Policy Gambles and Valence in Elections*

Peter Bils† Federica Izzo‡

Abstract

We study strategic policy experimentation by an incumbent politician when voters care about both policy and the candidates' valence. In our model, the voter does not know the location of her ideal policy and learns via experience, in turn, the officeholder uses policymaking to control voter learning. The incumbent thus faces a trade-off between implementing a policy close to his own ideal point, or one that induces the optimal amount of voter learning to win reelection. In equilibrium, how the incumbent solves the trade-off depends on his expected valence. We find that a trailing incumbent sometimes implements a safer policy than he would absent electoral incentives, despite needing to generate new information to win the election. Furthermore, increasing the incumbent's expected valence (and thus electoral advantage) can motivate him to gamble more in equilibrium. However, this relationship between electoral security and experimentation depends crucially on the amount of uncertainty on the valence dimension.

*We thank an Advisory Editor, three anonymous referees, Brenton Kenkel, Greg Martin, Mattias Polborn, Greg Sasso, and numerous seminar audiences for helpful comments and suggestions. This paper was previously circulated with the title 'Policymaking in Times of Crisis'.

†Dept. of Economics and Dept. of Political Science, Vanderbilt University.

‡Dept. of Economics and Dept. of Political Science, Vanderbilt University.

Introduction

Voters and politicians navigate a world filled with uncertainty, where the outcomes of different policies are often unpredictable. In response, voters use their past experiences to guide their decisions at the ballot box, drawing inferences and adjusting their preferences based on the outcomes of policymaking by different parties (Fiorina, 1981). In this context, officeholders can use policy to shape voters' experiences and learning. Bolder policy experiments have more uncertain outcomes, and therefore facilitate voter learning by generating more information. Safer policies, on the other hand, tend to produce more predictable outcomes that are less informative to voters. Depending on their electoral prospects, policymakers have strategic incentives to either pursue risky choices or avoid policy gambles.

Although policy choices significantly impact electoral outcomes, they are not the only factor that influence an officeholder's chances of winning reelection. The literature, dating back to Stokes (1963), highlights that certain 'valence' characteristics—such as honesty, charisma, experience, and competence—make candidates more appealing to voters, regardless of their policy positions or the voters' ideological leanings.

In this paper, we build on these observations and ask: How does an officeholder's valence influence his incentives to experiment with risky policies? To address this question, we develop a model of electoral accountability with valence-heterogeneous candidates in which there is uncertainty about the voter's optimal policy. In this setting, the incumbent's policy choice influences the amount of voter learning in equilibrium.

Intuition from existing theories suggests that improving the incumbent's electoral standing should induce him to implement safer policies that limit voter learning (e.g., Rose-Ackerman, 1980; Downs and Rocke, 1994), as doing so helps him preserve his strengthened electoral position. Likewise, electorally disadvantaged incumbents should always pursue riskier policies than their static preferences would dictate. In contrast with these intuitions, we find that increasing the incumbent's expected valence can encourage more policy experimentation. Furthermore, an electorally disadvantaged incumbent sometimes implements a safer policy than he would absent electoral incentives, despite needing to generate information to win reelection. However, the equilibrium relationship between expected valence and experimentation depends crucially on how much uncertainty the incumbent faces about his valence advantage over the challenger on election day.

Related Literature

Our model complements the political economy literature on policy experimentation (Biglaiser and Mezzetti, 1997; Majumdar and Mukand, 2004; Fu and Li, 2014; Dewan and Hortala-Vallve, 2019;

Cai and Treisman, 2009; Cheng and Li, 2019).[1] We contribute to this body of work by considering how an orthogonal valence dimension influences incentives for electorally accountable politicians to engage in experimentation. Furthermore, while most of the papers in the literature study policy experimentation with a binary policy space, we consider a continuous space. This allows us to analyze not only *when* a policymaker has incentives to experiment with risky policies, but also how the *intensity* of these dynamic incentives influences the *amount* of policy experimentation emerging in equilibrium. Importantly, a binary policy space would conceal the non-monotonic effect of valence we uncover in our framework, as well as the mediating effect of uncertainty.

The learning technology we use relates to the models introduced in Ashworth, Bueno de Mesquita and Friedenberg (2017) and Izzo (2023). However, there are technical differences between these works and our approach. In Ashworth, Bueno de Mesquita and Friedenberg (2017), the incumbent's action is unobserved, while in our model the voter updates her beliefs (and thus ideological preferences) based on the implemented policy as well as the outcome of the experiment. In Izzo (2023), voter learning is stark, with each realized outcome being either fully informative or completely uninformative, whereas learning in our setting is smooth, and outcomes are always somewhat informative but never fully so. This is an important change, as stark learning also obfuscates much of the impact of valence on policymaking, similar to assuming a binary policy space.

Callander (2011) also studies experimentation in a continuous policy space, albeit under different technical and substantive assumptions over the nature of policy uncertainty that generate different learning dynamics. More importantly, the focus of Callander's work fundamentally differs from our own. Focusing on the *statically* optimal choice for a policy maker, Callander (2011) assumes myopic parties. In contrast, our theory focuses on the dynamic incentives of politicians to control information due to electoral concerns.[2]

A prominent literature on elections has studied models in which voters care about fixed characteristics of the candidates, which are orthogonal to the policy dimension (e.g., Ansolabehere and Snyder, 2000; Groseclose, 2001; Bernhardt, Câmara and Squintani, 2011; Krasa and Polborn, 2012). Our contribution is to show that these valence characteristics can have surprising and nuanced effects on policy experimentation. Furthermore, we show that these effects can vary significantly with degree of uncertainty about the incumbent's relative valence.[3]

In this perspective, our paper complements Alonso and Câmara (2016), who also study an

[1]Others study how accountability influences an agent's incentives to exert (unobservable) effort improving the outcome of policy experiments (Hirsch, 2016; Yu, 2022). Another strand of this literature studies how decentralization affects policy experimentation, but abstract from electoral incentives (Strumpf, 2002; Volden, Ting and Carpenter, 2008; Callander and Harstad, 2015).

[2]Callander and Hummel (2014) considers forward-looking parties, but assumes exogenous retention probabilities.

[3]Other works study how the amount of information that is revealed about valence can impact the selection of candidates (Carrillo and Mariotti, 2001; Boleslavsky and Cotton, 2015) and the divergence of party platforms (Carrillo and Castanheira, 2008).

incumbent's incentives to control policy-relevant information in a setting with valence-heterogeneous candidates. In their model, the incumbent's objective is to maximize his probability of retention and the experiment has no direct effect on his payoff. In our model, the incumbent cares about both policy and office. Therefore, he directly incurs costs from using policy to control information, rather than following his ideological preferences. As such, Alonso and Câmara's setup is best interpreted as one where the experiment is a "small-scale policy trial", while our paper studies "full-scale policy experimentation" (Alonso and Câmara, 2016, p. 393).[4] This is a crucial feature of our model, since now the incumbent must account for the *intensity* of his incentives to control information, and not just the direction. Furthermore, Alonso and Câmara (2016) assume the incumbent's relative valence is always revealed before the election. Instead, we parametrize the probability of the incumbent's type being revealed, and show that this variable, i.e., the amount of uncertainty the incumbent faces about what his net valence will be at election day, is a crucial mediator of the relationship between valence and experimentation.

In Alonso and Câmara (2016), the usual "gambling for resurrection" intuition holds when the challenger's valence is drawn from a log-concave distribution: policy experimentation decreases with the incumbent's valence. However, this result is reversed when the distribution is log-convex. This reversal arises because switching to a log-convex distribution changes the curvature of the incumbent's winning probability, causing a high-valence incumbent to electorally benefit from a risky gamble (and vice versa for a low-valence incumbent). In our model, instead, the gambling for resurrection intuition fails to hold due to the interaction of the incumbent's static policy preferences with her dynamic electoral incentives, and can emerge even absent uncertainty over valence.

Model

Players and Actions: We consider a two-period model of electoral accountability. There is an incumbent (I), a challenger (C), and a representative voter (V). In the first period, the incumbent chooses a policy $x_1 \in \mathbb{R}$. At the end of the first period, the voter observes the policy choice, a noisy signal of her policy utility (described below) and, possibly, the candidates' relative valence. Next, the voter chooses whether to reelect the incumbent or replace him with the challenger. Finally, the winner of the election chooses policy $x_2 \in \mathbb{R}$ and the game ends.

[4]Motivated by our focus on experimentation via policymaking, we also adopt a different learning technology. Specifically, the incumbent can only manipulate the location of the implemented policy and, thus, the informativeness of outcomes. In contrast, in the Bayesian persuasion framework used by Alonso and Câmara (2016) the incumbent can choose any signal structure mapping the state to realized outcomes.

Payoffs and Information: Politicians are motivated by both ideology and winning office. Specifically, given implemented policy x_t, the period-t utility of politician $i \in \{I, C\}$ is

$$-(x_t - \hat{x}_i)^2 + \mathbb{I}_t^i \beta,$$

where $\hat{x}_i \in \mathbb{R}$ is i's ideal point, $\beta \geq 0$ represents the value of holding office, and $\mathbb{I}_t^i = 1$ if i is in office at time t and $\mathbb{I}_t^i = 0$ otherwise. For simplicity, we assume $0 < \hat{x}_I = -\hat{x}_C \equiv \hat{x}$, thus, the incumbent is right-wing and the challenger is symmetrically left-wing. Furthermore, the candidates' ideal points are common knowledge.

As for the voter, she cares about both the policy dimension and the candidates' relative valences. Specifically, in each period t her utility is given by:

$$-(x_t - \omega)^2 + \mathbb{I}_t^I \theta, \tag{1}$$

where ω is the voter's ideal point and θ is the net valence advantage (or disadvantage) of the incumbent over the challenger. Thus, consistent with previous models, we treat valence as a dimension that is orthogonal to policy.

We assume that the location of ω and the net valence θ are initially unknown to all players. However, the voter may learn about both ω and θ before the election.

First, ω can take one of two values, $\omega \in \{-1, 1\}$, normalized for simplicity. Players have a common prior belief that $Pr(\omega = 1) = \gamma \in (0, 1)$. Before making her retention decision, the voter observes a noisy signal of her policy utility that is derived from the implemented policy:

$$s_\omega = -(x_1 - \omega)^2 + \varepsilon, \tag{2}$$

where ε is a shock drawn from the standard normal distribution, with CDF Φ and PDF ϕ. This shock can capture idiosyncratic factors influencing voter's payoffs, or frictions in her learning process. The assumption that ω is unknown to all players can be interpreted as the voter learning about her own preferences by experience. Alternatively, uncertainty may stem from the underlying state of the world, which determines the voter's optimal policy.[5] In contrast, because politicians represent constituencies with more clearly established policy interests, their ideal points are state-independent.

Second, for the valence term, we assume θ is ex-ante unknown to all, $\theta \in \{-1, 1\}$. Thus, θ captures the net valence advantage ($\theta = 1$) or disadvantage ($\theta = -1$) of the incumbent relative to the challenger. Valence is drawn independently of ω, and players share common prior that $Pr(\theta = 1) = \pi \in (0, 1)$. To parametrize the amount of uncertainty the incumbent faces about his valence

[5] Here, the state ω may indicate the optimal direction of reform for the voter. In this interpretation, the model then normalizes the status quo to 0.

(dis)advantage at the time of policymaking, we assume that, after the incumbent chooses policy, θ is revealed with probability $\rho \in [0, 1]$. Thus, the voter observes a signal $s_\theta \in \{\emptyset, -1, 1\}$, where $Pr(s_\theta = 1|\theta = 1) = Pr(s_\theta = -1|\theta = -1) = \rho$ and $Pr(s_\theta = \emptyset|\theta = -1) = Pr(s_\theta = \emptyset|\theta = 1) = 1 - \rho$.[6]

Our notion of valence is broad, encompassing electorally relevant factors outside the candidate's control. If valence reflects unobserved attributes, such as competence, the probability ρ may capture the possibility that an exogenous shock, e.g., a crisis or scandal, will reveal θ and reshape voters' perceptions. Alternatively, it can represent uncertainty about how, on election day, voters will perceive the two candidates in terms of charisma, likeability, integrity, etc. However, the case of $\rho = 0$, where the incumbent's relative valence remains fixed throughout the game, is a a special case of particular substantive interest. This case connects our model to a large literature on elections which emphasizes the relevance of candidates' *observed* characteristics such as age, gender, or ethnicity. This interpretation also aligns with several models featuring valence-heterogeneous candidates with a known advantage (see, e.g., Ansolabehere and Snyder, 2000; Groseclose, 2001; Bernhardt, Câmara and Squintani, 2011).

Finally, we impose the following assumption on office benefit:

Assumption 1. *Office benefit is sufficiently large:* $\beta > \frac{\hat{x}^2(4\gamma-3)}{1-\gamma}$.

This assumption implies that the policy choice of an incumbent who is only slightly ahead of the challenger is driven mostly by reelection concerns. This does not alter our qualitative results, but simplifies the analysis and statement of the propositions.

Timing: To sum up, the game proceeds as follows:

1. Nature draws $\omega \in \{-1, 1\}$ and $\theta \in \{-1, 1\}$.
2. I chooses policy $x_1 \in \mathbb{R}$.
3. V observes the policy choice x_1, signals s_ω and s_θ, and updates her beliefs about ω and θ.
4. The voter makes her reelection decision.
5. The second period officeholder chooses policy $x_2 \in \mathbb{R}$.
6. Utilities are realized and the game ends.

[6] In Appendix B, we consider versions of the model where the incumbent has private information about θ. For simplicity, we consider separately the two limiting cases of $\rho = 0$ and $\rho = 1$. We show that, across the two models, information asymmetries do not change our qualitative findings from the baseline setup.

