently confident that the optimal policy is right-wing, and not on the exact location of the candidates' ideal points (see Remark 1).[5] As a consequence, any inference the voter may draw about the incumbent's ideal point from his policy choice would be electorally irrelevant, and the richer setting would be essentially isomorphic to the model presented here.

Timing. The timing is as follows.

1. Nature draws $\omega_d \in \{-\alpha, \alpha\}$ for each dimension $D \in \{X, Z\}$.
2. 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.
3. V observes I's choice and the realized outcomes on each dimension.
4. V chooses whether to re-elect I or replace him with C.
5. 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.[6]

[5]The model analyzed in the main body of the paper (under Assumption 1 below) is isomorphic to the asymmetric information version. In the generalized model analyzed in the Appendix (where Assumption 1 is relaxed), the voter's beliefs about the incumbent's ideal point directly enter her retention rule. However, our key result (i.e., the substitution effect across dimensions), continues to emerge for λ_v^x and ρ sufficiently high.

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

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. This allows us to assume away the possibility that the incumbent's policy choice directly provides information to the voter and instead focus on what the voter learns from her lived experiences (i.e., the inferences she draws upon observing realized outcomes).

In order to isolate the electoral incentives to pursue a multidimensional policy agenda, we begin by assuming that politicians do not intrinsically care about the secondary dimension Z:

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

We relax this assumption in Appendix B, where we consider politicians who have ideological preferences on both dimensions. We then build on this general model to discuss empirical implications in Section 7.

3.1 Model Discussion

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

Voter Uncertainty and Retrospective Voting. Voters, in our model, grapple with a fundamental problem: they do not know, with certainty, how different policies would affect their welfare (Callander, 2011; Tavits, 2007). For instance, higher taxation may be good for a voter, as it improves the provision of public goods, or bad for her, if it hampers economic growth. This uncertainty is neither novel nor unique to our model. Tavits (2007), for example, distinguishes between 'principled' and 'pragmatic' issue domains. The latter is defined by uncertainty over 'what types of policies are related to what sorts of outcomes' (ibid: 155), which is precisely the kind of uncertainty we capture.

Our model assumes that voters respond to these informational challenges by looking at their experiences with different policy choices (Stimson, 2018). That is, voters look back at the incumbent's actions and how these impact observed outcomes. If the incumbent's policy choices produce favorable outcomes, voters' evaluation and thus propensity to reelect him improve. Conversely, negative outcomes damage the incumbent's electoral chances. This mechanism is resonant of the literature on retrospective voting (Fiorina, 1981; Achen, 1992). We build on this framework, extending the underlying logic so that voters not only evaluate individual policies and candidates, but also use their experiences to update their broader beliefs over how different policies work (i.e., the state of the world).

Correlations and Multidimensionality. Within the retrospective voting framework, 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 do not assess each issue in isolation; rather, they see issue dimensions 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, voters recognize genuine correlations between policy domains. While Converse (1964) argued that the public lacks political sophistication, more recent research suggests a growing ideological consistency among voters (Levendusky, 2010). Increasingly, voters perceive connections across policy areas (Zingher and Flynn, 2019; Hare, 2022) and, as a result, adopt more coherent belief systems (Jewitt and Goren, 2016).

Second, while voters may not be as sophisticated as this perspective suggests, they may still rely on heuristics, responding to positive outcomes in one policy domain by developing a more favorable attitude toward ideologically aligned policies in other areas. This aligns with research on motivated reasoning (Taber and Lodge, 2006; Kunda, 1987), which shows that people interpret evidence in ways that reinforce their prior attitudes. In our framework, an initially undecided voter who comes to see the merits of, say, liberal economic policies may consequently adopt consistent views across other policy dimensions.

Both perspectives point to the same conclusion: voters, whether through rational inference or cognitive shortcuts, treat policy domains as interconnected. Our model captures this intuition through the correlation parameter ρ, which governs the extent to which voters experiences inform beliefs across issue areas. In the analysis that follows, we take these cross-issue connections as given[7] and examine their implications for electoral behavior and policymakers' incentives.

4 Voter Learning

Before moving to equilibrium analysis, we first characterize how voters learn about their preferences in our model. In our framework, learning occurs through two channels. First, direct learning—the voter updates her beliefs on a specific policy dimension by observing outcomes on that dimension. Second, indirect learning—beliefs about one dimension spill over to others when policy domains are correlated. These mechanisms shape how voters process political information, form policy preferences, and ultimately, make electoral decisions. As we now show, the incumbent's policy choice crucially impacts this process. In particular, the distance of the implemented policy from the status quo is paramount. Since we normalize the status quo on each dimension to 0, for ease of exposition in what follows we refer to the distance from the status quo as extremism or moderation in policy choice.

Definition 1. *For any two policies d and d', we say that d is more extreme than d' if $|d - d_{sq}| = |d| > |d'| = |d' - d_{sq}|$, and more moderate otherwise.*

We begin by considering direct learning. We characterize the voter's *interim* posterior beliefs on a given dimension D, i.e., her beliefs on ω_d as a function of the realized outcome on dimension D only. Formally, let $\tilde{\mu}^d$ the voter's interim belief that $\omega_d = \alpha$. Then:

Lemma 1. *Direct voter learning satisfies:*

[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.

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

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

(iii) If the incumbent acts on dimension D, the probability that $\tilde{\mu}^d \neq \pi_d$ *is* $\min\{2\alpha\psi_d|d_1|, 1\}$.

This result illustrates an intuitive property: when policies remain at the status quo, no new information is generated, and the voter does not update her beliefs. But when an incumbent enacts policy change, the voter can learn about the optimal direction of reform on that dimension by observing the realized outcome. Crucially, the informativeness of policy outcomes depends on the magnitude of the reform. This property emerges starkly in our model: each policy outcome is either completely uninformative or fully informative, and the likelihood of generating an informative outcome increases in the extremism of the policy choice on that dimension. While the starkness of the learning process follows from the assumption that shocks are uniformly distributed, the result that new and more extreme policies are more informative holds more generally, and is at the core of our mechanism. We further discuss this point in section 6.3 below.

The intuition behind this result is as follows. In *expectation*, a reform that moves policy in the optimal direction for the voter generates a better outcome for the voter than one that moves in the wrong direction. However, the *realized* outcome on each dimension is also a function of a random period-specific shock, making it difficult to distinguish information from noise. Crucially, this inference problem is easier to solve when the implemented policy is more extreme. For example, if the incumbent implements radical tax changes, variations in the voter's economic welfare are increasingly likely to reflect the policy choice rather than random economic fluctuations. In contrast, modest policy adjustments make it difficult to distinguish the policy effect from chance, preventing voter learning. Figure 1 provides a graphical illustration of these results.

