

Policy Gambles and Valence in Elections*

Peter Bils[†] Federica Izzo[‡]

May 14, 2024

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 the 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 crucially on his 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 valence (and thus electoral advantage) can motivate him to gamble more in equilibrium. However, this relationship between valence and experimentation depends crucially on the nature of the valence dimension.

*This paper was previously circulated with the title 'Policymaking in Times of Crisis'.

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

[‡]Department of Political Science, University of California San Diego.

Introduction

Voters and politicians navigate a world filled with uncertainty, where the outcomes of different policy choices 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.

While these strategic choices significantly impact electoral outcomes, they are not the only factors 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 heterogeneous candidates in which the incumbent can use his policy choice to influence voter learning. 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, and this only occurs if valence is electorally relevant. Furthermore, increasing the incumbent’s valence (and thus electoral advantage) can motivate him to gamble more in equilibrium. However, this relationship between valence and experimentation depends crucially on the nature of the valence dimension.

Our Approach: We analyze a two-period model of electoral accountability. In each period, the officeholder sets a one-dimensional ideological policy, and after the first period a voter decides whether to elect the incumbent or challenger. Players face common uncertainty about the location of the voter’s ideal point, but can learn about it from the outcome of the incumbent’s policymaking. Specifically, the voter observes a noisy signal of her utility from the implemented policy. Thus, policies directly impact utility *and* act as an experiment about the voter’s ideal point. This produces an environment that resonates with intuitions from the retrospective voting literature, where voters learn by experience. Voters update their beliefs after observing past outcomes, and this shifts their preferences over policies and candidates (as described, for example, in Fiorina, 1981).

In addition to policy, however, the voter also cares about the valence of the incumbent and

challenger. Following the literature, we consider two versions of the model. In the first version, each candidate's valence is fixed and known ex-ante (as in, e.g., Groseclose, 2001). This setting can capture known characteristics of the candidates, such as experience, name recognition, or even campaign funds. In the second version, the candidates' relative valence is initially unknown, but is revealed right before the election (as in, e.g., Aragonès and Xefteris, 2017). This may model environments where the officeholder's ability is put to the test (such as during a crisis), or settings where the voter has a private evaluation of the candidates' valence, which is ex-ante unknown and only revealed on election day.

In our model, the amount of voter learning on the policy dimension depends on the extremity of the incumbent's policy choice. More extreme policies increase the distance in the expected outcome as a function of the state and, thus, enhance opportunities for learning. To see this, consider a scenario where a voter experiences a favorable outcome from an extreme leftist policy. In this case, it becomes evident that the policy aligns well with the voter's interests. Outcomes of moderate policies are however less informative. Even if the policy moves slightly in a direction that is not ideal for voters, random fluctuations in the economy may still allow them to experience relatively high welfare.

The politicians in our model have policy preferences and care about winning office. Consequently, when implementing a policy, the officeholder considers both his own ideological preferences and the impact of policymaking on voter learning. In particular, the incumbent faces a tradeoff between choosing a policy closer to his ideal point versus one that induces the optimal amount of information for reelection. In turn, the way the incumbent solves this tradeoff depends on the voter's initial beliefs on the policy dimension *and* on the incumbent's valence relative to the challenger's.

Preview of Results: Absent electoral incentives, the incumbent in our model simply follows his policy preferences and implements his ideal point. However, more extreme policies reveal more information to the voter, thus altering the incumbent's electoral prospects. Consequently, our first results characterize whether the incumbent's reelection incentives prompt him to pursue a policy that is more or less informative than his ideal point. Intuition from existing theories suggests that career concerns should induce the incumbent to *gamble* by implementing a more extreme policy when electorally trailing, and instead choose a safer, more moderate policy, when he is ex ante ahead. Across both versions of our model, we show that this intuition does not always apply when the incumbent's initial (dis)advantage is moderate.

Consider first the fixed-valence model. For incumbents who are very far ahead the probability of re-election is decreasing in the informativeness (or extremism) of the implemented policy. If the incumbent is instead very far behind, then a more informative signal is always beneficial. Thus, consistent with earlier intuition, incumbents who are far ahead never gamble, and those who are

far behind always do. Instead, when the election is ex-ante more competitive, the incumbent's probability of being reelected is non-monotonic in the implemented policy. To see why, suppose that the incumbent is barely trailing ex ante, but his valence is higher than the challenger's valence. This trailing incumbent needs the voter to update positively on the policy dimension to win reelection. However, because he is advantaged on the valence dimension, the policy-relevant state must be very unlikely to favor him; as such, more information is highly likely to hurt his prospects. As a consequence, the incumbent's retention probability is highest at intermediate policies because these maximize the probability of generating a false positive. This trailing incumbent thus behaves like as if he is leading, and implements a policy more moderate than his ideological preferences in order to hamper voter learning. Importantly, this effect does not emerge when the voter is indifferent to the candidates' valence.

In the uncertain-valence model, the incumbent's incentive to gamble is not determined by his ex-ante leading or trailing status. Instead, it is determined by the probability of acquiring an *ex-post* electoral advantage or disadvantage due to the revelation of valence before the election. Consider an incumbent who unlikely to be high valence, but is nonetheless electorally leading because the voter's prior on the ideological dimension is favorable to him. This incumbent expects to lose his initial advantage. As such, he is incentivized to gamble and implement a more extreme policy than he would absent reelection concerns. The symmetric logic applies to a trailing incumbent who is nonetheless likely to have high valence. Again, this implies that the intuition for when we should observe officeholders gamble fails if the election is ex ante highly competitive.