Preliminaries

Our solution concept is Perfect Bayesian Equilibrium, and we restrict attention to a selection of equilibria that satisfy a differentiability requirement (see Appendix for details), which we henceforth refer to as "equilibrium". Before moving to the analysis of equilibrium policymaking, it is useful to establish some preliminary results about voter behavior.

The Voter's Problem. In the last period, the officeholder has no reelection concerns and thus always implements his ideal policy. Therefore, if I is re-elected then $x_2^* = \hat{x}$, otherwise, $x_2^* = -\hat{x}$. Given that the remainder of our analysis focuses on first-period policymaking, moving forward we drop the time subscripts on policy choices.

The voter wants to elect the candidate whose ideal point provides her with the highest expected utility, but she also prefers the candidate with the greatest valence. Thus, in equilibrium, her decision depends on her posterior beliefs over both the policy and valence dimensions.

Lemma 1. *The voter reelects the incumbent if:*

$$\mathbb{E}[\omega|x, s_\omega] > -\frac{\mathbb{E}[\theta|s_\theta]}{4\hat{x}}. \tag{3}$$

If $\mathbb{E}[\omega|x, s_\omega] < -\frac{\mathbb{E}[\theta|s_\theta]}{4\hat{x}}$*, then the voter elects the challenger. Otherwise, if the voter is indifferent, then each candidate wins the election with probability* $1/2$.[7]

Lemma 1 shows that, intuitively, increasing the incumbent's expected valence relative to the challenger makes the voter more lenient on the policy dimension. Notice further that the effect of valence on the voter's reelection decision is stronger when the candidates are less polarized.

Building on this result, we can identify the conditions under which the voter would ex ante choose to reelect the incumbent. This depends on the initial expectations of valence, $\mathbb{E}[\theta]$, and the optimal policy, $\mathbb{E}[\omega]$. To simplify notation moving forward, we define *I's ex ante expected valence* as:

$$v \equiv \mathbb{E}[\theta].$$

From equation 3, the ex ante indifference threshold for the voter can be then expressed as:

$$\overline{v} = -4\hat{x}\mathbb{E}[\omega].$$

Leveraging this notation, we introduce the following definition:

[7] The indifference-breaking assumption makes the incumbent's problem continuous at $x = 0$ for all v. It is only consequential for a measure 0 set of parameters and does not otherwise affect our results.

Definition 1. *The incumbent has an electoral advantage if* $v > \overline{v}$*, otherwise the incumbent has an electoral disadvantage.*

In our framework, the *composition* of the incumbent's overall electoral advantage (or disadvantage) is also important for understanding his strategic choices in policymaking. In particular, the sign of v captures whether the valence dimension initially favors the incumbent or challenger, while the sign of $\mathbb{E}[\omega]$ determines whether the voter initially prefers the ideal point of the incumbent or the challenger. For ease of exposition, we define two more terms:

Definition 2. *The incumbent has*

- *a valence advantage if* $v > 0$*, and a valence disadvantage if* $v < 0$*;*
- *a policy advantage if* $\mathbb{E}[\omega] > 0$*, and policy disadvantage if* $\mathbb{E}[\omega] < 0$*.*

Note that $v = 2\pi - 1$ and $\mathbb{E}[\omega] = 2\gamma - 1$. Thus, the incumbent may have a policy advantage but a valence disadvantage, or vice versa. In such cases, the incumbent may have an electoral advantage or disadvantage, depending on the exact value of the voter's initial beliefs.

To streamline the analysis, we impose the following assumption on the level of ideological polarization between incumbent and challenger:

Assumption 2. $\frac{1}{4} < \hat{x} < \frac{1}{4|\mathbb{E}[\omega]|}$.

The assumption that $\frac{1}{4} < \hat{x}$ ensures that $-1 < -\frac{\mathbb{E}[\theta|s_\theta]}{4\hat{x}} < 1$ for any valence signal s_θ. As such, even if the voter learns the candidates' relative valence, there always exist policy signals, s_ω, that are decisive for her electoral decision. The second inequality, $\hat{x} < \frac{1}{4|\mathbb{E}[\omega]|}$, guarantees $\overline{v} \in (-1, 1)$. Thus, the incumbent may be ex-ante electorally advantaged or disadvantaged depending on his initial expected valence.

Policy Experimentation and Voter Learning. Next, we characterize the features of the voter's learning on the policy dimension, which is endogenous to the amount of experimentation. Using Bayes' rule, if the incumbent chooses policy x and this generates signal s_ω, then the voter's posterior belief that $\omega = 1$, is given by:

$$Pr\big(\omega = 1|x, s_\omega\big) = \frac{\gamma\phi\big(s_\omega + (x-1)^2\big)}{\gamma\phi\big(s_\omega + (x-1)^2\big) + (1-\gamma)\phi\big(s_\omega + (x+1)^2\big)}. \tag{4}$$

Equation 4 highlights two crucial properties of the learning process. First, because the shock distribution satisfies the Monotone Likelihood Ratio Property, $Pr(\omega = 1|x_1, s_\omega)$ is increasing in s_ω when $x > 0$, and decreasing in s_ω otherwise. Thus, when a right-wing (left-wing) policy generates a higher signal s_ω, the voter believes it is more likely that her ideal policy is also right-wing (left-wing).

Second, even fixing the signal s_ω, the inference that the voter draws depends on the implemented policy. As x becomes more extreme, the signal distributions conditional on the state ω move farther apart. As a consequence, the voter is better able to filter out information from noise and draws a more precise inference. Given these properties, more extreme policies reduce the variance in the posterior distribution. In the Appendix, we formalize this discussion and show that outcomes are more (Blackwell) informative as $|x|$ increases.[8]

Policy Gambles and Valence in Elections

We now characterize equilibrium policymaking. Absent reelection incentives, the incumbent always implements his ideal point (as in the second period). However, in our model, any policy (other than 0) is informative about ω, and more extreme policies are more informative for the voter. Thus, if we observe I implementing a policy more extreme than $\hat{x}$, then it must be because his electoral incentives compel him to generate a more informative signal than he would in the absence of career concerns. Conversely, if the incumbent moderates his policy away from $\hat{x}$ and towards 0, then it is due to electoral incentives to prevent information generation relative to $\hat{x}$.

Definition 3. *The incumbent gambles if $x > \hat{x}$.*

The Fixed-Valence Case

We begin by studying the case where $\rho = 0$, so that the voter never updates about θ and the extent of the incumbent's valence (dis)advantage remains fixed throughout the game. As noted earlier, this case holds particular substantive significance and captures how valence has often been modeled in previous work. Furthermore, the core incentives driving our results are clearest in this scenario, making it especially useful for illustrating the underlying intuition.

The equilibrium policy choice depends on I's ideological preferences and his reelection incentives. Thus, the first step is to characterize the effect of the policy on I's probability of winning.

Lemma 2. *Assume $\rho = 0$. Let $x > 0$.*

1. *Suppose I has an electoral advantage.*
   - *(a) If I has a valence disadvantage and a policy advantage, then the probability of winning is single-dipped in x;*
   - *(b) Otherwise, the probability of winning is strictly decreasing in x.*

[8] Note, under the interpretation where ω captures the optimal direction of reform, policies that are more extreme *relative to the status quo* induce more informative signals.

2. *Suppose I has an electoral disadvantage.*

    (a) *If I has a valence advantage and a policy disadvantage, then the probability of winning is single-peaked in x;*

    (b) *Otherwise, the probability of winning is strictly increasing in x.*

Symmetric results hold for $x < 0$.[9]

Whether increasing the informativeness of the policy signal improves the incumbent's probability of winning depends on two factors. The first is whether the incumbent would lose or win absent new information — i.e., does the incumbent have an electoral advantage ($v > \overline{v}$) or disadvantage ($v < \overline{v}$) — which we refer to as the *extensive margin* effect. This effect underpins the conventional intuition that an electorally disadvantaged incumbent has incentives to gamble, whereas an advantaged one seeks to limit new information. The extensive margin is stronger when $|v - \overline{v}|$ increases, as a highly disadvantaged (advantaged) incumbent benefits the most (least) from generating information. The second factor is whether information is more likely to favor the incumbent — i.e., whether γ is greater or less than $\frac{1}{2}$ — which we call the *intensive margin* effect. This effect depends on the distance of γ from $\frac{1}{2}$, which captures how much uncertainty there is ex ante about the outcome of the policy signal. Since $\overline{v} = -4\hat{x}(2\gamma - 1)$, the intensive margin is stronger when $|0 - \overline{v}|$ increases.

If the incumbent's electoral and policy advantages have opposite signs, then the extensive and intensive margin effects work in the same direction. As such, the probability of winning is monotonic in x. To see why, suppose the incumbent has an electoral advantage ($\overline{v} < v$) and policy disadvantage ($\mathbb{E}[\omega] < 0$ which implies $\gamma < \frac{1}{2}$). Then, new information can only hurt I's initial electoral prospects *and* more accurate information is likely to be unfavorable, since $\gamma < \frac{1}{2}$. Consequently, the probability of winning is decreasing in x. A symmetric logic implies that the probability of winning is increasing in x for an incumbent with an electoral disadvantage and policy advantage.

Now, suppose the electoral and policy advantages have the same sign. In this case, the extensive and intensive margin effects push in opposite directions, potentially making the probability of winning non-monotonic. For example, if the incumbent enjoys both an electoral and policy advantage, then he is guaranteed re-election in the absence of new information, however, new information is also likely to be favorable. We find that whether these competing effects generate a non-monotonicity depends on the mismatch between the incumbent's policy and valence advantages. If the incumbent's policy and valence advantages have the same sign, then v is relatively farther from $\overline{v}$ than 0 is from $\overline{v}$, $|v - \overline{v}| > |0 - \overline{v}|$. Thus, the extensive margin dominates the intensive one, and the probability of winning is again monotonic in x. In contrast, if the policy and valence advantages

[9] If v is exactly equal to $\overline{v}$, then I's probability of winning is increasing when I has a policy advantage, and instead decreasing when he has a policy disadvantage.

Figure 1: Probability of winning as a function of policy choice

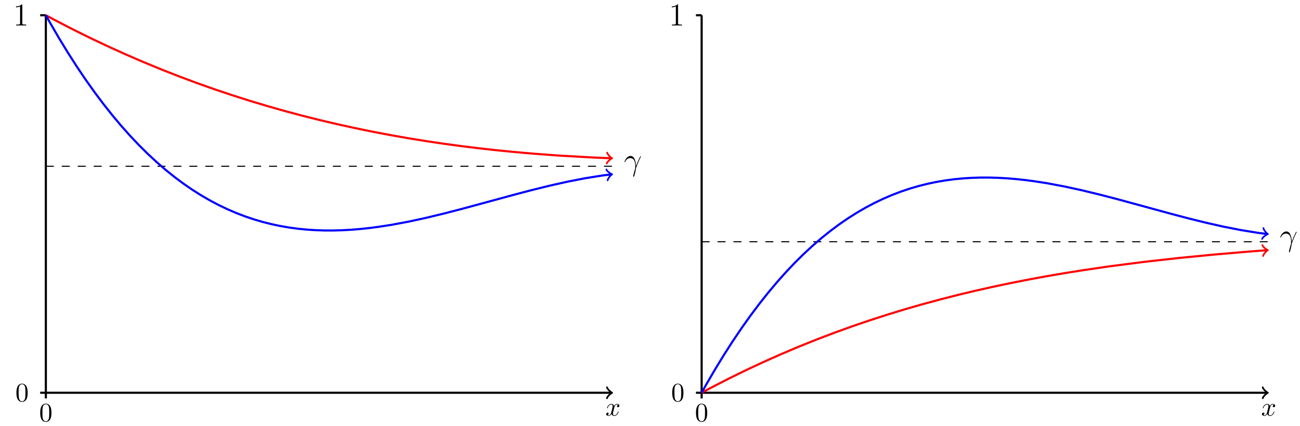

(a) electorally advantaged incumbent: with valence advantage (red), with valence disadvantage (blue)

(b) electorally disadvantaged incumbent: with valence disadvantage (red), with valence advantage (blue)

have opposite signs, so that $|v - \overline{v}| < |0 - \overline{v}|$, then the intensive margin can come to dominate, generating a non-monotonicity in the probability of winning.

Therefore, there are two cases in which I's probability of winning is non-monotonic. First, consider an incumbent with an electoral and policy advantage, but a valence disadvantage. When $\gamma > \frac{1}{2}$, the intensive margin effect implies that increasing the informativeness of the policy reduces the likelihood of a 'false negative', where the policy yields a low voter payoff despite aligning with the state ω. When x is very close to 0 this effect is negligible, as the incumbent is effectively guaranteed re-election. Consequently, the probability of winning decreases in x. However, as x increases, the intensive margin strengthens and eventually dominates because $\overline{v} < v < 0$, and so $|v - \overline{v}| < |0 - \overline{v}|$. Thus, the probability of winning is single-dipped in x, first decreasing and then increasing.

Finally, consider an incumbent who faces both an electoral and policy disadvantage but has a valence advantage. Here, dampening information may be beneficial, as it increases the likelihood of a 'false positive', where the policy yields a high payoff despite not aligning with the state ω. The condition $|v - \overline{v}| < |0 - \overline{v}|$ ensures that this effect is strong enough to outweigh the extensive margin when x is sufficiently large. However, for x close to 0 the extensive margin still dominates, as generating some information is necessary to win. Thus, the probability of winning is single-peaked, maximized at an intermediate policy that amplifies the likelihood of a false positive. Figure 1 depicts how the incumbent's probability of winning as a function of x depends on I's policy and valence advantages (or disadvantages).

If I implemented policy solely based on reelection incentives to control information, then Lemma 2 would fully characterize how the optimal policy depends on v. However, in equilibrium, the policy choice is also influenced by I's policy preferences. Lemma 3 then specifies the conditions under which the incumbent chooses to gamble, implementing a policy more extreme than his ideal point.