In our framework, however, the voter does not assess policies in isolation. Instead, she forms beliefs across multiple dimensions—filling in gaps in one domain by extrapolating from another. This leads to indirect learning, where the voter updates her beliefs on a policy dimension based on the signal she receives on another, correlated dimension. Formally, denote by $\mu^x(\tilde{\mu}^z)$ the voter's

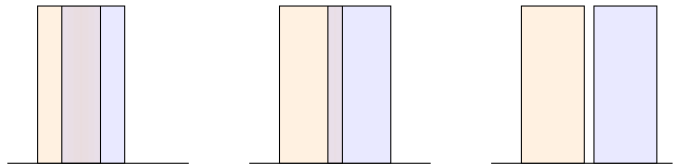

Figure 1: Voter Learning. The plots display the realized outcome on dimension D under a positive (blue function) and negative (orange function) state of the world. In the left panel, a moderate right-wing policy produces a partial overlap in the conditional distributions, leading to either no learning (within the overlap) or complete learning (outside the overlap). As policies become more extreme (middle and right panels), the conditional distributions are pulled further apart, shrinking the region of overlap until the voter always learns the state.

final posterior on ω_x, as a function of the realized outcome on *both* dimensions X and Z (via her interim beliefs $\tilde{\mu}^z \in \{0, \pi_z, 1\}$). The next result shows how information spillovers work.

Lemma 2.

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

In short, if the voter learns nothing from direct experience on a policy dimension, learning spillovers via outcomes on other correlated dimensions determine her posterior beliefs. Suppose she is uncertain about economic policy but learns that she benefits from progressive healthcare policies. If she perceives economic and healthcare policies as correlated, she will infer that progressive economic policies are likely beneficial as well. The strength of this effect depends on the correlation parameter ρ.

5 Voter Decision

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

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

Proof. All Proofs are collected in the Appendix. □

The voter re-elects the right-wing incumbent if and only if she is sufficiently confident that right-wing policies are optimal for her.[8]

6 The Incumbent's Problem

The findings from the previous sections highlight two key effects that influence the incumbent's decisions and expected payoff. The first is a static ideological effect, which is relatively straightforward. When the incumbent's policy choice aligns more closely with his own ideological preferences, his payoff in the first period increases. The second is a dynamic information effect. This effect is driven by how the incumbent's first-period decisions influence his chances of retention and, consequently, his expected second-period payoff. This, in turn, depends on voter learning, with the implemented policy influencing the likelihood of voter learning both directly and indirectly (via learning spillovers).

These two effects, ideological and informational, generate a potential trade-off for the incumbent. While he prefers to set a policy close to his ideal point, this policy may generate suboptimal levels of voter learning—either too little or too much information—from his electoral perspective.

From Lemma 1, we know that more extreme policies—those that deviate more significantly from the status quo—are more likely to generate informative outcomes for the voter. Therefore, depending on whether information is electorally advantageous, the incumbent may have incentives to distort his policy choices either towards extremes or towards the status quo.

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

As we will see in the following sections, whether the incumbent chooses to move to the extremes or stick with the status quo depends on two crucial factors: the incumbent's initial electoral standing and his decision to focus on a single policy dimension or to broaden the scope of his policymaking. To clearly describe the incumbent's incentives, it is then useful to introduce the following definition:

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

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

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

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

Recall that $\frac{1}{2\alpha\psi_x}$ is the smallest policy that guarantees direct learning on dimension d (from Lemma 1). This assumptions guarantees that the incumbent's electoral incentives influence his policy choice,[9] and that the incumbent always chooses to act on the primary dimension X.

6.1 A Unidimensional Benchmark

To sharpen our understanding of the strategic incentives at play, we begin by shutting down learning spillovers (formally, setting $\rho = \frac{1}{2}$). Under assumption 1, this implies that the incumbent's policy choice on the secondary dimension Z has no bearing on his retention chances: the voter's posterior beliefs and electoral decision depend exclusively on the outcomes in dimension X. Denote d_u the equilibrium policy choice of dimension d in this unidimensional benchmark. We have:

Proposition 1. *Let $\lambda_I^z = \lambda_C^z = 0$ and $\rho = \frac{1}{2}$. In equilibrium:*

The incumbent does not act on the secondary dimension Z. On the primary dimension X:

[9] Absent this assumption, a trailing incumbent's *statically* and *dynamically* optimal policies sometimes coincide.

(i) A leading incumbent implements a policy more moderate than his bliss point, $x_u(\pi) < x_I$*;*

(ii) A trailing incumbent implements a policy more extreme than his bliss point, $x_u(\pi) > x_I$*.*

With no spillovers from dimension Z, the incumbent's policy on that dimension is irrelevant for electoral outcomes. Since the incumbent is also indifferent about Z from an ideological perspective, he simply leaves the status quo unchanged. In contrast, the implemented policy on X determines the probability of the voter observing an informative outcome and, thus, the incumbent being reelected. As a consequence, the incumbent's policy choice on the primary dimension is distorted away from his ideological preference. The direction of this distortion depends on whether the incumbent is ex-ante leading or trailing.

Consider an incumbent who is ex-ante trailing—that is, the voter starts off preferring the challenger. If the voter fails to receive new information, she defaults to replacing the incumbent. This means the incumbent's only path to reelection is if the voter learns something new—specifically, that right-wing policies are optimal for her. Recall that more extreme policies facilitate learning. Then, the incumbent has incentives to gamble and distorts policy to the extreme, away from both his static optimum and the status quo (normalized to 0).

For an incumbent who is ex-ante leading, the logic flips. If the voter learns nothing new, the incumbent is reelected with certainty. Any new information introduces risk. Thus, leading incumbents distort policy towards the status quo 0 so as to minimize information.

This unidimensional benchmark provides a useful baseline, which aligns with the conventional intuition on "gambling on resurrection". In the next sections, we introduce multidimensionality and explore how the presence of policy spillovers fundamentally reshapes strategic incentives.

6.2 A Multidimensional World

We now allow for $\rho > \frac{1}{2}$. This introduces learning spillovers: the voter's beliefs about X can be influenced by policy outcomes on Z. Even though the incumbent's ideological preferences remain unidimensional (under Assumption 1), his strategic incentives become multidimensional. Our goal is to determine when the incumbent has an incentive to act on the secondary dimension and how this affects his choices on the primary dimension.[10] To guide this analysis, we introduce the following:

Definition 3. *Let*

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

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

This definition captures cases where learning spillovers are strong enough to meaningfully influence the voter's electoral decision. Specifically, if the correlation ρ is low, then outcomes on Z are electorally inconsequential. This is because the spillovers effects are too weak, and the realized outcome on Z has a negligible impact on the voter's beliefs on X. Conversely, if ρ is high, outcomes on the secondary dimension have a sufficient impact on the voter's beliefs about the primary issue that, in the absence of direct learning, they determine her electoral decision. 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 (i.e., the voter learns that $\omega_z = -\alpha$). Conversely, in a symmetrical fashion, if a trailing incumbent successfully generates favorable information regarding the secondary dimension (i.e., an outcome revealing that $\omega_z = \alpha$), the spillover effects of learning become pivotal in propelling him toward re-election.