Lemma 3. *Assume $\rho = 0$. There exists a $\hat{v}_G$ such that: if $v < \hat{v}_G$ then the incumbent gambles in equilibrium, $x^* > \hat{x}$, and if $v > \hat{v}_G$ then the incumbent does not gamble, $x^* < x$.*[10] *If the incumbent has a policy disadvantage, then $\hat{v}_G \in (0, \overline{v})$. Otherwise, $\hat{v}_G = \overline{v}$.*

When the incumbent has an electoral advantage, his probability of winning may be non-monotonic but is always maximized if no new information is revealed. Thus, limiting information is relatively more beneficial than gambling, and he always chooses a policy more moderate than $\hat{x}$.

Consider instead an electorally disadvantaged incumbent. When this incumbent does not have a disadvantage on the policy dimension, Lemma 2 indicates that his probability of winning is strictly increasing in x. Therefore, this incumbent always gambles and $\hat{v}_G = \overline{v}$. Suppose instead the incumbent has a policy disadvantage. If he also has a disadvantage over the valence dimension, then his probability of winning is always increasing in x and $x^* > \hat{x}$ (implying that $\hat{v}_G$ is strictly larger than 0). When instead this incumbent has a valence advantage, so $v \in (0, \overline{v})$, his probability is maximized at an intermediate value of x. In particular, for v sufficiently large in this range, $x < \hat{x}$ maximizes the incumbent's ability to exploit false negatives, and such an incumbent is electorally motivated to distort policy to a more moderate position, despite needing to generate information. Consequently, this incumbent's policymaking resembles that of an electorally advantaged incumbent, and thus $\hat{v}_G \in (0, \overline{v})$.

Having characterized how v influences the direction in which the incumbent distorts policy away from $\hat{x}$, we complete the analysis by studying how the equilibrium policy x^* changes in v.

Proposition 1. *Assume $\rho = 0$. There exists a cut-point $\underline{v} \geq -1$, with $\underline{v} < \overline{v}$, such that, if $v \in (\underline{v}, \overline{v})$ then x^* is decreasing in v. Otherwise, x^* is increasing in v.*

Changing v has two effects on the incumbent's incentives. First, increasing v makes the voter more willing to reelect the incumbent, which implies I can survive more negative policy information when advantaged, and needs less positive policy information when disadvantaged. Second, moving v towards $\overline{v}$ creates an ex ante more competitive electoral environment, which makes new policy information more relevant for the voter's decision. Importantly, these two effects imply that increasing v can induce the incumbent to implement either a more moderate or more extreme policy.

If the incumbent has an electoral advantage, $v > \overline{v}$, these two effects go in the same direction. Increasing v away from $\overline{v}$ increases how much negative policy information the incumbent can reveal while still winning reelection *and* makes policy information less relevant. Both effects weaken I's incentives to control information, moving x^* to the extreme towards the incumbent's ideal point.

[10] At $v = \hat{v}_G$ the incumbent may gamble or not depending on the parameters.

In contrast, if the incumbent is electorally disadvantaged, then these two effects compete. First, suppose that $\hat{v}_G = \overline{v}$, so that the electorally disadvantaged incumbent always gambles in equilibrium. Increasing v towards $\overline{v}$ increases the salience of the policy dimension (making information control more relevant), but also makes the voter more lenient towards I (making information control less relevant). When v is nearer to -1, the former effect can dominate, intensifying the incumbent's incentives to experiment. Conversely, when v approaches $\overline{v}$, the latter effect prevails, dampening incentives to control information. Therefore, as v increases away from -1, initially x^* can move more extreme as the incumbent is incentivized to gambles more. However, as v gets closer to $\overline{v}$, the incentives flip and x^* begins to moderate back towards $\hat{x}$.

Second, suppose that $\hat{v}_G < \overline{v}$, so that an electorally disadvantaged incumbent gambles for low values of v and does not gamble for values of v close to $\overline{v}$. In the gambling region, the incentives discussed above continue to hold, and the effect of v on x^* is non-monotonic, with extremism increasing then decreasing. As v moves into the no-gambling region, further increasing v towards $\overline{v}$ strengthens incentives to control information by exploiting false positives, as previously discussed. In turn, this induces the incumbent to further moderate his policy, to maximize the likelihood of generating a false positive outcome. Thus, regardless of whether $\hat{v}_g$ is equal or strictly below $\overline{v}$, a disadvantaged incumbent's policy choice may be non-monotonic in v, first increasing and then decreasing. Proposition 2 provides conditions under which this non-monotonicity must emerge.

Proposition 2. *If $\hat{x}$ is sufficiently small then $\underline{v} > -1$.*

Recall that $\hat{x}$ captures the degree of polarization between the candidates. When $\hat{x}$ is low the candidates deliver similar policies in the second period, and thus electoral outcomes are mostly determined by valence. In this case, when v is close to -1 the electorally disadvantaged incumbent's probability of winning is very low. Thus, such an incumbent has no incentive to distort policy very far from his ideal point, which implies that policy must initially be increasing in v, and $\underline{v} > -1$.[11]

Figure 2 pulls together Lemma 3 and Proposition 1 to depict how the equilibrium policy changes as a function of v when $\rho = 0$. Notice that the least amount of experimentation is by incumbents with a very small electoral advantage.[12] Furthermore, if $\underline{v} > -1$ (as in Figure 2), then the maximum amount of policy experimentation is by incumbents who face only a moderate disadvantage. These results highlight the importance of considering both ideological preferences and reelection motives for understanding incentives to engage in policy experimentation.

In concluding this section, we note that in our model presented above the probability of winning is not jointly continuous in x and v at $(x, v) = (0, \overline{v})$, which generates the discontinuity in the equilibrium policy. However, introducing an idiosyncratic shock to the voter's payoff to smooth the

[11] It is easy to verify numerically that $\underline{v} = -1$ can emerge for some parameters when $\hat{x}$ is larger.

[12] Indeed, such incumbents engage in almost no experimentation. In the Appendix, under our assumption on β we show that $\lim_{v \to \overline{v}^+} x^* = 0$.

Figure 2: Equilibrium Policymaking and Expected Valence

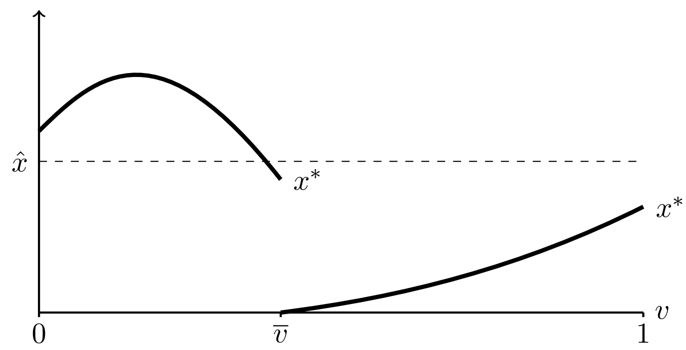

Note: Figure 2 depicts the incumbent's equilibrium policy choice, x^*, as a function of his valence when $\rho = 0$.

probability of winning would not change the effect of v on policy experimentation. Even with an idiosyncratic shock, an incumbent with very low ex-ante valence has poor electoral prospects, as the likelihood of the shock pushing him above the retention threshold remains low. This weakens his incentive to take risks, keeping policy close to the static optimum $\hat{x}$. As v increases, two opposing forces emerge, similar to what is described above. Improved electoral prospects encourage gambling on resurrection, while the reduced need for positive policy signals discourages extreme policies. This tension drives the non-monotonicity seen in our model for an electorally disadvantaged incumbent. When v instead is high, the incumbent is likely to win even without new policy information, as the probability the idiosyncratic shock causes electoral defeat is low. This pushes policy toward moderation. However, as v increases further, the stronger valence advantage shields the incumbent from electoral risk, weakening incentives to control information and shifting policy back toward $\hat{x}$.

Policymaking with Uncertainty about Valence

Finally, we move to the case where $\rho > 0$, so there is a possibility that the incumbent's type is revealed before the election. Thus, an incumbent who has an ex ante valence advantage may lose it, or vice versa. The incumbent, then, takes this possibility into account when balancing the trade-off between ideological preferences and the need to control voter learning to secure reelection.

Here, we show that the amount of uncertainty on the valence dimension, ρ, is a key mediator in the relationship between the incumbent's ex ante valence and the amount of policy experimentation. Recall that, in our framework, $Pr(\theta = 1) = \pi$ and $v = E[\theta] = 2\pi - 1$. This dual role of π, as probability and expectation, does not arise when $\rho = 0$, but it is central to the results of this section. Therefore, it is important to note that, when taking comparative statics on v we are effectively studying how π shapes equilibrium policymaking.

Proposition 3. *There exist $\underline{\rho}$ and $\overline{\rho}$, with $0 < \underline{\rho} \leq \overline{\rho} < 1$, such that if $\rho > \overline{\rho}$ then x^* is strictly decreasing in v. Instead, if $\rho < \underline{\rho}$ then x^* is non-monotonic in v.*

To understand this result, compare the case of $\rho = 1$ to the incentives emerging under $\rho = 0$. Differently from what we described in the previous section, when $\rho = 1$ the incumbent's ex-ante valence v does not directly enter the voter's retention decision. Increasing v increases the probability that I obtains (or consolidates) an ex-post electoral advantage. However, a higher v does not make the voter more lenient towards the incumbent on election day, since the voter's new information supersedes her prior beliefs.[13] This substantially alters how I balances his policy preferences against his incentive to control information compared to the fixed-valence model. A higher v only influences the incumbent's policy by increasing the probability he benefits from limiting the informativeness of the policy signal, and thus amplifying his incentives to protect an increasingly likely ex-post advantage. In turn, this induces him to implement a more moderate policy.

When ρ is interior, v not only determines the incumbent's ex-ante valence (as in the $\rho = 0$ case) but also captures the likelihood of the incumbent gaining (or solidifying) an ex-post advantage if the voter acquires new information about θ (as in the $\rho = 1$ case). Therefore, the incumbent's incentives are a combination of those emerging in the two limiting cases of $\rho = 0$ and $\rho = 1$. When ρ is high, the incumbent's type is likely to be revealed. As v increases the incumbent becomes more incentivized to protect an increasingly likely valence advantage. Thus, the equilibrium policy is decreasing in v. When ρ is low, the voter is unlikely to observe the incumbent's type. As v increases, the incumbent's ex-ante electoral standing improves. As illustrated in the previous section, this has an ambiguous impact on the incumbent's incentives, leading to a non-monotonic effect of v on x^*.

Proposition 4 provides conditions where the effect of v on x^* is always non-monotonic.

Proposition 4. *If $\hat{x} \to \frac{1}{4}$ then $\underline{\rho} \to 1$.*

If $\hat{x}$ is small then the candidates deliver similar policies in the second period. Thus, valence considerations are most prominent in the voter's calculus. Proposition 4 shows that it is precisely in this case that the conventional intuition about electoral security and policy experimentation always fails. When valence considerations are prominent for the voter, the incumbent's equilibrium policy choice under a fixed valence is very sensitive to changes in v and, as a consequence, the policy is always non-monotonic in v (as shown in Proposition 2). Thus, even in the general model, the non-monotonicity driven by these incentives continues to emerge for any value of ρ.

So far we have treated ρ as exogenous. Our final result explores the incumbent's incentives to endogenously reveal information about his valence, for example, by initiating a crisis that tests

[13] Necessary for our qualitative results is that the information the voter exogenously receives about valence is sufficiently more precise than her prior beliefs. Whether this exogenous signal is fully informative is less important.

his unknown competence (as in Downs and Rocke, 1994). Specifically, we allow the incumbent to generate a valence test, $\rho = 1$, or not, $\rho = 0$.

Proposition 5. *Assume office benefit is sufficiently large. If the incumbent has an electoral advantage then he chooses $\rho^* = 0$. If the incumbent has a policy disadvantage, a valence disadvantage, and $\mathbb{E}[\omega] < v$, then he chooses $\rho^* = 1$.*

If the incumbent is electorally advantaged, he has no reason to take the risk and always prefers $\rho = 0$. If instead the incumbent is electorally disadvantaged, then new information on either policy or valence is necessary for reelection. Suppose, in particular, that the incumbent is disadvantaged on both dimensions, valence and policy. Then, his probability of winning under $\rho = 0$ is strictly increasing in x, and thus takes a maximum value of γ. Instead, if the incumbent chooses to generate a valence test, $\rho = 1$, his probability of winning is at least π, since he is always guaranteed reelection if he implements policy $x = 0$ and the valence test succeeds. Thus, if if $\mathbb{E}[\omega] < v$ (i.e., $\gamma < \pi$), then only gambling on policy is always worse in expectation than a valence test, and hence $\rho^* = 1$.

Discussion and Conclusion

This paper examines how policymakers' valence influences their willingness to take risks in policymaking. Our findings show that these effects depend on the level of uncertainty over the valence dimension. Furthermore, our results challenge conventional wisdom, revealing that the effect of expected valence on policy experimentation may be non-monotonic, and disadvantaged incumbents may be unwilling to gamble despite needing to generate information to secure reelection.

Our results can thus help explain situations where, contrary to the conventional intuition, an incumbent engages in more rather than less policy experimentation after her valence — and likewise electoral security — increases. The privatization of many UK industries by Margaret Thatcher in the 1980s, following an increase in voters' perception of her valence, provides an example of such a case. The increase in Thatcher's electoral standing occurs after the Falklands War in 1982. In the early 1980s Thatcher's prospects for winning reelection are bleak. However, Thatcher achieves a swift and decisive victory in the Falklands that boosts her electoral standing (Dunleavy and Husbands, 1985), and she becomes seen as "tough and resourceful".[14] Thus, the war acts as a semi-exogenous shock that increases Thatcher's valence on a dimension that is orthogonal to privatization.[15]

Despite the victory improving her perceived valence, it is only after the war that Thatcher begins substantial privatization reform (Marsh, 1991).[16] Importantly, these policies were not im-

[14] `https://www.history.com/news/margaret-thatcher-falklands-war`

[15] The invasion of the Falkland Islands by Argentina took Thatcher by surprise (Biles, 2012).

[16] Note that this change in behavior is not due to end period effects. Many of the reforms start in 1983 during her second term, and Thatcher runs for reelection again in 1987.

plemented due to demand from the voters. Rather, consistent with the logic of our model, they were elite driven (McAllister and Studlar, 1989) and "inspired by considerations of political ideology" (Swann, 1988, p. 316). Furthermore, privatization significantly altered the status quo and its effects were far from certain. Indeed, these reforms are commonly regarded as a policy experiment by both contemporaneous (e.g. Young, 1986) and later accounts (e.g. Parker, 2004). John Moore himself, the Financial Secretary to the Treasury under Thatcher who spearheaded many of the reforms, referred to privatization as a "radical experiment." (Moore, 1992). Our theory suggests that Thatcher's increased valence made her more willing to take risks, and allowed her to move closer to her ideological preferences.

Before concluding, we discuss how our theory can inform future empirical research. Our findings highlight that analyzing the effect of valence on policy experimentation requires accounting for the incumbent's uncertainty about what voters' perception of his relative valence will be on election day. One possible strategy is to leverage exogenous crises, such as global recessions or the COVID-19 pandemic, which a recent political economy literature conceptualizes as valence-revealing events (e.g., Ashworth, Bueno de Mesquita and Friedenberg, 2018; Izzo, Forthcoming). Thus, we may treat policymaking during a crisis as happening in a period of high uncertainty—where the incumbent anticipates voters updating their beliefs before election day. Normal times are periods of low uncertainty—where voters are less likely to revise their prior beliefs. Our theory then predicts a fundamentally different effect of the incumbent's ex-ante valence on policy experimentation in periods of crisis and normal times.[17]

Since our theory points to the possible non-monotonic effect of valence, a flexible specification is essential. Under low uncertainty (normal times), a cubic polynomial provides a good approximation of the predicted relationship between valence and policy experimentation. Alternatively, researchers can use previous margins of victory or higher-level election results to distinguish between electorally advantaged incumbents—for whom the effect should be positive—and disadvantaged ones—for whom the effect is expected to be non-monotonic. In contrast, during crises, our theory predicts a negative effect for all incumbents. Causally identifying these effects is challenging. However, we hope our theory offers a useful framework for empirical research, highlighting the need to account for both the mediating role of uncertainty and potential non-monotonic effects to better understand the conditions that drive variation in policy experimentation.

[17] In Appendix B we compare our model under fixed ($\rho = 0$) and uncertain ($\rho = 1$) valence to derive further insights into how the risk of a crisis influences an incumbent's incentives to engage in policy experimentation.

A Proofs

Recalling that $v = 2\pi - 1$, and thus $\pi = \frac{1+v}{2}$, we write the incumbent's maximization problem as a function of v as follows:

$$\max_x -(x-\hat{x})^2 + (\beta + 4\hat{x}^2)\left(\rho\Big(\frac{1+v}{2}\cdot P_1(x) + \frac{1-v}{2}\cdot P_{-1}(x)\Big) + (1-\rho)P_v(x)\right) - 4\hat{x}^2,$$

where $P_{\mathbb{E}[\theta|s_\theta]}(x)$ denotes the incumbent's probability of winning the election conditional on the voter observing the signal s_θ and policy x. Before proceeding, two comments are in order.

First, we focus on a selection of equilibrium such that x^* is differentiable in v for $v < \overline{v}$ and $v > \overline{v}$. We note that the equilibrium policy can be discontinuous in v_I at $v_I = \overline{v}$, despite including noise in the voter's utility. This is because at $v = \overline{v}$ and $x = 0$ the signal is uninformative and, thus, the voter remains indifferent between I and C. Instead, the voter has a strict preference at $x = 0$ whenever $v \neq \overline{v}$. However, the incumbent's problem is well-behaved for $v > \overline{v}$ and $v \leq \overline{v}$, delivering existence of a differentiable x^* on each side of $\overline{v}$. Furthermore, numerical examples suggest that the equilibrium policy is unique and that the discontinuity does not emerge when ρ is large.

Second, given the symmetry in our setup, any pair of policies x and $-x$ induces the same posterior distribution in expectation, therefore $P_v(x) = P_v(-x)$. Recall that the incumbent's ideal point is $\hat{x} > 0$. Thus, in equilibrium, the incumbent never implements a policy $x < 0$.

Lemma (A1). *If $|x| > |x'|$ then policy experiment x is Blackwell more informative than x'.*

Proof. The noise term is distributed normally and thus satisfies the MLRP property. Furthermore, fixing an x on either side of zero, the policy choice and the state of the world are strict complements. This can be verified by noting that, for any $z > y > 0$, we have

$$-(z-1)^2 + (z+1)^2 > -(y-1)^2 + (y+1)^2,$$

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

Lemma 1. *The voter reelects the incumbent if:*

$$\mathbb{E}[\omega|x, s_\omega] > -\frac{\mathbb{E}[\theta|s_\theta]}{4\hat{x}}. \tag{3}$$

If $\mathbb{E}[\omega|x, s_\omega] < -\frac{\mathbb{E}[\theta|s_\theta]}{4\hat{x}}$, then the voter elects the challenger. Otherwise, if the voter is indifferent, then each candidate wins the election with probability $1/2$.[18]

[18] The indifference-breaking assumption makes the incumbent's problem continuous at $x = 0$ for all v. It is only

Proof. We define μ as the voter's posterior belief that $\omega = 1$, $Pr(\omega = 1|x, s_\omega)$. The voter's expected utility from re-electing the incumbent is greater than her utility from electing the challenger if

$$-\mu(\hat{x}_C - 1)^2 - (1-\mu)(\hat{x}_C + 1)^2 \leq -\mu(\hat{x}_I - 1)^2 - (1-\mu)(\hat{x}_I + 1)^2 + \mathbb{E}[\theta|s_\theta].$$

Substituting $\hat{x}_C = -\hat{x}_I = \hat{x}$, the above reduces to

$$-\frac{\mathbb{E}[\theta|s_\theta]}{4\hat{x}} \leq 2\mu - 1 = \mathbb{E}[\omega|x, s_\omega].$$

□

Definition 4. *Let* $\overline{\mu}_v = \frac{1}{2} - \frac{v}{8\hat{x}}$ *and* $\lambda_v = \ln\left(\frac{(1-\gamma)\overline{\mu}_v}{\gamma(1-\overline{\mu}_v)}\right)$.

Lemma (A2). *Assume* $\rho = 0$. *The probability of reelection for the incumbent is*

$$P_v(x) = \gamma\Big(1 - \Phi\big(\frac{\lambda_v}{4|x|} - 2|x|\big)\Big) + (1-\gamma)\Big(1 - \Phi\big(\frac{\lambda_v}{4|x|} + 2|x|\big)\Big). \tag{5}$$

Proof. Recall from the proof of Lemma 1 that μ is the voter's updated belief that $\omega = 1$. By Bayes rule we have:

$$\mu(x, s_\omega) = \frac{\gamma\phi\big(s_\omega + (x-1)^2\big)}{\gamma\phi\big(s_\omega + (x-1)^2\big) + (1-\gamma)\phi\big(s_\omega + (x+1)^2\big)}.$$

Thus, using Lemma 1, the incumbent's probability of being re-elected is given by

$$Pr\left(\frac{\gamma\phi\big(s_\omega + (1-x)^2\big)}{\gamma\phi\big(s_\omega + (1-x)^2\big) + (1-\gamma)\phi\big(s_\omega + (-1-x)^2\big)} > \overline{\mu}_v\right), \tag{6}$$

where ϕ is the PDF of the standard normal distribution.

From the incumbent's perspective, s_ω is probabilistic, therefore 6 can be rewritten as

$$\begin{aligned} &\gamma \cdot Pr\left(\frac{\gamma\phi\big(-(\hat{x}-1)^2 + \epsilon + (x-1)^2\big)}{\gamma\phi\big(-(\hat{x}-1)^2 + \epsilon + (x-1)^2\big) + (1-\gamma)\phi\big(-(\hat{x}+1)^2 + \epsilon + (x+1)^2\big)} > \overline{\mu}_v\right) \quad (7)\\ &+(1-\gamma) \cdot Pr\left(\frac{\gamma\phi\big(-(\hat{x}+1)^2 + \epsilon + (x-1)^2\big)}{\gamma\phi\big(-(\hat{x}+1)^2 + \epsilon + (x-1)^2\big) + (1-\gamma)\phi\big(-(\hat{x}+1)^2 + \epsilon + (x+1)^2\big)} > \overline{\mu}_v\right). \end{aligned}$$

Equation 7 further reduces to

$$\gamma \cdot Pr\left(\frac{\gamma\phi(\epsilon)}{\gamma\phi(\epsilon) + (1-\gamma)\phi(4x+\epsilon)} > \overline{\mu}_v\right) + (1-\gamma) \cdot Pr\left(\frac{\gamma\phi(-4x+\epsilon)}{\gamma\phi(-4x+\epsilon) + (1-\gamma)\phi(\epsilon)} > \overline{\mu}_v\right), \tag{8}$$

consequential for a measure 0 set of parameters and does not otherwise affect our results.

and we can rewrite this probability as

$$\gamma \cdot Pr\left(e^{-\frac{\epsilon^2}{2}+\frac{(4x+\epsilon)^2}{2}} > \frac{\overline{\mu}_v(1-\gamma)}{\gamma(1-\overline{\mu}_v)}\right) + (1-\gamma)\cdot Pr\left(e^{-\frac{-(-4x+\epsilon)^2}{2}+\frac{\epsilon^2}{2}} > \frac{\overline{\mu}_v(1-\gamma)}{\gamma(1-\overline{\mu}_v)}\right). \tag{9}$$

Suppose that $x > 0$. After rearranging and applying a logarithmic transformation, the above obligingly reduces to

$$\gamma \cdot Pr\Big(\epsilon > \frac{\lambda_v}{4x} - 2x\Big) + (1-\gamma)\cdot Pr\Big(\epsilon > \frac{\lambda_v}{4x} + 2x\Big), \tag{10}$$

as claimed. A similar derivation yields the expression for $x < 0$. □

Definition 5. *Let* $\Delta_v^- = \frac{\lambda_v}{4x} - 2x$ *and* $\Delta_v^+ = \frac{\lambda_v}{4x} + 2x$.

Lemma (A3).

1. $\frac{\partial \Delta_v^+}{\partial x} = \frac{-\lambda_v}{4x^2} + 2 = -\frac{1}{x}\Delta_v^-$, *and*
2. $\frac{\partial \Delta_v^-}{\partial x} = \frac{-\lambda_v}{4x^2} - 2 = -\frac{1}{x}\Delta_v^+$.

Proof. Follows immediately by differentiating. □

Definition 6. *Let* $\Gamma_v(x) = \gamma\Delta_v^+\phi(\Delta_v^-) + (1-\gamma)\Delta_v^-\phi(\Delta_v^+)$ *and* $\Omega_v(x) = \gamma\phi(\Delta_v^-) + (1-\gamma)\phi(\Delta_v^+)$.

Lemma (A4). *We have* $\frac{\partial P_v}{\partial x} = \frac{1}{x}\Gamma_v(x)$ *and* $\frac{\partial^2 P_v}{\partial x^2} = -\frac{1}{x^2}\Gamma_v(x) + \frac{1}{x}\frac{\partial \Gamma_v}{\partial x}$, *where* $\frac{\partial \Gamma_v}{\partial x} = \frac{1}{x}\Big[\Delta_v^-\Delta_v^+\Gamma_v(x) - \Omega_v(x)\Big]$.

Proof. Taking the derivative of $P_v(x)$ with respect to x yields:

$$\begin{aligned}\frac{\partial P_v}{\partial x} &= \gamma\phi\big(\frac{\lambda_v}{4x} - 2x\big)\big[\frac{\lambda_v}{4x^2} + 2\big] + (1-\gamma)\phi\big(\frac{\lambda_v}{4x} + 2x\big)\big[\frac{\lambda_v}{4x^2} - 2\big] \\ &= \frac{1}{x}\Gamma_v(x).\end{aligned} \tag{11}$$

Next, we differentiate Γ_v and obtain:

$$\begin{aligned}\frac{\partial \Gamma_v}{\partial x} &= \gamma\frac{\partial \Delta_v^+}{\partial x}\phi(\Delta_v^-) - \gamma\Delta_v^+\frac{\partial \Delta_v^-}{\partial x}\Delta_v^-\phi(\Delta_v^-) + (1-\gamma)\frac{\partial \Delta_v^-}{\partial x}\phi(\Delta_v^+) - (1-\gamma)\Delta_v^-\frac{\partial \Delta_v^+}{\partial x}\Delta_v^+\phi(\Delta_v^+) \\ &= \frac{1}{x}\gamma\phi(\Delta_v^-)(\Delta_v^-(\Delta_v^+)^2 - 1) + \frac{1}{x}(1-\gamma)\phi(\Delta_v^+)(\Delta_v^+(\Delta_v^-)^2 - 1) \\ &= \frac{1}{x}\Big[\Delta_v^-\Delta_v^+\Gamma_v(x) - \Omega_v(x)\Big].\end{aligned}$$

The second derivative of $P_v(x)$ then follows straightforwardly by differentiating. □

Lemma (A5). *Any equilibrium policy x^* must solve:*

$$2(\hat{x} - x) + (\beta + 4\hat{x}^2)\Big[\rho\Big(\frac{1+v}{2} \cdot \frac{\partial P_1}{\partial x} + \frac{1-v}{2} \cdot \frac{\partial P_{-1}}{\partial x}\Big) + (1-\rho)\frac{\partial P_v}{\partial x}\Big] = 0.$$

Proof. We show that any equilibrium policy x^* must solve the first-order condition. First, note that the objective function is continuously differentiable in x. Thus, if there is an interior maximizer it must solve the first-order condition. Second, $\beta + 4\hat{x}^2 < \infty$ implies $x^* < \infty$. Hence, recalling that $x^* \geq 0$, a maximizer exists.