Our next result, characterizing the conditions under which the incumbent pursues policymaking on the secondary dimension despite having no ideological preferences for doing so, follows straightforwardly from this discussion:

Proposition 2. *Let $\lambda_I^z = \lambda_C^z = 0$. In equilibrium, the incumbent acts on Z if and only if he is trailing and the correlation with the primary dimension is high.*

A leading incumbent never has strategic incentives to act on Z, since he wants to prevent the voter from obtaining any new information. For a trailing incumbent, the logic flips. His only path to

[10] Buisseret and Van Weelden (2022) also examine multidimensional incentives, although their focus is on an incumbent's decision to call a referendum on a secondary policy issue to reveal information about the distribution of voters to parties, rather than manipulate voter learning.

reelection is through shifting voter beliefs—and to do that, he needs to create learning opportunities. If correlation is high, then spillovers from Z can meaningfully influence voter beliefs on X. In this case, the incumbent expands the scope of policymaking, hoping that a favorable policy outcome in Z will indirectly shift voter beliefs in his favor on X, pushing him above the reelection threshold.

However, if correlation is low, then spillovers are too weak to be electorally relevant. In this scenario, the trailing incumbent gains nothing from acting on Z and remains indifferent between intervention and inaction. By assumption, he then chooses the status quo.

Building on this result, we next characterize how the incumbent's policy choice on X in this multidimensional setting compares to the equilibrium policy in the unidimensional benchmark. Our central result uncovers a strategic substitution effect, where a trailing incumbent pursues extreme policies on the secondary dimension and simultaneously moderates his policy choice on the primary dimension. Recall that $x_u(\pi)$ is the incumbent's optimal policy choice in the unidimensional benchmark. Then, we have:

Proposition 3 (Substitution Effect). *Let $\lambda_I^z = \lambda_C^z = 0$. Suppose that the correlation between the two dimensions is high and the incumbent is trailing, so that in equilibrium he chooses to act on the secondary dimension Z. Then:*

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

(ii) on the secondary dimension, $z_1^*(\pi) \geq \frac{1}{2\alpha\psi_z}$.

Recall that $\frac{1}{2\alpha\psi_z\lambda_v^z}$ is the smallest reform that guarantees full learning on Z. Then, Proposition 3(ii) is intuitive. Even though the incumbent has no ideological preference over Z, his strategic incentives to facilitate voter learning drive him to pursue extreme policies on this dimension. This result highlights a striking implication: in our setting, extreme policymaking is not necessarily driven by extreme ideological preferences. Instead, it can emerge purely as an electoral strategy, even on issues where the incumbent has no intrinsic ideological attachment.

At first glance, one might expect that a trailing incumbent, facing reelection pressure, would pursue extreme policies on both dimensions—maximizing the chance of a favorable voter update.

Surprisingly, Proposition 3(i) reveals the opposite: high correlation between dimensions actually induces moderation on the primary dimension. That is, while in the unidimensional benchmark a trailing incumbent pushes policy to the extreme ($x_u(\pi) > x_I$), in the multidimensional case, he moderates instead ($x_1^*(\pi) < x_I$). While this result is in sharp contrast with the conventional logic on "gambling on resurrection", it still stems from the incumbent's incentive to shape voter learning. Specifically, the reason underlying this strategic substitution effect is that distorting policy to the extreme on both dimensions would undermine the very spillover effect the incumbent is trying to exploit.

To understand the logic, recall that the outcome on Z can influence the voter's retention decision only if she does not learn about ω_x directly (as otherwise she reaches a degenerate interim posterior). 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 contrast with the results of the unidimensional baseline, then, this scenario generates incentives for the trailing incumbent to pursue *moderate* policies on the primary dimension.

Notice that, in principle, the incumbent could find it optimal to forgo the learning spillovers and gamble on X instead. Put differently, a trailing incumbent might want to skip healthcare policymaking altogether and instead gamble on radical economic policy. However, we know from Proposition 2 that this never occurs in equilibrium. The reason lies in the asymmetry of information across dimensions. Ex ante, players have less information on the secondary dimension than on the primary one—they know better about the possible consequences of specific economic policies than they do about policies on healthcare.

Because of this asymmetry, when the incumbent is trailing (i.e., $\pi_x < \frac{1}{2}$), the ex-ante probability of generating a favorable outcome is higher on the secondary dimension than on the primary one (i.e., $\pi_z > \pi_x$). Consequently, a trailing 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 ω_x as well, even when the optimal policy on the primary dimension is actually a left-wing one.

Given these incentives, in equilibrium the incumbent will never gamble on both dimensions, and will never choose to forego gambling on Z when he can benefit from it. Rather, if the correlation is too low to exploit learning spillovers, he will have no strategic incentives to act on Z and will continue gambling on X. If instead the correlation is high, he will gamble on Z but prefer to avoid risky choices on X.

In concluding this section, we note that one may worry that the substitution effect arises solely because, under Assumption 1, the voter's beliefs about the secondary dimension do not directly influence her retention choice (i.e., outcomes on Z only matter via learning spillovers). In Appendix B we show that this is not the case: the effect emerges even when this assumption is relaxed. We discuss the results of this general model in Section 7 below.

6.3 Scope Conditions and Robustness

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

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

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

Voter Learning with Non-Uniform Shocks. As emphasized above, an essential element of our theory is that the amount of voter learning is endogenous to the implemented policy. On each dimension, more extreme policies facilitate voter learning. In our model, the voter's learning technology is stark. Due to the assumption that the shock $\varepsilon_{d,t}$ is drawn from a uniform distribution, the voter's realized utility is either fully informative or completely uninformative. The more extreme the policy, the higher the likelihood that it generates an informative outcome realization. It is important to emphasize that this information property of extreme policies holds more generally, beyond this specific distributional assumption.[11]

7 Empirical Implications

To explore the theory's empirical implications, in Appendix B we extend the model to a more general setting where both the incumbent and challenger hold ideological preferences over both dimensions, relaxing Assumption 1. For this generalized model, we need to make assumptions on how the voter weighs competing information when forming her electoral decision. Specifically, if she learns that the state is α on one dimension and $-\alpha$ on the other, which dimension drives her choice? Here, we assume the voter prioritizes the primary dimension (see Assumption 4 in the Appendix). Thus, the primary dimension is also the more salient electorally. This aligns with our premise that voters are initially more informed about this dimension than the secondary one.

In contrast to the model analyzed in the previous sections, the incumbent now has both ideological preferences and strategic incentives to control information over the secondary dimension. For a trailing incumbent, both forces align: he benefits from pursuing reforms on this dimension. A leading incumbent, however, faces a tradeoff. If his ideological preferences on the secondary dimension are sufficiently strong, he will embrace a multidimensional agenda, implementing reforms

[11] In particular, extreme policies are more informative whenever the distribution of $\varepsilon_{d,t}$ satisfies the monotone likelihood ratio property (MLRP). For example, with normally distributed shocks, learning occurs more gradually, but extreme policies still provide more information by increasing the distance between expected outcomes under different states. We provide a formal proof of this result in Lemma A-1 in the online Appendix.

on both dimensions (see Proposition B-1 in the Appendix).

Proposition B-2 in the Appendix then establishes that, when the incumbent acts on both dimensions, the substitution effect we identify above persists under high correlation, emerging for both trailing and leading incumbents. Compared to his choice in the unidimensional benchmark, a trailing incumbent implements a more moderate policy on on the primary dimension, $x_1^* < x_u$, simiarly to what described in the previous section. In contrast, a leading one adopts the opposite approach: To mitigate electoral risks from negative spillovers from the secondary dimension, leading incumbents pursue more informative, and thus extreme, policies on the primary dimension, $x_1^* > x_u$. Furthermore, in sharp contrast with conventional intuition and in line with Proposition 3, under some conditions this substitution effect is so pronounced that a leading incumbent adopts a primary-dimension reform that is even more extreme (relative to the status quo) than his static ideological bliss point, $x_1^* > x_I$, actively *facilitating* direct learning on X. Symmetrically, a trailing incumbent sometimes becomes risk-averse on the primary dimension and distorts policy towards the status quo, $x_1^* < x_I$,—a sharp departure from their typically bold approach in a unidimensional world.

While these predictions compare multidimensional policy choices to a theoretical unidimensional benchmark that may be difficult to observe in practice, the logic of the substitution effect generates distinct, testable implications. The key insight is that the relationship between issue salience and policy extremism depends critically on the incumbent's electoral standing. All else equal, we should expect *trailing* incumbents to adopt more moderate stances on primary (more salient) issues, and go more extreme on less salient ones. For *leading* incumbents, the relation should have the opposite sign. The next result formally establishes this conditional effect, fixing the value of the parameters such that $x_I = z_I$, $\psi_z = \psi_x$ and $\lambda_x^I = \lambda_z^I$.[12]

Implication 1. *If the incumbent is leading, then the implemented policy is farther from the status quo on the primary dimension, i.e., the electorally most salient one. If instead the incumbent is*

[12]Implication 1 focuses on dimensions that the voter perceives as sufficiently correlated, so that learning spillovers are significant. Empirically, we believe this to be the case for most issues that are on the policy agenda: as discussed in section 3.1, voters increasingly see different issues are connected, and form their policy preferences accordingly.

trailing, then the policy is farther from the status quo on the secondary dimension, i.e., the electorally less salient one.

We note that although our model focuses on two dimensions, the result applies more generally. Specifically, when considering more than two dimensions, the extremism of the implemented policy relative to the status quo decreases for dimensions that are ranked lower in salience when the incumbent is leading, and increases for those ranked lower when the incumbent is trailing. Highlighting this is important, as our data cover more than two issue dimensions.

To our knowledge, the relationship between policy extremism, issue salience, and the incumbent's electoral prospects has not been systematically examined in empirical research. To bridge this gap, we outline a roadmap for empirical analysis, providing a guide for scholars interested in testing the key implications of our model. As an initial step, we discuss findings from a preliminary "reality check," leveraging existing data from the U.S. context.

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

The dependent variable related to our central result is the extremity of policy change on different policy dimensions. To measure it, we must compute the distance between the location of the status quo policy on each policy dimension, and the location of the implemented policy (if any). For this purpose, we use data from Lowande et al. (2021), who uses legislative candidate responses to policy questions—indicating whether policies in specific issue areas was too liberal or conservative—to generate estimates of the status quo policy locations in each of the surveyed issue areas for the 103rd through 114th Congresses. Then, we compare the location of the status quo at the end of one congressional session with the status quo in that same issue area at the beginning of the session.[13] Matching our theoretical quantity, this implies that reforms are coded as more extreme the larger the distance from the inherited status quo.

The next key variable of our theory captures primary and secondary policy dimensions. Here, empirical tests should use measures of issue salience. We make use of Heffington, Park and Williams's "Most Important Problem" Dataset (2019). Public opinion scholars and survey firms have asked U.S. citizens about the the "most important problem facing the United States" for decades, and this dataset constitutes the largest compilation of these questions available. Much as political scientists turn to these questions to measure the issue priorities of American over time, politicians also track results of such surveys as they are regularly published by news agencies. We incorporate the survey responses by ranking, for each Congressional session, the issues according to how frequently respondents identify them as their top concern.

Finally, our model requires a measurement of whether an incumbent—in this case, an incumbent president—is "leading" or "trailing." Ideally, we would have a measure of incumbents' beliefs about whether they were likely to win or lose in the upcoming election. Such data do exists for U.S. presidents—the Iowa Electronic Markets provide a strong example—but do not extend beyond a few months prior to an election.

Lacking this data, we adopt two approaches. First, we use presidential approval ratings compiled by the Roper Center. We proceed by subtracting presidential disapproval rates from overall approval rates. Using this 'net approval' measure, we then code a president as "leading" if his approval rating is a net-positive, and trailing otherwise. While presidential approval is a good proxy of the president's own popularity, it does not capture broader features of the environment that would, according to our theory, influence the incumbent's strategic incentives. More specifically, the president may be concerned not only about his own popularity, but about the chances of his party obtaining a majority in congress (so that even term-limited incumbents may face incentives similar to the ones highlighted in our theory). To capture these concerns, we also provide a quali-

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

tative measurement that classifies presidents as "leading" or "trailing" also taking into account the popularity of their party for each Congress. This allows us to capture cases when the president retains fairly strong popularity due to outside factors like the honeymoon effect, even when his party is faring quite poorly. During Pres. Obama's first term, for example, Democrats' political struggles eventually led to historic losses in the 2010 election—despite the fact that Obama retained relatively high approval ratings throughout.

One specific reason why the president may care about the composition of Congress is that an aligned legislature facilitates his action in office. Straightforwardly, these considerations are moot for a second-term president during his last two years in office, since no other congressional election will occur during his tenure. Thus, we consider the full sample of observations (N = 106) as well as restricted sample where we drop the last two year of a president's second term (N = 78).[14] We expect our results to be stronger in our restricted sample, where strategic incentives to control information should be more pronounced.

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

These results, of course, are noisy, and we do not aim to over-interpret the associations we depict. However, they do underscore how our theory may help empirical researchers better understand

[14]These sample sizes are not exact multiples of the number of Congresses in our sample (12) due to some issue-level missingness in Lowande's replication data.