Finally, we show that $x^* > 0$. To do so, we show that the objective function is increasing as x increases away from 0. We have $\lim_{x \to 0} \frac{\partial P_{\mathbb{E}[\theta|s_\theta]}}{\partial x} = \lim_{x \to 0} \frac{1}{x}\Gamma_v(x) = 0$ because the normal PDF goes to 0 faster than any polynomial goes to ∞. Therefore, if $x \to 0$ then the first-order condition goes to $2\hat{x} > 0$, as required. □

Lemma 2. *Assume $\rho = 0$. Let $x > 0$.*

1. *Suppose I has an electoral advantage.*
   - *(a) If I has a valence disadvantage and a policy advantage, then the probability of winning is single-dipped in x;*
   - *(b) Otherwise, the probability of winning is strictly decreasing in x.*
2. *Suppose I has an electoral disadvantage.*
   - *(a) If I has a valence advantage and a policy disadvantage, then the probability of winning is single-peaked in x;*
   - *(b) Otherwise, the probability of winning is strictly increasing in x.*

Symmetric results hold for $x < 0$.[19]

Proof. We break our analysis into cases depending on I's valence and electoral advantages. We begin with the cases where the probability of winning is monotonic in x. Additionally, here, we prove the result stated in the footnote for an incumbent without an electoral advantage or disadvantage. Note, from Lemma (A4) $\frac{\partial P_v}{\partial x}$ is negative if and only if equation (11) is negative.

Case 1: We first prove the result for part 1(b). Assume the incumbent has an electoral advantage but does not have a valence disadvantage.[20]

[19] If v is exactly equal to $\overline{v}$, then I's probability of winning is increasing when I has a policy advantage, and instead decreasing when he has a policy disadvantage.

[20] Note this covers all cases when I has an electoral advantage besides 1(a), since it impossible to have an electoral advantage without having an advantage on at least one dimension.

By definition, $v \geq 0$ and $v > -4\hat{x}\mathbb{E}[\omega]$. The incumbent having an electoral advantage, $v > -4\hat{x}\mathbb{E}[\omega]$, implies that $\lambda_v < 0$. Hence, (11) is always negative if $\frac{\lambda_v}{2x^2} + 1 < 0$, which holds for all $x \in [0, \sqrt{\frac{-\lambda_v}{2}}]$.

To finish proving part 1(b) of the lemma, we show that (11) is also negative for $x > \sqrt{\frac{-\lambda_v}{2}}$. If $x > \sqrt{\frac{-\lambda_v}{2}}$ then $\frac{\lambda_v}{2x^2} + 1 > 0$. Therefore, (11) is negative if and only if:

$$\frac{\phi(\frac{\lambda_v}{4x} - 2x)}{\phi(\frac{\lambda_v}{4x} + 2x)} < \frac{1-\gamma}{\gamma} \left(\frac{-\frac{\lambda_v}{8x^2} + 1}{\frac{\lambda_v}{8x^2} + 1} \right). \tag{12}$$

Which we rewrite as:

$$e^{-\frac{1}{2}(\frac{\lambda_v}{4x} - 2x)^2 + \frac{1}{2}(\frac{\lambda_v}{4x} + 2x)^2} < \frac{1-\gamma}{\gamma} \left(\frac{-\frac{\lambda_v}{8x^2} + 1}{\frac{\lambda_v}{8x^2} + 1} \right). \tag{13}$$

Applying a logarithmic transformation to both sides the above reduces to:

$$\lambda_v < \ln \left(\frac{1-\gamma}{\gamma} \frac{\frac{-\lambda_v}{8x^2} + 1}{\frac{\lambda_v}{8x^2} + 1} \right), \tag{14}$$

which holds if and only if:

$$\frac{\overline{\mu}_v(1-\gamma)}{\gamma(1-\overline{\mu}_v)} < \frac{1-\gamma}{\gamma} \frac{\frac{-\lambda_v}{8x^2} + 1}{\frac{\lambda_v}{8x^2} + 1}. \tag{15}$$

Finally, this condition further simplifies to:

$$2\overline{\mu}_v + \frac{\lambda_v}{8x^2} < 1. \tag{16}$$

Recall that $\lambda_v < 0$ by $v > \overline{v}$. Thus, a sufficient condition for (16) to always hold is that $\overline{\mu}_v \leq \frac{1}{2}$. Expanding, $\overline{\mu}_v = \frac{1}{2} - \frac{v}{8\hat{x}} \leq \frac{1}{2}$ by assumption that $v \geq 0$, which completes the argument.

Case 2: Next, we prove part 2(b). Assume that the incumbent has an electoral disadvantage and does not have a valence advantage. Then, by definition, $v \leq 0$ and $v < \overline{v}$. By $v < \overline{v}$ we have $\lambda_v > 0$. Thus, the same rearrangement of (11) as in the previous case yields that $\frac{\partial P_v}{\partial x} > 0$ if and only if:

$$2\overline{\mu}_v + \frac{\lambda_v}{2x^2} > 1. \tag{17}$$

By assumption $v \leq 0$, hence, $\overline{\mu}_v \geq \frac{1}{2}$. Together with $\lambda_v > 0$ this yields that 17 always holds.

Case 3: Next, we prove part 2(a). Assume the incumbent has an electoral disadvantage, a valence advantage, and policy disadvantage. Thus, $0 < v < \overline{v}$ and, by definition of $\overline{v}$, $0 < \overline{v}$ implies that $\gamma < 1/2$. By $v < \overline{v}$ we have $\lambda_v > 0$. Therefore, $\frac{\partial P_v}{\partial x} < 0$ if and only if (16) holds. Thus, to prove the result we show there is a unique cutoff such that (16) fails for all x below the cutoff and holds for all x above. First, $\lim_{x\to 0}$LHS (16) $= \infty > 1$, because $\lambda_v > 0$. Second $\lim_{x\to\infty}$LHS (16) $= 2\overline{\mu}_v < 1$, where the inequality holds because $\overline{\mu}_v < \frac{1}{2}$ by $v > 0$. To complete the argument, notice that LHS (16) is clearly strictly decreasing in x because $\lambda_v > 0$. Furthermore, this implies that I's probability of winning is maximized at the unique $x > 0$ that solves (16) at equality.

Case 4: To conclude the proof of the lemma we now prove part 1(a). Assume the incumbent has an electoral advantage, a valence disadvantage, and a policy advantage. Therefore, $v \in (\overline{v}, 0)$, which further imposes $\gamma > \frac{1}{2}$ by $\overline{v} < 0$. Recall that $v < \overline{v}$ implies $\lambda_v < 0$. Thus, the same argument as in Case 1 yields that (11) is negative for all $x \in [0, \sqrt{\frac{-\lambda_v}{2}}]$. As before, if $x > \sqrt{\frac{-\lambda_v}{2}}$ then (11) is negative if and only if (16) holds. By $\lambda_v < 0$ the LHS (16) is strictly increasing in x. Furthermore, $\lim_{x\to\infty}$LHS (16) $= 2\overline{\mu}_v > 1$ because $v < 0$ implies $\overline{\mu}_v > 1/2$. Hence, there exists a unique cutoff such that (11) is negative if and only if x is below the cutoff. Specifically, this cutoff is given by the unique $x > 0$ that solves (16) at equality.

Case 5: Finally, we consider the case where the incumbent does not have an electoral advantage or disadvantage, $v = \overline{v}$. In this case, $\overline{\mu}_v = \gamma$ which implies $\lambda_v = 0$, and (11) simplifies to:

$$2\gamma\phi(-2x) + (-2)(1-\gamma)\phi(2x) \tag{18}$$

$$= 2\mathbb{E}[\omega]\phi(2x), \tag{19}$$

where the second line follows from symmetry of the pdf of the normal distribution and $\mathbb{E}[\omega] = 2\gamma - 1$. Thus, (19) is strictly positive if I has a policy advantage, strictly negative if I has policy disadvantage, and 0 if $\mathbb{E}[\omega] = 0$.

□

Lemma 3. *Assume $\rho = 0$. There exists a $\hat{v}_G$ such that: if $v < \hat{v}_G$ then the incumbent gambles in equilibrium, $x^* > \hat{x}$, and if $v > \hat{v}_G$ then the incumbent does not gamble, $x^* < x$.*[21] *If the incumbent has a policy disadvantage, then $\hat{v}_G \in (0, \overline{v})$. Otherwise, $\hat{v}_G = \overline{v}$.*

Proof. We break the argument in two cases, depending on whether I has a policy advantage or not.

Case 1: I has a policy disadvantage. Thus, $\mathbb{E}[\omega] < 0$, which implies $0 < \overline{v}$.

First, we show that if $v \geq \overline{v}$ then the incumbent does not gamble. In this case, $v \geq \overline{v} > 0$ and, therefore, I's probability of winning is decreasing in x by Lemma 2. Consequently, for any policy $x > \hat{x}$ deviating to $\hat{x}$ gives I higher policy utility and a greater probability of winning. Thus, if $v \geq \overline{v}$ then $x^* \leq \hat{x}$.

[21] At $v = \hat{v}_G$ the incumbent may gamble or not depending on the parameters.

Second, consider $v \leq 0 < \overline{v}$. By Lemma 2 the incumbent's probability of winning is increasing in x. Therefore, for any $x < \hat{x}$ I deviating to $\hat{x}$ gives I higher policy utility and a greater probability of winning. Thus, if $v < 0$ then $x^* \geq \hat{x}$. Because the incumbent has quadratic utility of policy, the marginal static loss from decreasing the policy away from $\hat{x}$ is zero at $x = \hat{x}$, but the marginal increase in the probability of winning is strictly positive. As such, it must be the case that $x^* > \hat{x}$.

Finally, assume $v \in (0, \overline{v})$. From Lemma 2, the incumbent's probability of winning is maximized at the positive solution to:

$$2\overline{\mu}_v + \frac{\lambda_v}{2x^2} = 1, \tag{20}$$

which we denote as x_v^w. If $x_v^w < \hat{x}$ then the incumbent's probability of winning is decreasing in x for $x > \hat{x}$. Thus, $x^* \leq \hat{x}$. On the other hand, if $x_v^w > \hat{x}$ then the incumbent's probability of winning is increasing in x for $x < \hat{x}$. Thus, $x^* \geq \hat{x}$. Solving (20) explicitly for x_v^w yields:

$$x_v^w = \sqrt{\frac{\lambda_v}{2 - 4\overline{\mu}_v}}.$$

Recall that $\lambda_v > 0$ for $v < \overline{v}$ and that $\overline{\mu}_v = 1/2$ when evaluated at $v = 0$. Thus, $\lim_{v \to 0} x_v^w = \infty$. Additionally, $\lim_{v \to \overline{v}} x_v^w = 0$, because $\lambda_v = 0$ at $v = \overline{v}$.

Differentiating we have $\frac{\partial \overline{\mu}_v}{\partial v} = -\frac{1}{8\hat{x}} < 0$ and $\frac{\partial \lambda_v}{\partial v} = \frac{1}{\overline{\mu}_v(1-\overline{\mu}_v)} \frac{\partial \overline{\mu}_v}{\partial v} < 0$. Therefore, $\frac{\partial x_v^w}{\partial v} < 0$. This implies there is a unique cutoff $\hat{v}_G \in (0, \overline{v})$ such that $x^* > \hat{x}$ if and only if $v < \hat{v}_G$, as required. Specifically, $\hat{v}_G$ is given by the v that solves

$$\hat{x} = \sqrt{\frac{\lambda_v}{2 - 4\overline{\mu}_v}}.$$

Note, this implies that at $\hat{v}_G$ the incumbent does not gamble, specifically $x^* = \hat{x}$.

Case 2: I has does not have a policy disadvantage. In this case, $0 \geq \overline{v}$. To prove the lemma we need to show that if $v > \overline{v}$ then the incumbent chooses $x^* \leq \hat{x}$, and if $v < \overline{v}$ then instead $x^* > \hat{x}$.

First, assume $v > \overline{v}$ and $v \geq 0 \geq \overline{v}$. Since at least one inequality in the second condition must hold strictly, then the probability of winning is strictly decreasing in x by Lemma 2. Therefore, $x^* \leq \hat{x}$, and the incumbent never gambles.

Second, if $v < \overline{v} \leq 0$ then by Lemma 2 the incumbent's probability of winning is strictly increasing in x. Therefore, by the same argument earlier, I's marginal cost of moving x away from $\hat{x}$ is 0,which implies that $x^* > \hat{x}$ and the incumbent always gambles when $v < \overline{v}$, as required.

Finally, we need to show that I does not gamble when $\overline{v} < v < 0$. From the proof of Lemma

2 there exists a cutoff x_v^w such that I's probability of winning is increasing in x for $x > x_v^w$ and decreasing in x for $x < x_v^w$.

We first prove that the optimal policy cannot be above x_v^w. Because the voter's signal becomes perfectly informative about ω when $x \to \infty$, we have $\lim_{x\to\infty} P_v(x) = \gamma$. Thus, the incumbent's expected utility from choosing any $x \geq x_v^w$, is bound above by $-4\hat{x}^2(1-\gamma) + \gamma\beta$, which is the payoff from getting policy $x = \hat{x}$ and winning with probability γ. On the other hand, because $v > \overline{v}$, if $x = 0$ then the incumbent wins with probability 1. Thus, the incumbent's equilibrium utility is bound below by the expected utility from choosing $x = 0$ and winning for sure: $-\hat{x}^2 + \beta$. Consequently, a sufficient condition to ensure $x^* < \hat{x}_v^w$ is that

$$-\hat{x}^2 + \beta > -4\hat{x}^2(1-\gamma) + \gamma\beta$$
$$\Leftrightarrow \beta > \frac{\hat{x}^2(4\gamma - 3)}{1-\gamma},$$

which holds by Assumption 1.