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

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

important outcomes like policy extremism, by disentangling interrelated factors like issue salience and politicians' electoral security. That being said, more research is doubtlessly needed to probe the robustness of these results, both within and outside the US.

Additional Implications. To conclude this section, we note that our theory's implications may shed light on a puzzling empirical pattern: the mixed evidence on the relationship between issue salience and policy congruence—that is, the extent to which policymakers implement policies aligned with the majority's preferences. Intuition suggests a positive correlation: the more salient an issue, the stronger the incentive for policymakers to align with voter preferences. Yet, empirical findings

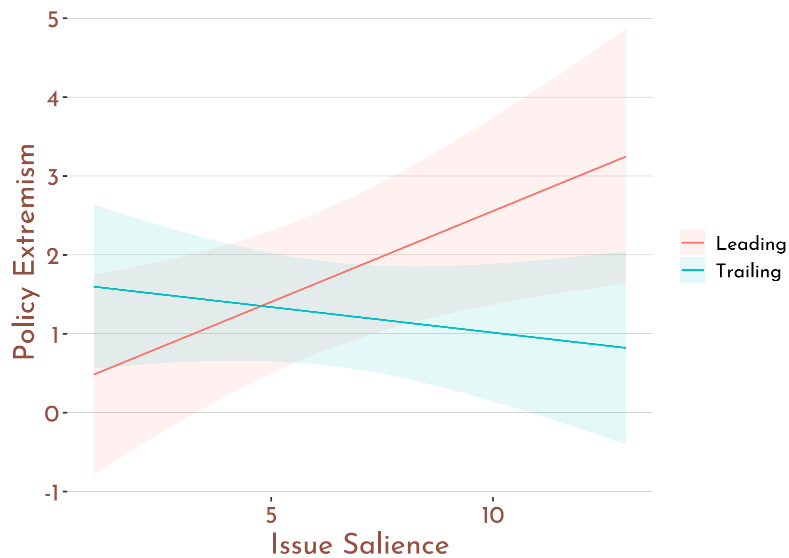

Figure 2: Conditional Relationship between Issue Salience and Policy Extremism

are inconsistent (Canes-Wrone and Shotts, 2004; Rasmussen, Reher and Toshkov, 2019).

Our model offers a potential explanation. The substitution effect implies that the link between issue salience and policy congruence is not uniform but instead shaped by the incumbent's electoral standing. To see why, consider a case where the voter's preferred policy (given her *prior* beliefs) is relatively moderate. Implication 1 suggests that leading incumbents, seeking to suppress voter learning, should exhibit greater policy congruence on secondary, low-salience dimensions, while trailing incumbents—who need to generate voter learning—should be more congruent on high-salience issues. This conditional relationship means that empirical studies must account for the incumbent's electoral position when assessing how issue salience influences policy extremism.

8 Conclusion

This paper develops a theory of policymaking in a multidimensional world, where voters update their preferences upon experiencing the consequences of different policies, and incumbents balance ideological goals against electoral survival. Conventional wisdom suggests trailing incumbents pursue radical reforms while leading ones moderate. Our model refines this by showing that with correlated issues, voter learning creates a strategic substitution effect: incumbents balance extremism in some areas with moderation in others. This effect is mediated by the incumbent's electoral standing and the electoral salience of each policy area.

Leading incumbents, typically expected to pursue cautious, centrist policies, may instead take bold positions on high-salience issues while moderating on secondary ones—an attempt to avoid negative spillovers from voter learning on these secondary issues. Conversely, trailing incumbents, needing to shift voter beliefs, may gamble on extremism in less salient domains while moderating their positions on the most politically central issues. We bring our predictions to the data and find preliminary supporting evidence, and outline directions for future research to systematically evaluate our theory's implications.

Beyond its immediate empirical implications, this paper contributes to a broader reconsideration of multidimensional politics. Traditional models often reduce policymaking to a single dominant dimension. This perspective is typically justified by the observation 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 a multidimensional world may be fundamentally different from the unidimensional case.

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.

Blumenthal, Benjamin. 2025. Environmental Policymaking with Political Learning. Technical report.

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

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

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

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

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

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

Downs, George W and David M Rocke. 1994. "Conflict, agency, and gambling for resurrection: The principal-agent problem goes to war." *American Journal of Political Science* pp. 362–380.

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

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

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

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

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

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

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

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

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

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

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

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

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

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

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

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

Rose-Ackerman, Susan. 1980. "Risk taking and reelection: Does federalism promote innovation?" *The Journal of Legal Studies* 9(3):593–616.

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

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

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

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

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

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

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

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

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

Table of Contents

Appendix A: Main Results

Proof of Lemma 1. We prove the statements for dimension X. Let $\mu^x \in [0,1]$ denote V's posterior that the state of the world on dimension X is positive, i.e., that $\omega_x = \alpha$. (i) A possible outcome realization o_t^x given the incumbent's choice (x_t), and conditioning on the true state ω_x has to fall within:

$$\left[\omega_x x_t - \frac{1}{2\psi_x}, \omega_x x_t + \frac{1}{2\psi_x}\right]. \tag{A-1}$$

Suppose $x_t > 0$. A symmetric reasoning applies for $x_t < 0$. We can immediately see that if V observes $o_t^x > -\alpha x_t + \frac{1}{2\psi_x}$, she knows for sure that $\omega_x = \alpha$, i.e., $\mu^x = 1$. Similarly, if V observes $o_t^x < \alpha x_t - \frac{1}{2\psi_x}$, then $\mu^x = 0$. The last case to consider is when o_t^x falls within the interval $\left[\alpha x_t - \frac{1}{2\psi_x}, -\alpha x_t + \frac{1}{2\psi_x}\right]$. Denote by $f(\cdot)$ the PDF of the error term $\varepsilon_{x,t}$. When o_t^x falls within this interval we have that:

$$\Pr(\omega_x = \alpha | o_t^x) = \frac{f(o_t^x - \alpha x_t)\pi_x}{f(o_t^x - \alpha x_t)\pi_x + f(o_t^x + \alpha x_t)(1-\pi_x)}. \tag{A-2}$$

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

$$\Pr(\omega_x = \alpha | o_t^x) = \Pr(\omega_x = \alpha) = \pi_x.$$

(ii)-(iii) Now, denote by $L_x \in \{0,1\}$ players' learning of ω_x. There exists a value of policy x_t' such that, for any $x_t > x_t'$, the realization of o_t^x is fully informative, i.e., the interval (A-1) is empty. This requires:

$$\alpha x_t - \frac{1}{2\psi_x} \geq -\alpha x_t + \frac{1}{2\psi_x} \tag{A-3}$$

which rearranged yields:

$$x_t \geq \frac{1}{2\alpha\psi_x}. \tag{A-4}$$

Define $x' \equiv \frac{1}{2\alpha\psi_x}$, and assume $x_t \in [0, x']$. We have:

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

Since the two probabilities are symmetric, we have

$$\begin{aligned}\Pr(L_x = 1|\pi_x, 0 < x_t < x') &= \Pr(\alpha x_t + \varepsilon_{x,t} > -\alpha x_t + \frac{1}{2\psi_x}) \\ &= \Pr(\varepsilon_{x,t} > -2\alpha x_t + \frac{1}{2\psi_x}) \\ &= \Pr(\varepsilon_{x,t} < 2\alpha x_t - \frac{1}{2\psi_x}) \\ &= 2\alpha\psi_x x_t. \end{aligned} \tag{A-5}$$

The proof for dimension Z is analogous therefore omitted. □

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

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

$$\alpha d' - (-\alpha)d' > \alpha d - (-\alpha)d,$$

which simplifies to

$$2\alpha d' > 2\alpha d,$$

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

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

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

Similarly, suppose that the voter learns that $\omega_z = \alpha$ on dimension Z. By Bayes' rule, we have:

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

□

Proof of Remark 1. The voter re-elects the right-wing incumbent if and only if the expected payoff from voting for I given the information received in $t = 1$ is greater than that of voting for C. Under Assumption 1, the voter expects both challenger and incumbent to leave policy on dimension Z unchanged if elected in period 2. Thus, we need only compare expected utility on dimension X. Given that outcomes are determined by $o_t^x = \omega_x x_t + \varepsilon_{x,t}$, the voter's expected utility from politician i implementing policy x_i in period 2 is:

$$\begin{aligned} E[u_{v,2}(i)] &= \lambda_v^x E[o_2^x] \\ &= \lambda_v^x E[\omega_x x_i + \varepsilon_{x,2}] \\ &= \lambda_v^x[\mu^x \alpha x_i + (1-\mu^x)(-\alpha)x_i] \\ &= \lambda_v^x x_i \alpha[2\mu^x - 1] \end{aligned} \tag{A-8}$$

Thus, the voter prefers the incumbent if and only if:

$$\lambda_v^x x_I \alpha[2\mu^x - 1] > \lambda_v^x x_C \alpha[2\mu^x - 1] \tag{A-9}$$

which, given $x_I > x_C$, simplifies to $\mu^x > \frac{1}{2}$. □

Proof of Proposition 1. By Remark 1, the incumbent is re-elected if and only if $\mu_x > 1/2$. Suppose the incumbent is leading. By Definition 2, 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 $\omega_x = \alpha$, then $\mu^x = 1$ and the incumbent is re-elected. It is only when V learns that $\omega_x = -\alpha$ that $\mu^x = 0$. In this case, the incumbent is ousted. Suppose now the incumbent is trailing ($\pi_x \leq 1/2$). Analogously to the argument above, it is only when when V learns that $\omega_x = \alpha$ that $\mu_x > 1/2$ and the incumbent is re-elected. Next, notice that given the assumption that $|x_I| < |\frac{1}{2\alpha\psi_x}| = |x'|$ and the fact that any policy $|x| \geq |x'|$ guarantees $L_x = 1$ with probability 1, in equilibrium the incumbent never implements a policy $|x_1| > |x'|$. Consider now the incumbent's problem:

$$U_I(x_1) = \lambda_I^x \left[f_I^x(x_1) + \mathbb{P}(x_1, z_1)(f_I^x(x_I) - f_I^x(x_C)) + f_I^x(x_C) \right], \tag{A-10}$$

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

$$U_I(x_1) = \lambda_I^x \left\{ f_I^x(x_1) + \Big[1 - 2\alpha\psi_x x_1(1 - \pi_x)\Big]\Big[f_I^x(x_I) - f_I^x(x_C)\Big] + f_I^x(x_C) \right\}. \tag{A-11}$$

Letting $K = \lambda_I^x(f_I^x(x_I) - f_I^x(x_C))$, we can write the first-order condition as:

$$\lambda_I^x \frac{\partial f_I^x(x_1)}{\partial x_1} - 2\alpha\psi_x(1 - \pi_x)K = 0. \tag{A-12}$$

Then, x_1^* is equal to $max \in \{0, \widehat{x}\}$, where $\widehat{x}$ is the policy that solved (A-12). Recall that $K > 0$. Because f_I^x is single-peaked, $\frac{\partial f_I^x(x_1)}{\partial x_1} > 0$ iff $x_1 < x_I$. Thus, $x_1^* < x_I$.

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

$$U_I(x_1) = \lambda_I^x f_I^x(x_1) - (1 - 2\alpha\psi_x x_1 \pi_x)K. \tag{A-13}$$

which yields the following first-order condition:

$$\lambda_I^x \frac{\partial f_I^x(x_1)}{\partial x_1} + 2\alpha\psi_x\pi_x K = 0.$$

Thus, it must be the case that $x_1^* > x_I$. □

Lemma A-2. *Let $\lambda_I^z = \lambda_C^z = 0$ and $\lambda_v^z > 0$.*

If the correlation across dimensions (ρ) is low, the incumbent's probability of winning is

- $1 - 2\alpha\psi_x(1-\pi_x)x_1$ *when* $\pi_x > \frac{1}{2}$;
- $2\alpha\psi_x\pi_x x_1$ *when* $\pi_x < \frac{1}{2}$.

If ρ is high, the incumbent's probability of winning is

- $1 - 2\alpha\psi_x(1-\pi_x)x_1 - (1-2\alpha\psi_x x_1)2\alpha\psi_z(1-\pi_z)z_1$ *when* $\pi_x > \frac{1}{2}$;
- $2\alpha\psi_x\pi_x x_1 + (1-2\alpha\psi_x x_1)2\alpha\psi_z\pi_z z_1$ *when* $\pi_x < \frac{1}{2}$.

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

$$\mu^x(\varnothing, -\alpha, \rho) = \frac{1}{2}, \tag{A-14}$$

where

$$\mu^x(\varnothing, -\alpha, \rho) = \frac{(1-\rho)\pi_x}{(1-\rho)\pi_x + \rho(1-\pi_x)}, \tag{A-15}$$

which yields:

$$\widehat{\rho} = \pi_x. \tag{A-16}$$

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

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

$$\mu^x(\varnothing, \alpha, \rho) = \frac{1}{2}, \tag{A-17}$$

where

$$\mu^x(\varnothing, \alpha, \rho) = \frac{\pi_x \rho}{\pi_x \rho + (1-\pi_x)(1-\rho)}. \tag{A-18}$$

Combining the above, we have

$$\widehat{\rho} = 1 - \pi_x. \tag{A-19}$$

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

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

Plugging in the values of $\Pr(L_x(x_1) = 1)$ and $\Pr(L_z(z_1) = 1)$, we obtain the results. □

Proof of Proposition 2. The proof follows from Lemma A-2 and the assumption that when indifferent the incumbent prefers not to act on Z. □