To conclude the proof, we argue that $x^* \notin (\hat{x}, x_\theta^w]$. To see this, consider any $x' \in (\hat{x}, x_\theta^w]$. In this case, deviating to $\hat{x}$ yields a greater policy utility and a higher probability of winning, because $P_\theta(x)$ is decreasing for $x < x_\theta^w$.

For completeness, we now also characterize behavior at $v = \overline{v}$. From the proof of Lemma 2, if I has a policy advantage ($\overline{v} < 0$) then the probability of winning is strictly increasing in x, and thus I gambles. If instead $v = \overline{v} = 0$ then I's probability of winning is flat in x, and thus $x^* = x$. □

Lemma (A6). *Assume $\rho = 0$. Then* $\lim_{v\to\overline{v}^+} x^* = 0$.

Proof. We break the argument in to three steps. The first step derives a bound on I's probability of winning when $v = \overline{v}$. The second step uses this to derive a bound on I's equilibrium payoff when $v = \overline{v}$. Finally, the third step uses this bound to prove the lemma.

Step 1. We first establish that the $\overline{v}$ politician's probability of winning is bound above by $\max\{\gamma, 1/2\}$. Furthermore, from the proof of Lemma 2, at $v = \overline{v}$ I's probability of winning is monotonic in $x > 0$. With these observations in hand we now prove the bound by splitting the argument into two cases, depending on γ.

First, suppose that $\gamma > 1/2$. From Case 5 of the proof of Lemma 2 if $v = \overline{v}$ then I's probability is strictly increasing in x. As $x \to \infty$ the voter's signal is perfectly informative. Thus, I's maximum probability of winning is γ. Second, consider $\gamma \leq 1/2$. In this case, the proof of Lemma 2 yields that I's probability of winning is weakly decreasing in x. Thus, I's probability of winning is maximized at $x = 0$. At $v = \overline{v}$ the voter is indifferent between I and C ex ante, and if $x = 0$ then s_ω is fully uninformative. Thus, the voter remains indifferent between I and C and by assumption the voter reelects I with probability 1/2, as required.

Step 2. We now show that when $v = \overline{v}$ the incumbent's equilibrium payoff must be lower than his payoff if he could choose policy $x = 0$ and win the election with probability 1. Let $U_v(x) = -(x-\hat{x})^2 + P_v(x)\beta - \big(1 - P_v(x)\big)4\hat{x}^2$ be the expected utility to an incumbent with valence v from choosing policy x. Additionally, define $U^0 = \beta - \hat{x}^2$ as I's payoff from getting $x = 0$ and winning with probability 1. The above argument yields that for $v = \overline{v}$ we have $U_{\overline{v}}(x) \leq \max\{\frac{1}{2}, \gamma\}\beta - 4\hat{x}^2\big(1 - \max\{\frac{1}{2}, \gamma\}\big)$, which is the utility from choosing $\hat{x}$ and winning with the highest possible probability. Thus, a sufficient condition for $U^0 > U_v(x)$ to hold for all x is that:

$$\beta - \hat{x}^2 > \max\left\{\frac{1}{2}, \gamma\right\}\beta - 4\hat{x}^2\left(1 - \max\left\{\frac{1}{2}, \gamma\right\}\right).$$

If $1/2 \geq \gamma$ then this inequality clearly holds. If $\gamma > 1/2$ then the above inequality reduces to $\beta > \frac{\hat{x}^2(4\gamma-3)}{1-\gamma}$, which is true by Assumption 1.

Step 3. Finally, we show that for any $x > 0$ there exists δ_x such that if $|v - \overline{v}| < \delta_x$ then $U^0 - U_v(x) > 0$. If $x > 0$ then $U_v(x)$ is continuous in v. Hence,

$$\lim_{v \to \overline{v}} U^0 - U_v(x) = U^0 - U_{\overline{v}}(x).$$

By step 1 of the proof $U^0 - U_{\overline{v}}(x) > 0$ for any $x > 0$. For $v > \overline{v}$ we have $U_v(0) = U^0$. Thus, if $\lim_{v \to \overline{v}^+} x^* > 0$, then there exists v arbitrarily close to $\overline{v}$ such that $U^0 > U_v(x^*)$, a condtradiction. □

Proposition 1. *Assume $\rho = 0$. There exists a cut-point $\underline{v} \geq -1$, with $\underline{v} < \overline{v}$, such that, if $v \in (\underline{v}, \overline{v})$ then x^* is decreasing in v. Otherwise, x^* is increasing in v.*

Proof. By Lemma (A5) any equilibrium policy must solve the first-order condition. Applying the implicit function theorem we have

$$\frac{\partial x^*}{\partial v} = -(\beta + 4\hat{x}^2)\frac{\frac{\partial^2 P_v}{\partial v \partial x}}{-2 + \frac{\partial^2 P_v}{\partial x^2}}.$$

Thus, $\frac{\partial x^*}{\partial v} < 0$ if and only if:

$$\frac{\partial^2 P_v}{\partial v \partial x} < 0.$$

$$\Leftrightarrow \frac{1}{2x^2}\frac{\partial \lambda_v}{\partial v}\Big(\gamma\phi(\Delta_v^-) + (1-\gamma)\phi(\Delta_v^+) + \gamma\phi'(\Delta_v^-)\Delta_v^+ + (1-\gamma)\phi'(\Delta_v^+)\Delta_v^-\Big) < 0,$$

Recall that $\frac{\partial \lambda_v}{\partial v} < 0$, as such, the above inequality holds if and only if:

$$\begin{aligned}
&\gamma\phi(\Delta_v^-) + (1-\gamma)\phi(\Delta_v^+) + \gamma\phi'(\Delta_v^-)\Delta_v^+ + (1-\gamma)\phi'(\Delta_v^+)\Delta_v^- > 0 \\
&\Leftrightarrow \gamma\phi(\Delta_v^-) + (1-\gamma)\phi(\Delta_v^+) - \gamma\phi(\Delta_v^-)\Delta_v^-\Delta_v^+ - (1-\gamma)\phi(\Delta_v^+)\Delta_v^+\Delta_v^- > 0 \\
&\Leftrightarrow (1-\Delta_v^-\Delta_v^+)\big(\gamma\phi(\Delta_v^-) + (1-\gamma)\phi(\Delta_v^+)\big) > 0.
\end{aligned}$$

Therefore, $\frac{\partial x^*}{\partial v} < 0$ if and only if $1 - \Delta_v^-\Delta_v^+ > 0$, which can be rewritten as:

$$4(x^*)^2\big(1 + (x^*)^2\big) > \lambda_v^2. \tag{21}$$

First, we show that if $v > \overline{v}$ then inequality (21) fails, and thus $\frac{\partial x^*}{\partial v} > 0$ for all $v > \overline{v}$. Lemma (A6) shows that $\lim_{v \to \overline{v}^+} x^* = 0$. Therefore, x^* must be increasing in v for $v > \overline{v}$ sufficiently close to $\overline{v}$. Hence, (21) fails for all v sufficiently close to $\overline{v}$. Suppose that eventually (21) holds at some $v > \overline{v}$. Since x^* and λ_v are continuous in v there exists a point $v' > \overline{v}$ such that (21) fails for $v < v'$, and holds with equality at $v = v'$.

Because (21) fails for $v < v'$ it must be that LHS (21) is increasing faster in v than RHS (21) at $v = v'$. To prove that (21) fails for all $v > \overline{v}$ we show that the existence of such a v' where LHS (21) is increasing faster than RHS (21) yields a contradiction. Differentiating both sides:

$$\begin{aligned}
\frac{\partial LHS(21)}{\partial v} &= 8\frac{\partial x^*}{\partial v}x^*\Big[1 + 2(x^*)^2\Big] \\
\frac{\partial RHS(21)}{\partial v} &= \frac{\partial \lambda_v}{\partial v}2\lambda_v.
\end{aligned}$$

Since $\frac{\partial \lambda_v}{\partial v} < 0$ and $\lambda_v < 0$ for all $v > \overline{v}$, this yields $\frac{\partial RHS(21)}{\partial v} > 0$. However, by construction, if (21) holds with equality at $v = v'$ then $\frac{\partial x^*}{\partial v}|_{v=v'} = 0$, contradicting that $\frac{\partial LHS(21)}{\partial v} > \frac{\partial RHS(21)}{\partial v}$ at $v = v'$.

Second, suppose $v \leq \overline{v}$. We show there exists $\underline{v} < \overline{v}$ such that $\frac{\partial x^*}{\partial v} < 0$ if and only if $v \in (\underline{v}, \overline{v})$. By construction, $\lambda_v = 0$ at $v = \overline{v}$. Thus, inequality (21) must hold because $x^* > 0$. Since x^* is continuous in v in this range, it must also hold for all v sufficiently close to $\overline{v}$. That is, there must exist $\underline{v}$ such that (21) holds for all $v \in (\underline{v}, \overline{v}]$. We now prove that (21) can only hold over this interval. Specifically, we show that once (21) holds at some v' it must also hold for all $v \in (v', \overline{v})$.

Suppose (21) holds for some $v < v'$ and fails for some $v > v'$. Because both sides of (21) are continuous in v for $v < \overline{v}$, there must exist some v' such that (21) holds with equality at $v = v'$. Recall that $\lambda_v > 0$ for $v < \overline{v}$ and $\frac{\partial \lambda_v}{\partial v} < 0$. Thus, λ_v^2 is decreasing in v for $v < \overline{v}$. Additionally for $v < v'$ we have that $\frac{\partial x_f^*}{\partial v} < 0$. Hence, for $v < v'$ both sides of (21) are decreasing in v. Since (21) holds for $v < v'$, for (21) to hold with equality at $v = v'$ it must be that $\frac{\partial LHS(21)}{\partial v}|_{v=\overline{v}} < \frac{\partial RHS(21)}{\partial v}|_{v=\overline{v}}$. By the same argument as in the previous case, $\frac{\partial LHS(21)}{\partial v}|_{v=\overline{v}} = 0$. In contrast, $\frac{\partial RHS(21)}{\partial v} < 0$ for all

$v < \overline{v}$, a contradiction. Thus, such a v' cannot exist, as required. $\square$

Proposition 2. *If $\hat{x}$ is sufficiently small then $\underline{v} > -1$.*

Proof. When $\hat{x} = \frac{1}{4}$ the voter kicks out the incumbent if $v = -1$, regardless of the signal on the policy dimension. Thus, $x^* = \hat{x}$ when $v = -1$. However, for $v > -1$ the incumbent has a strictly positive probability of winning. Thus, by Lemma 2, we have $\frac{\partial P_v}{\partial x} > 0$ for v sufficiently small. Inspecting I's first-order condition, clearly $x^* = \hat{x}$ cannot be optimal if $\frac{\partial P_v}{\partial x} > 0$. Therefore, for v sufficiently close to -1 we have $x^* > \hat{x}$, which implies that x^* is increasing in v and hence $\underline{v} > 0$ at $\hat{x} = \frac{1}{4}$. As $\underline{v}$ is continuous in $\hat{x}$, this yields $\underline{v} > 0$ for all $\hat{x}$ sufficiently close to $\frac{1}{4}$. $\square$

Lemma (A7). *Assume $\rho = 1$. The equilibrium policy x^* is strictly decreasing in v.*

Proof. From Lemma (A5) any equilibrium policy must solve the first-order condition. Setting $\rho = 1$, and applying the implicit function theorem:

$$\frac{\partial x^*}{\partial v} = -\frac{\beta + 4\hat{x}^2}{2} \frac{\frac{\partial P_1}{\partial x} - \frac{\partial P_{-1}}{\partial x}}{-2 + \frac{1+v}{2}\frac{\partial^2 P_1}{\partial x^2} + \frac{1-v}{2}\frac{\partial^2 P_{-1}}{\partial x^2}}.$$

Therefore, $\frac{\partial x^*}{\partial v} \leq 0$ if and only if $\frac{\partial P_1}{\partial x} - \frac{\partial P_{-1}}{\partial x} \leq 0$. From assumption 2 we have $-1 < \overline{v} < 1$, and so Lemma 2 yields $\frac{\partial P_1}{\partial x} < 0$ and $\frac{\partial P_{-1}}{\partial x} > 0$. Thus, $\frac{\partial x^*}{\partial v} \leq 0$, as required. $\square$

Proposition 3. *There exist $\underline{\rho}$ and $\overline{\rho}$, with $0 < \underline{\rho} \leq \overline{\rho} < 1$, such that if $\rho > \overline{\rho}$ then x^* is strictly decreasing in v. Instead, if $\rho < \underline{\rho}$ then x^* is non-monotonic in v.*

Proof. By Lemma (A5) any equilibrium policy must solve the first-order condition, and the implicit function theorem yields:.

$$\frac{\partial x^*}{\partial v} \propto \rho\left(\frac{\partial P_1}{\partial x} - \frac{\partial P_{-1}}{\partial x}\right) + (1-\rho)\frac{\partial^2 P_v}{\partial v \partial x}. \tag{22}$$

Note that each derivative is continuous in x and by the implicit function theorem x^* is continuous in ρ. The result then follows from Proposition 1, Lemma A7, and continuity of the RHS of 22 in ρ. $\square$

Proposition 4. *If $\hat{x} \to \frac{1}{4}$ then $\underline{\rho} \to 1$.*

Proof. To start, define $M = \sup_{v \in [-,1,1]} \frac{\partial^2 P_v}{\partial v \partial x}$. Therefore, for a given ρ, $\frac{\partial x^*}{\partial v} < 0$ for all v if and only if:

$$\rho\left(\frac{\partial P_1}{\partial x} - \frac{\partial P_{-1}}{\partial x}\right) + (1-\rho)M < 0. \tag{23}$$

By Proposition 3 $\rho\big(\frac{\partial P_1}{\partial x} - \frac{\partial P_{-1}}{\partial x}\big) + (1-\rho)M > 0$ for all ρ sufficiently small. Thus, there exists a $\underline{\rho}$ that solves (23) at equality such that x^* is non-monotonic in v for all $\rho < \underline{\rho}$. To prove the result we show that $\lim_{\hat{x}\to 1/4} \underline{\rho} = 1$. To do so we verify the following limits:

$$\lim_{\hat{x}\to 1/4} \frac{\partial P_1}{\partial x} = \lim_{\hat{x}\to 1/4} \frac{\partial P_{-1}}{\partial x} = 0, \text{ and}$$
$$\lim_{\hat{x}\to 1/4} M > 0.$$

We first prove that $\lim_{\hat{x}\to 1/4} \frac{\partial P_1}{\partial x} = 0$. To start, note that $\lim_{\hat{x}\to 1/4} \overline{\mu}_1 = 0$. Therefore, $\lim_{\hat{x}\to 1/4} \lambda_1 = \ln(0) = -\infty$. This yields

$$\lim_{\hat{x}\to 1/4} \frac{\partial P_1}{\partial x} = \lim_{\hat{x}\to 1/4} \gamma\phi\Big(\frac{\lambda_1}{4x} - 2x\Big)\Big[\frac{\lambda_1}{4x^2} + 2\Big] + (1-\gamma)\phi\Big(\frac{\lambda_1}{4x} + 2x\Big)\Big[\frac{\lambda_1}{4x^2} - 2\Big] = 0,$$

since $\phi\big(\ln(z)\big)$ goes to zero faster than $\ln(z)$ goes to $-\infty$ as $z \to 0$, as desired. The same argument shows that $\lim_{\hat{x}\to 1/4} \frac{\partial P_{-1}}{\partial x} = 0$.

Now fix $\rho < 1$ and consider M. First, we argue that if $\hat{x} \to 1/4$ then $M \geq 0$. Suppose not, so $M < 0$. Thus, x^* is strictly decreasing in v. However, $\lim_{\hat{x}\to 1/4} \frac{\partial P_1}{\partial x} = \lim_{\hat{x}\to 1/4} \frac{\partial P_{-1}}{\partial x} = 0$ implies that $x^* = \hat{x}$ at $v = -1$ and at $v = 1$, contradicting that x^* is strictly decreasing in v. Therefore, $M \geq 0$.

Second, we show that $M > 0$. Suppose not, so $M = 0$. Since $x^* = \hat{x}$ for $v = 1$ and $v = -1$ then $M = 0$ implies that $\frac{\partial^2 P_v}{\partial v \partial x} = 0$ for all v and, hence, $\lim_{\hat{x}\to 1/4} x^* = \hat{x}$ for all v. However, from the proof of Lemma 1 and $\hat{x} \geq 1/4$, we have that $\frac{\partial P}{\partial x}|_{x=\hat{x}} \neq 0$ for $v \in (-1, 1)$ while the marginal cost of moving x away from $\hat{x}$ is 0, contradicting that $x^* = \hat{x}$ for all v. □

Proposition 5. *Assume office benefit is sufficiently large. If the incumbent has an electoral advantage then he chooses $\rho^* = 0$. If the incumbent has a policy disadvantage, a valence disadvantage, and $\mathbb{E}[\omega] < v$, then he chooses $\rho^* = 1$.*

Proof. Note that if β is sufficiently high then I places ≈ 0 weight on policy incentives and chooses policy to maximize his probability of winning. Thus, $x^* \approx \operatorname{argmax}_x P_v(x)$, if $\rho = 0$, and $x^* \approx \operatorname{argmax} \pi P_1(x) + (1-\pi)P_{-1}(x)$. As such, clearly for β sufficiently high the incumbent decision to have a valence test or not is based only on which action can generate a higher probability of winning.

First, we show that if $v > \overline{v}$ then the incumbent never wants to generate a valence test. If $\rho = 0$ then I can win with probability 1 by choosing $x = 0$. Instead, if $\rho = 1$ then I's probability of winning is bound above by $\pi + (1-\pi)\gamma$. Thus, as $\beta \to \infty$, the incumbent's equilibrium payoff is higher with no valence test.

To finish proving the result, next assume that $\mathbb{E}[\omega] < v < 0 < \overline{v}$. If the incumbent chooses $\rho = 0$ then, by Lemma 2, his probability of winning is strictly increasing in x. Thus, if $\rho = 1$ then I's probability of winning is γ. Suppose instead the incumbent chooses $\rho = 1$ and $x = 0$, then he wins for certain if $\theta = 1$ and loses for certain if $\theta = 0$. Thus, his probability of winning from this strategy is π, which is strictly greater than γ by assumption that $\mathbb{E}[\omega] = 2\gamma - 1 < 2\pi - 1 = \mathbb{E}[\theta] = v$. $\square$

B Extensions

Asymmetric Information

In this section, we allow the incumbent to have private information about his valence. Now, the incumbent's policy choice may impact his reelection probability via two channels. As in the baseline model, the implemented policy acts as an experiment and its outcome influences the voter's beliefs about ω. Additionally, the information asymmetry implies that the policy choice may also directly provide the voter with a signal about the incumbent's valence. We show that our qualitative results survive in this richer information setting.

For simplicity, we consider the two limiting cases of $\rho = 1$ and $\rho = 0$. At the beginning of the game nature draws $\theta \in \{-1, 1\}$, according to the commonly known distributions $Pr(\theta = 1) = \pi \in (0, 1)$. Next, the incumbent observes a private signal $s \in \{-1, 1\}$, where $Pr(s = 1|\theta = 1) = Pr(s = -1|\theta = -1) \in (1/2, 1)$. The game then proceeds as before. In the uncertain-valence model, $\rho = 1$, the incumbent's relative valence is fully revealed prior to the election. Instead, in the fixed-valence model, $\rho = 0$, the voter receives no exogenous information about θ. Our solution concept is perfect Bayesian equilibrium, henceforth "equilibrium".

After observing the signal, the incumbent updates his beliefs about θ according to Bayes' rule. Let ψ_s be the incumbent's (interim) posterior belief that $\theta = 1$ conditional on the realization of his private signal. Thus, $0 < \psi_{-1} < \pi < \psi_1 < 1$.

We show that, across the two models, $\rho = 1$ ad $\rho = 0$, information asymmetries do not change our qualitative results from the baseline setup.

1. $\rho = 0$

We begin analyzing the $\rho = 0$ case, and define $x^b_f(v)$ as the equilibrium policy in the baseline fixed-valence model with symmetric uncertainty. In the asymmetric information setting, we let $x^a_f(s)$ denote the equilibrium policy choice of the incumbent after observing the signal s. Finally, let $\mu_\theta(x_1)$ be the voter's updated interim belief about the incumbent's ability after observing his policy choice.

First, we verify that there always exists a pooling equilibrium in which the incumbent implements $x^b_f(v)$ following either signal $s \in \{-1, 1\}$. That is, both types of the incumbent choose the equilibrium policy from the baseline model without asymmetric information.

Lemma 1B. *Suppose $\rho = 0$. There always exists an equilibrium where the incumbent chooses $x^b_f(v)$ following either signal, $x^a_f(-1) = x^a_f(1) = x^b_f(v)$.*

Proof. Suppose that $x^a_f(0) = x^a_f(1) = x^b_f(v)$, which implies $\mu_\theta\big(x^b_f(v)\big) = \pi$. For any x_1 off the path of play assume that $\mu_\theta(x_1) = \pi$. Then the expected utility to the $s \in \{-1, 1\}$ type from any policy x is:

$$-(\hat{x} - x)^2 + P_\pi(x) \cdot \beta - \big(1 - P_\pi(x)\big) \cdot 4\hat{x}^2. \tag{24}$$

By definition $x^b_f(v)$ maximizes equation (24), thus, neither type of I has a profitable deviation. □

In the fixed-valence model, the voter will not observe exogenous information about the incumbent's valence prior to the election. Consequently, fixing the voter's interim posterior $\mu_\theta(x_1)$, the incumbent's dynamically optimal policy is not a function of his own beliefs. In other words, the optimal policy does not depend on the incumbent's private signal. Notice this implies that standard refinements (intuitive criterion, D1, etc.) do not have bite in this setting. For example, following a deviation off the equilibrium path, suppose the voter believes that $s = -1$. Then, neither type has an incentive to deviate, and the conjectured on-path behavior can always be sustained in equilibrium.

Next, we demonstrate that this is the best equilibrium for both types of the incumbent. As a first step, we establish an indifference result.

Lemma 2B. *Suppose $\rho = 0$. In any equilibrium, both types of the incumbent are always indifferent between all policies on the equilibrium path.*

Proof. Consider any two on path policy, x' and x''. For a contradiction, suppose — wlog — that the $s = 1$ type strictly prefers policy x' over policy x''. Recall that both types have the same expected utility for any policy x_1. Thus, the $s = -1$ type also strictly prefers x' over x'', contradicting that x'' is on the path of play. Therefore, in any equilibrium both types must be indifferent between all policies on the equilibrium path. □

To illustrate why our results hold, consider a separating equilibrium. As noted above, fixing the voter's interim posterior, the incumbent's expected dynamic utility from any policy x is not a function of his private information. Therefore, if separation can be sustained in equilibrium, it must be that the incumbent is always indifferent between the policies on the equilibrium path. Furthermore, in any separating equilibrium, an incumbent who observes $s = -1$ must be locating at his dynamically optimal policy from the baseline model. Thus, in a separating equilibrium both types (at best) receive the payoff from the fixed-valence complete information model with $\theta = \psi_{-1}$. In contrast, in the pooling equilibrium described earlier both types receive the payoff from the complete information game with $\theta = \pi > \psi_{-1}$. From here, our next result follows from a standard envelope argument.

Lemma 3B. *Suppose $\rho = 0$. Among all pooling equilibria, the one where both types of the incumbent choose $x_f^*(v)$ maximizes the equilibrium payoff of both types of the incumbent.*

Proof. Consider a pooling equilibrium in which both types of I choose policy x_1. Because the types pool on the same policy the voter's belief about the incumbent's ability after observing x_1 is $\mu_\theta(x_1) = \pi$. Thus, in any such equilibrium the payoff of type s is given by:

$$-(\hat{x} - x_1)^2 + P_\pi(x_1) \cdot \beta - \big(1 - P_\pi(x_1)\big) \cdot 4\hat{x}^2. \tag{25}$$

By construction, $x_f^*(v)$ maximizes (25), as required. □

Proposition 1B. *Suppose $\rho = 0$. Among all equilibria, the equilibrium where $x_f^a(-1) = x_f^a(1) = x_f^b(v)$ maximizes the incumbent's expected utility under each signal.*

Proof. By Lemma 3B the proposed equilibrium is better than any other pooling equilibrium for both types of the incumbent. Now consider an equilibrium with multiple policies on the path of play. By the law of total expectation there must be an on-path policy x' such that $\pi \geq \mu_\theta(x') \equiv \mu'_\theta$. By Lemma 2B both types of the incumbent must be indifferent over all policies on the equilibrium path, thus, the equilibrium payoff of both types of the incumbent is at most:

$$\max_x -(\hat{x} - x)^2 + P_{\mu'_\theta}(x) \cdot \beta - \big(1 - P_{\mu'_\theta}(x)\big) \cdot 4\hat{x}^2. \tag{26}$$

In contrast, in the equilibrium where $x_f^a(-1) = x_f^a(1) = x_f^b(v)$, the equilibrium payoff of both types is given by:

$$\max_x -(\hat{x} - x)^2 + P_\pi(x) \cdot \beta - \big(1 - P_\pi(x)\big) \cdot 4\hat{x}^2. \tag{27}$$

As $\pi \geq \mu'_\theta$, the envelope theorem delivers that (27) is greater than (26), completing the argument. □

This result shows that there is no equilibrium in which the incumbent can do better than the one where he ignores his private information, even if he learns that he is almost certainly the competent type. An important implication follows immediately. If we focus on the equilibrium that provides all types of the incumbent with their highest expected utility, asymmetric information has no impact on the implemented policy in this fixed-valence model. The incumbent acts as if he had no private information, and conditions his choice on the ex-ante valence v. Thus, the unique equilibrium policy choice surviving this refinement is the same as the complete information fixed-valence model.

2. $\rho = 1$

Next, we move to the $\rho = 1$ case. Denote $x_u^b(v)$ the equilibrium policy in the baseline uncertain-valence model with symmetric information. We let $x_u^a(s)$ denote the equilibrium policy choice of the incumbent after observing the signal s in the asymmetric information setting with uncertain valence.

Lemma 4B. *Suppose* $\rho = 1$. *In every equilibrium, the incumbent chooses different policies following each signal,* $x_u^a(-1) \neq x_u^a(1)$. *Furthermore,* $x_u^a(s = -1) = x_u^b(v = \psi_{-1})$ *and* $x_u^a(s = 1) = x_u^b(v = \psi_1)$.