Proof of Proposition 3. (i) Consider the incumbent's choice on X. Recall that in equilibrium we must have $x_1 \leq x' = \frac{1}{2\alpha\psi_x}$. Then, when I is trailing and $\rho > \widehat{\rho}$, we have

$$\mathbb{P} = 2\alpha\psi_x x_1 \pi_x + (1 - 2\alpha\psi_x x_1)p^\dagger \pi_z, \tag{A-20}$$

where $p^\dagger = \min\{1, 2\alpha\psi_z z_1\}$. 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, $z_1^* \geq \frac{1}{2\alpha\psi_z}$. (ii) Given point (i), the trailing incumbent's retention probability reduces to

$$\mathbb{P} = 2\alpha\psi_x\pi_x x_1 + (1 - 2\alpha\psi_x x_1)\pi_z. \tag{A-21}$$

Plugging (A-21) into the incumbent's problem yields the following first-order condition:

$$(x_1) \quad \lambda_I^x \frac{\partial f_I^x(x_1)}{\partial x_1} + 2\alpha\psi_x(\pi_x - \pi_z)K = 0. \tag{A-22}$$

Note that, given $\pi_x < \frac{1}{2}$, $\pi_z = \pi_x\rho + (1-\pi_x)(1-\rho) > \pi_x$. Since $K > 0$ and $f_I^x(\cdot)$ is single-peaked, it follows that $x_1^* < x_I$. □

Appendix B: General Model

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

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

which rearranged yields:

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

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

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

Assumption 3. $\lambda_v^x > \lambda_v^z \frac{z_I - z_C}{x_I - x_C}$.

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

Next, we characterize the incumbent's probability of winning. Suppose that the incumbent is leading ($\pi_x > \frac{1}{2}$). Then, generalizing Definition 3, denote $\widehat{\rho}_g$ the value that solves:

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

where

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

which yields:

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

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

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

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

where

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

Combining the above, we have

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

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

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

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

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

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

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:

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

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

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

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

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

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

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

We can now use these preliminary results to characterize the equilibrium policies. In particular, we are interested in establishing two results. First, that there exist conditions under which the incumbent reforms policy on Z in the first period. Second, that under some conditions the substitution effect identifies in the previous section continues to emerge, and for both trailing and leading incumbents.

Characterization of Equilibrium Policies.

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

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

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

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

   - *If the precision of the* Z*-dimension shock* ε_z (ψ_z) *is sufficiently low, then a leading incumbent acts on dimension* Z*.*

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

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

$$\lambda_I^z \frac{\partial f_I^z(z_1)}{\partial z_1}|_{z_1=0} + K(1 - 4\alpha\psi_x\lambda_v^x x_1\pi_x)4\alpha\psi_z\lambda_v^z\pi_z < 0, \tag{B-16}$$

which can never be satisfied since $\lambda_I^z \frac{\partial f_I^z(z_1)}{\partial z_1}|_{z_1=0} > 0$ and $1 - 4\alpha\psi_x\lambda_v^x x_1\pi_x > 0$ in equilibrium. Thus a trailing incumbent must always set $z_1^* > 0$.

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

$$\lambda_I^z \frac{\partial f_I^z(z_1)}{\partial z_1}|_{z_1=0} - (1 - \pi_z)(1 - 4\alpha\widehat{x}\psi_x\lambda_v^x)4\alpha\psi_z\lambda_v^z K < 0, \tag{B-17}$$

where $\widehat{x}$ is the optimal policy on X when $z_1 = 0$, and is not a function of ψ_z.

We can immediately see that the LHS of (B-17) 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 $z_1^* > 0$. □

Having established (sufficient) conditions under which incumbents engage in policymaking on the secondary dimension, we now turn to characterizing the strategic substitution effect in this generalized setting. The previous analysis in Proposition 3 compared the equilibrium policy on the primary dimension when $\rho > \frac{1}{2}$ to the policy emerging in the unidimensional benchmark ($\rho = \frac{1}{2}$), while holding $\lambda_I^z = 0$ fixed. This assumption ensured that any differences in equilibrium policy resulted solely from variations in voter learning rather than from changes in the incumbent's ideological incentives. In the general model, where $\lambda_I^z > 0$, the incumbent's value from re-election naturally increases compared to the unidimensional benchmark, as re-election allows the incumbent to avoid ideological costs on both dimensions. To ensure a meaningful comparison that isolates the learning effect, we need to account for this difference in the value of re-election. Specifically, we compare equilibrium policies in the general model to those in a unidimensional benchmark with an augmented value of re-election that equalizes the overall electoral stakes. Formally, the augmented value of re-election in the unidimensional scenario is defined as $K' = \lambda_I^x\big(f_I^x(x_I) - f_I^x(x_C)\big) + \Delta$, whereas the value of re-election in the general scenario is $K = \lambda_I^x\big(f_I^x(x_I) - f_I^x(x_C)\big) + \lambda_I^z\big(f_I^z(z_I) - f_I^z(z_C)\big)$, with Δ set such that $K = K'$. Importantly, while adding Δ alters equilibrium policy in the unidimensional benchmark, it does not change the *direction* of the distortion relative to the static optimum. Thus, Proposition 1 continues to hold under this augmented value of re-election.

Proposition B-2 (Substitution Effect, General Model). *Let $\lambda_v^z, \lambda_I^z, \lambda_C^z > 0$. Suppose that the incumbent chooses to act on the secondary dimension in equilibrium. Everything else being equal, when the correlation with the primary dimension is low we have $z_1^* = z_I$ and $x_1^*(\pi) = x_u(\pi)$. Suppose instead the correlation is high. Then, everything else being equal:*

- *when the incumbent is trailing, we have $z_1^* \geq z_I$ and $x_1^* \leq x_u$;*
- *when the incumbent is leading, we have $z_1^* \leq z_I$ and $x_1^* \geq x_u$.*

Proof of Proposition B-2. If $\rho < \widehat{\rho}_g$ then the incumbent's probability of winning does not depend on z_1, making his strategic problem on dimension X identical to the unidimensional case. On Z, he simply implements his ideologically preferred policy since z_1 does not influence his retention chances.

Suppose instead, $\rho > \widehat{\rho}_g$. Consider first a trailing incumbent. A sufficient condition to ensure that $z_1^* \geq z_I$ is that $\frac{\partial \mathbb{P}}{\partial z_1} > 0$, which is always true in equilibrium for any value of z_1 and x_1. From the KKT conditions, and given $z_1^* > 0$ and concavity of f_I^x, a sufficient condition to ensure that $x_1^* \leq x_u$ is that $\frac{\partial \mathbb{P}}{\partial x_1}|_{z_1=0} > \frac{\partial \mathbb{P}}{\partial x_1}|_{z_1>0}$. Plugging in the value of $\mathbb{P}$, this condition becomes

$$2\alpha\psi_x\pi_x > 2\alpha\psi_x\pi_x - 2\alpha\psi_x 2\alpha\psi_z\pi_z z_1, \tag{B-18}$$

which is always satisfied.

Next, consider a leading incumbent. A sufficient condition to ensure that $z_1^* \leq z_I$ is that $\frac{\partial \mathbb{P}}{\partial z_1} < 0$, which is always true in equilibrium for any value of z_1 and x_1. From the KKT conditions, and given

$z_1^* > 0$ and concavity of f_I^x, a sufficient condition to ensure that $x_1^* \geq x_u$ is that $\frac{\partial \mathbb{P}}{\partial x_1}|_{z_1=0} < \frac{\partial \mathbb{P}}{\partial x_1}|_{z_1>0}$. Plugging in the value of $\mathbb{P}$, this condition becomes

$$-2\alpha\psi_x(1-\pi_x) < -2\alpha\psi_x(1-\pi_x) + 2\alpha\psi_x 2\alpha\psi_z(1-\pi_z)z_1, \tag{B-19}$$

which is always satisfied. □

The next result establishes sufficient conditions under which, in sharp contrast to a unidimensional world, a leading incumbent implements a policy more extreme than his static optimum, and a trailing one a more moderate one. To facilitate the analysis, it is useful to introduce here a discount factor δ.

Proposition B-3. *Suppose the incumbent discounts second-period payoffs by a factor $\delta \in (0,1]$. If δ if sufficiently low and z_I is sufficiently large, then*

- *A leading incumbent implements $x_1^* > x_I$;*
- *A trailing incumbent implements $x_1^* < x_I$.*

Proof. First, consider a leading incumbent. Sufficient condition to ensure that $x_1^* > x_I$ is that the probability of winning is increasing in x_1, given the choice on the Z dimension. Plugging in the probability of winning, this requires

$$2\alpha\psi_x\Big[2\alpha\psi_z z^\dagger(1-\pi_z) - (1-\pi_x)\Big] > 0, \tag{B-20}$$

where $z^\dagger = \min\{z_1, \frac{1}{2\alpha\psi}\}$. Notice that this is always satisfied when $z_1 \geq \frac{1}{2\alpha\psi_z}$, since for a leading incumbent $\pi_x > \pi_z$. For an arbitrarily low δ and a sufficiently large z_I, we must have that in equilibrium $z_1 \geq \frac{1}{2\alpha\psi_z}$, therefore the condition holds.

Next, consider a trailing incumbent. Sufficient condition to ensure that $x_1^* < x_I$ is that the probability of winning is decreasing in x_1, given the choice on the Z dimension. Plugging in the probability of winning, this requires

$$2\alpha\psi_x\Big[\pi_x - 2\alpha\psi_z z^\dagger\pi_z\Big] < 0, \tag{B-21}$$

where $z^\dagger = \min\{z_1, \frac{1}{2\alpha\psi}\}$. Notice that this is always satisfied when $z_1 \geq \frac{1}{2\alpha\psi_z}$, since for a trailing incumbent $\pi_x < \pi_z$. For an arbitrarily low δ and a sufficiently large z_I, we must have that in equilibrium $z_1 \geq \frac{1}{2\alpha\psi_z}$, therefore the condition holds. □

Proof of Implication 1. Let $x_I = z_I$, $\psi_x = \psi_z$ and $\lambda_I^x = \lambda_I^z$.

First, suppose the incumbent is leading. To prove that $x_1^* > z_1^*$, we will proceed by contradiction. Suppose that $x_1^* < z_1^*$. Notice that, assuming $x_I = z_I$, $\psi_x = \psi_z$ and $\lambda_I^x = \lambda_I^z$, a necessary condition for this to hold in equilibrium is that $\frac{\partial \mathbb{P}}{\partial x} < \frac{\partial \mathbb{P}}{\partial z}$. For a leading incumbent, we have that

$$\mathbb{P} = 2\alpha\psi\pi_x x + (1 - 2\alpha\psi x)\Big(1 - 2\alpha\psi(1 - \pi_z)z\Big). \tag{B-22}$$

Therefore, $\frac{\partial \mathbb{P}}{\partial x} = 2\alpha\psi\Big(\pi_x - [1 - 2\alpha\psi(1 - \pi_z)z]\Big)$ and $\frac{\partial \mathbb{P}}{\partial z} = -2\alpha\psi(1 - \pi_z)\Big(1 - 2\alpha\psi x\Big)$.

Thus, the condition that $\frac{\partial \mathbb{P}}{\partial x} < \frac{\partial \mathbb{P}}{\partial z}$ reduces to

$$\pi_x - \pi_z < 2\alpha\psi(1 - \pi_z)(x - z). \tag{B-23}$$

Recall that, for a leading incumbent, $\pi_x > \pi_z$. Thus, if $x < z$, then the condition cannot be satisfied, and $\frac{\partial \mathbb{P}}{\partial x} > \frac{\partial \mathbb{P}}{\partial z}$, a contradiction.

Next, suppose the incumbent is trailing. Proceeding by contradiction, suppose that $x_1^* = p_h$ and $z_1^* = p_l$, where $p_l < p_h$. We will show that a deviation to $x = p_l$ and $z = p_h$ is strictly profitable. Notice that, assuming $x_I = z_I$, $\psi_x = \psi_z$ and $\lambda_I^x = \lambda_I^z$, the deviation is profitable iff $\mathbb{P}(x_1^* = p_h, z_1^* = p_l) < \mathbb{P}(x = p_l, z = p_h)$. This is because the static utility from the two pairs of policies is the same. Assuming interior policies (a similar argument proves the result for policies at a corner), for a trailing incumbent, we have that

$$\mathbb{P} = 2\alpha\psi\pi_x x + (1 - 2\alpha\psi x)2\alpha\psi\pi_z z. \tag{B-24}$$

Therefore, $\mathbb{P}(x = p_l, z = p_h) > \mathbb{P}(x_1^* = p_h, z_1^* = p_l)$ reduces to

$$2\alpha\psi p_h\pi_x + (1 - 2\alpha\psi p_h)2\alpha\psi p_l\pi_z < 2\alpha\psi p_l\pi_x + (1 - 2\alpha\psi p_l)2\alpha\psi p_h\pi_z \tag{B-25}$$

Rearranging:

$$2\alpha\psi(p_h - p_l)\pi_x < 2\alpha\psi\pi_z[p_h(1 - 2\alpha\psi p_l) - p_l(1 - 2\alpha\psi p_h)] \tag{B-26}$$

which simplifies to:

$$2\alpha\psi(p_h - p_l)\pi_x < 2\alpha\psi\pi_z(p_h - p_l) \tag{B-27}$$

Recall that for a trailing incumbent $\pi_x < \pi_z$. Thus, the condition is always satisfied. Therefore, in equilibrium it must be the case that $x_1 < z_1$. □