Proof. If $\rho = 1$ the incumbent's type is fully revealed. Therefore, in equilibrium, the voter's interim posterior $\mu_\theta(x_1)$ is electorally irrelevant. Thus, the incumbent's policy choice influences his reelection chances only via experimentation on the policy dimension. As such, in equilibrium the incumbent must act as if there is no asymmetry of information between him and the voter, and implement the dynamically optimal policy given the interim posterior ψ_s. □

The incumbent's relative valence will be fully revealed to the voter before the election. As a consequence, the voter's interim posterior $\mu_\theta(x_1)$ is electorally irrelevant. As such, the incumbent's policy choice influences his reelection chances only via experimentation and voter learning on the policy dimension. Consequently, the incumbent's strategic problem is identical to the baseline model and he always acts *as if* there was no information asymmetry between him and the voter. Thus, he implements the dynamically optimal policy given his interim posterior ψ_s. As a consequence, the *expected* equilibrium policy is decreasing in $v = \pi$. As in the symmetric uncertainty baseline, officeholders with higher valence in expectation enact more moderate policies, all else equal.

We note that the assumption that the incumbent's valence is fully revealed to the voter is not necessary for these results. All that is needed is that the information that is exogenously generated is more significant for the voter than the one signaled by the incumbent's actions (formally, the public signal revealed before the election is more informative than the incumbent's private signal).

Application to Crises

In this section, we apply our framework to examine the impact of an exogenous crisis on policy experimentation. Most of the work in the political economy literature conceptualizes crises as a shock to the actors' policy tastes (e.g., Drazen and Easterly, 2001; Prato and Wolton, 2018; Fernandez and Rodrik, 1991; Guiso et al., 2019; Levy and Razin, 2021; Bils, 2023). We complement this literature by conceptualizing crises as tests that reveal the officeholder's valence (as in Ashworth, Bueno de Mesquita and Friedenberg (2018)), and providing a theory of the supply-side effect on policy experimentation, one that applies even if the crisis is orthogonal to the policy dimension.

In Bils (2023), for example, the crisis and policy dimensions are not orthogonal. The voters' reelection decision does not depend on their beliefs about the state of the world, but rather their beliefs about the incumbent's ability to learn the state. Moreover, voters do not observe a signal of their policy payoffs before voting, and thus the strategic tension for politicians is about how to signal ability via policymaking. In contrast, in our model, the politician does not have an informational advantage over the voter and the policy choice acts as an experiment, with the resulting learning over ω directly altering the voter's electoral decision. In particular, in Bils (2023), unlike in our paper, the incumbent's initial observed electoral standing has no impact on his incentive to choose an extreme policy.

Within our framework, the fixed-valence model ($\rho = 0$) can be interpreted as describing the behavior of the incumbent and voters during a period of business as usual. In this scenario, the incumbent's competence is not tested, and the voter bases her retention decision on her prior beliefs, i.e., on the incumbent's ex-ante net valence $v = \pi$. Conversely, the case of $\rho = 1$ can represent a period of crisis, where the outcome reveals the incumbent's true competence. However, the incumbent must implement the policy experiment on the ideological dimension without knowing what the crisis outcome will be on election day.

By comparing our model under fixed ($\rho = 0$) and uncertain ($\rho = 1$) valence, we thus gain insight into how the risk of a crisis influences the incumbent's incentives to take endogenous risks in policymaking. Proposition 2B compares policymaking in the two versions of the model, and characterizes conditions under which a crisis induces more or less experimentation.

Proposition 2B. *If $v > \overline{v}$ then the incumbent enacts a more extreme policy during times of crisis than during normal times. If $v < \underline{v}$ then the incumbent implements a more moderate policy during times of crisis than during normal times. If $v \in (\underline{v}, \overline{v})$ then the crisis can lead to more or less extreme policies.*

Proof. Denote x_u^* the equilibrium policy under $\rho = 1$, i.e., under a crisis, and x_f^* the equilibrium policy under $\rho = 0$, i.e., under normal times.

First, suppose $v > \overline{v}$. By Proposition 1 x_f^* is increasing in v, whereas x_u^* is decreasing in v by Lemma (A7). Moreover, $x_f^* = x_u^*$ at $v = 1$, thus $x_f^* < x_u^*$, as required. Second, let $v < \underline{v}$. Again, by Proposition 1 x_f^* is increasing in v, while x_u^* is decreasing in v by Lemma (A7). Furthermore, $x_f^* = x_u^*$ at $v = -1$, hence $x_f^* > x_u^*$ for $v < \underline{v}$. □

The crisis can alter the extent of I's existing incentives to control information, that is, it can have a *quantitative* effect on policymaking. This effect always moves policy closer to $\hat{x}$, as the crisis makes new policy information less relevant to the voter's decision, thus reducing the incumbent's incentives to distort policy away from $\hat{x}$. Furthermore, the crisis can have a *qualitative* effect on policymaking, changing the effect of policymaking on the probability of winning, and thus incentivizing the incumbent to switch from gambling to not, or vice versa.

Proposition 2B highlights that the effect of the crisis on policy depends on the incumbent's electoral standing. More precisely, this effect is mediated by the officeholder's expected ability.

First, suppose the incumbent is electorally advantaged. Then, in the fixed-valence model he never gambles, and policy is increasing in v. In contrast, under $\rho = 1$, the incumbent never gambles for high enough v, and may or may not gamble for sufficiently low $v > \overline{v}$. However, policy is always decreasing in v. Thus, regardless of whether the crisis has a qualitative or quantitative effect, it must be that it induces the incumbent to implement more extreme policies.

Instead, if the incumbent is electorally disadvantaged, the crisis may induce more or less policy experimentation. When v is very low, the effect of the crisis must be quantitative, because the incumbent always gambles, both during normal times and during crises. As a consequence, the crisis induces the incumbent to move x closer to $\hat{x}$, with a moderating effect on policy. However, when v is higher the effect may be qualitative or quantitative, and the crisis can thus induce more or less experimentation, depending on the parameter values.

The results in this section describe how a crisis influences an officeholder's incentive to *supply* policy experimentation, even on a dimension largely unrelated to the crisis. Earlier theories instead focus on *demand*-side effects that emerge when the crisis influences voters' preferences over policy. These theories argue that crises should increase policy experimentation (e.g., Tommasi and Velasco, 1996). However, findings in the empirical literature are mixed. Some scholars do find that crises and reform are positively related (Pitlik and Wirth, 2003; Lora and Olivera, 2004; Alesina, Ardagna and Trebbi, 2006), while others find crises may lead to less reform (Pop-Eleches, 2008; Campos, Hsiao and Nugent, 2010; Castanheira, Nicodème and Profeta, 2012; Galasso, 2014; Mian, Sufi and Trebbi, 2014).

In our setting, whether crises induce more or less experimentation depends on the incumbent's ex-ante valence. Failing to account for this interaction, an empirical analysis of the effect of a crisis on reform may recover biased estimates. Furthermore, the bias can go in either direction, which implies that researchers may even recover a zero effect when averaging across different values of v. Considering the supply-side incentives of politicians can therefore help explain why a crisis may lead to less reform, and provide a potential framework to reinterpret the mixed results in the literature. Additionally, our model helps elucidate the exact channels through which this supply-side effect may materialize and how it may be mediated by other features of the competitive environment. These results provide additional implications that are unique to our theory and thus open several potential avenues for future research. In the Appendix, we show that our qualitative results from this section do not necessarily change even if the crisis also shifts the voter's beliefs on the policy dimension.

To conclude, notice that a direct implication of Proposition 2B is that some crises appear as unifying ones, pushing the incumbent's policies closer to the challenger's preferences (compared to the no-crisis counterfactual). Instead, other crises have a polarizing effect and push policy to the extreme. Importantly, in our framework, the difference between a unifying crisis and a polarizing one is not in the nature of a crisis itself. Rather, these are equilibrium effects that emerge, respectively, under incumbents of low and high expected ability even when they face identical crises.

References

Alesina, Alberto, Silvia Ardagna and Francesco Trebbi. 2006. "Who adjusts and when? The political economy of reforms." *IMF Staff Papers* 53(1):1–29.

Alonso, Ricardo and Odilon Câmara. 2016. "Political disagreement and information in elections." *Games and Economic Behavior* 100:390–412.

Ansolabehere, Stephen and James M Snyder. 2000. "Valence politics and equilibrium in spatial election models." *Public Choice* 103(3):327–336.

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

Ashworth, Scott, Ethan Bueno de Mesquita and Amanda Friedenberg. 2018. "Learning about voter rationality." *American Journal of Political Science* 62(1):37–54.

Bernhardt, Dan, Odilon Câmara and Francesco Squintani. 2011. "Competence and ideology." *Review of Economic Studies* 78(2):487–522.

Biglaiser, Gary and Claudio Mezzetti. 1997. "Politicians' decision making with re-election concerns." *Journal of Public Economics* 66(3):425–447.

Biles, Peter. 2012. "Falklands Invasion 'Surprised' Thatcher." *BBC News, December* 28.

Bils, Peter. 2023. "Overreacting and Posturing: How Accountability and Ideology Shape Executive Policies." *Quarterly Journal of Political Science* 18(2):153â–182.

Boleslavsky, Raphael and Christopher Cotton. 2015. "Information and extremism in elections." *American Economic Journal: Microeconomics* 7(1):165–207.

Cai, Hongbin and Daniel Treisman. 2009. "Political decentralization and policy experimentation." *Quarterly Journal of Political Science* 4(1):35–58.

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

Callander, Steven and Bård Harstad. 2015. "Experimentation in federal systems." *The Quarterly Journal of Economics* 130(2):951–1002.

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

Campos, Nauro F, Cheng Hsiao and Jeffrey B Nugent. 2010. "Crises, what crises? New evidence on the relative roles of political and economic crises in begetting reforms." *The Journal of Development Studies* 46(10):1670–1691.

Carrillo, Juan D and Micael Castanheira. 2008. "Information and strategic political polarisation." *The Economic Journal* 118(530):845–874.

Carrillo, Juan D and Thomas Mariotti. 2001. "Electoral competition and politician turnover." *European Economic Review* 45(1):1–25.

Castanheira, Micael, Gaëtan Nicodème and Paola Profeta. 2012. "On the political economics of tax reforms: survey and empirical assessment." *International Tax and Public Finance* 19(4):598–624.

Cheng, Chen and Christopher Li. 2019. "Laboratories of democracy: Policy experimentation under decentralization." *American Economic Journal: Microeconomics* 11(3):125–154.

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.

Drazen, Allan and William Easterly. 2001. "Do crises induce reform? Simple empirical tests of conventional wisdom." *Economics & Politics* 13(2):129–157.

Dunleavy, Patrick and Christopher T Husbands. 1985. *British democracy at the crossroads: voting and party competition in the 1980s.* Taylor & Francis.

Fernandez, Raquel and Dani Rodrik. 1991. "Resistance to reform: Status quo bias in the presence of individual-specific uncertainty." *American Economic Review* pp. 1146–1155.

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

Fu, Qiang and Ming Li. 2014. "Reputation-concerned policy makers and institutional status quo bias." *Journal of Public Economics* 110:15–25.

Galasso, Vincenzo. 2014. "The role of political partisanship during economic crises." *Public Choice* 158(1-2):143–165.

Groseclose, Tim. 2001. "A model of candidate location when one candidate has a valence advantage." *American Journal of Political Science* pp. 862–886.

Guiso, Luigi, Helios Herrera, Massimo Morelli and Tommaso Sonno. 2019. "Global crises and populism: the role of Eurozone institutions." *Economic Policy* 34(97):95–139.

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* 117(3):1089–1104.

Izzo, Federica. Forthcoming. "Who runs when?" *The Journal of Politics* .

Krasa, Stefan and Mattias K Polborn. 2012. "Political competition between differentiated candidates." *Games and Economic Behavior* 76(1):249–271.

Levy, Gilat and Ronny Razin. 2021. "Short-term political memory and the inevitability of polarisation." *London School of Economics working paper* .
URL: *https://drive.google.com/file/d/1a1rRJfGw9bEukdgVHdTmINCi2MF6EQMX/view*

Lora, Eduardo and Mauricio Olivera. 2004. "What makes reforms likely: Political economy determinants of reforms in Latin America." *Journal of Applied Economics* 7(1):99–135.

Majumdar, Sumon and Sharun W Mukand. 2004. "Policy gambles." *American Economic Review* 94(4):1207–1222.

Marsh, David. 1991. "Privatization under Mrs. Thatcher: A review of the literature." *Public Administration* 69(4):459–480.

McAllister, Ian and Donley T Studlar. 1989. "Popular versus elite views of privatization: The case of Britain." *Journal of Public Policy* 9(2):157–178.

Mian, Atif, Amir Sufi and Francesco Trebbi. 2014. "Resolving debt overhang: Political constraints in the aftermath of financial crises." *American Economic Journal: Macroeconomics* 6(2):1–28.

Moore, John. 1992. "British privatization–taking capitalism to the people." *Harvard Business Review* 70(1):115–124.

Parker, David. 2004. "The UK's privatisation experiment: The passage of time permits a sober assessment." *Available at SSRN 514224* .

Pitlik, Hans and Steffen Wirth. 2003. "Do crises promote the extent of economic liberalization?: an empirical test." *European Journal of Political Economy* 19(3):565–581.

Pop-Eleches, Grigore. 2008. *From Economic Crisis to Reform.* Princeton University Press.

Prato, Carlo and Stephane Wolton. 2018. "Rational ignorance, populism, and reform." *European Journal of Political Economy* 55:119–135.

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

Stokes, Donald E. 1963. "Spatial models of party competition." *American political science review* 57(2):368–377.

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

Swann, Dennis. 1988. *The Retreat of the State: Deregulation and Privatization in the UK and US.* University of Michigan Press.

Tommasi, Mariano and Andres Velasco. 1996. "Where are we in the political economy of reform?" *The Journal of Policy Reform* 1(2):187–238.

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.

Young, Stephen. 1986. "The nature of privatisation in Britain, 1979–85." *West European Politics* 9(2):235–252.

Yu, Tinghua. 2022. "Accountability and learning with motivated agents." *Journal of Theoretical Politics* 34(2):313–329.