Next, we characterize how the incumbent's valence influences the *intensity* of his strategic incentives to control information, and thus the equilibrium amount of policy experimentation. We show that the effect of valence is fundamentally different across the two versions of the model.

In the fixed-valence model, changing the incumbent's valence has two effects. On one hand, increasing his valence makes the incumbent more attractive, and thus learning on the policy dimension becomes less relevant for his reelection chances. On the other hand, if increasing the incumbent's valence moves him closer to the voter's indifference threshold, then his electoral prospects become more responsive to learning on the policy dimension. For a trailing incumbent, these two effects push in opposite directions, which generates a non-monotonic effect of valence on policy extremism. Instead, for leading incumbents these effects work in the same direction, thus higher valence incumbents choose more extreme policies. Given these forces, we find that incumbents with moderately low valence engage in the most experimentation, whereas incumbents with moderately high valence engage in the least amount of experimentation.

In contrast, in the uncertain-valence model, increasing the incumbent's expected valence increases the probability that he will acquire an ex-post advantage. This unambiguously increases his incentives to prevent voter learning. Therefore, incumbents with lower expected valence always

engage in more policy experimentation. Importantly, we also show that these results are robust to assuming the incumbent has private information about his valence.

Finally, we apply our model to study the effects of a crisis — such as a war, financial recession, or pandemic — on policy reform. To do so, we compare the results of the two models. We interpret the results of the uncertain-valence model as describing the players’ behavior in a time of crisis, which acts as an exogenous test of the incumbent’s competence. Instead, the fixed-valence model describes a period of business as usual where the voter does not learn more about the incumbent’s competence. In our setting, the valence dimension is orthogonal to the policy dimension. Thus, our results capture incentives for policymakers to *supply* experimentation during a crisis, whereas existing arguments emphasize that reforms occur because a crisis alters voters’ *demands*. We establish a conditional effect of crises. Compared to policymaking during normal times, a crisis prompts more cautious reforms when the incumbent’s expected valence is very low, but bolder reforms when the incumbent is very likely to solve the crisis.

Related Literature

Our model relates to the literature on experimentation and learning in strategic interactions (see Hörner and Skrzypacz (2017) for a review). Specifically, we contribute to the study of how reelection concerns impact a politician’s willingness to engage in policy experimentation. Several scholars have studied how such concerns can distort incentives to enact risky policies (Biglaiser and Mezzetti, 1997; Majumdar and Mukand, 2004; Fu and Li, 2014; Dewan and Hortala-Vallve, 2019).¹ Others in this literature have analyzed how decentralizing policymaking impacts incentives for experimentation when policymakers face elections (Rose-Ackerman, 1980; Cai and Treisman, 2009; Cheng and Li, 2019).² We contribute to this body of work by considering how an orthogonal valence dimension influences incentives for electorally accountable politicians to engage in experimentation.

Importantly, the above papers consider a binary policy space, with one risky option and one safe option.³ As such, these works can only analyze a decisionmaker’s choice to experiment or not. Instead, we consider policy experimentation with a continuous space. Doing so allows us to analyze the intensity of the policymaker’s dynamic incentives to take risks and study the equilibrium *amount* of policy experimentation. This is important because a binary policy choice would obfuscate many of the effects of valence on policymaking. In particular, a binary policy space would conceal the

¹Related, others have studied how accountability influences an agent’s incentives to exert (unobservable) effort improving the outcome of policy experiments (Hirsch, 2016; Yu, 2022).

²Another strand of this literature has studied how decentralization affects policy experimentation, but abstract from electoral incentives (Strumpf, 2002; Volden, Ting and Carpenter, 2008; Callander and Harstad, 2015).

³Cai and Treisman (2009) has multiple policies, but the rewards from each are independent and only one policy can be chosen at a time. Thus, the policymaker cannot alter the informativeness of the experiment.

non-monotonic effect of valence we uncover in our framework, as well as the differences across the fixed and uncertain-valence settings.

The learning technology we use relates more closely to the models introduced in Ashworth, Bueno de Mesquita and Friedenbergr (2017) and Izzo (2023). However, there are technical differences between these works and our approach. In both of these papers, as in ours, the policymaker has a continuum of actions, and his choice determines the informativeness of the resulting outcome for the voter. However, Ashworth, Bueno de Mesquita and Friedenbergr (2017) considers a continuous choice of effort which is *unobserved* by the voter, whereas we study an ideological policy choice which is *observed* by the voter. As a consequence, the voter in our model updates her beliefs (and thus ideological preferences) based on the implemented policy as well as the outcome of the experiment. Izzo (2023) also focuses on policymaking along an ideological dimension; however, in that model learning is stark because each policy outcome is either fully informative or completely uninformative, due to the assumption that the noise in the outcome realization is uniformly distributed. In contrast, in our setting policy outcomes are never fully informative, because noise in our model is drawn from a normal distribution, and this leads to less stark behavior by politicians in equilibrium. This is an important change, as a uniform-shock would also obfuscate 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. In our world, players learn about the expected consequences of the various policy reforms. Therefore, more extreme policy choices are more informative. In Callander (2011), players face no uncertainty about expected outcomes, but try to learn about the exact effects of each specific policy. Thus, it is small incremental changes that facilitate more learning about the consequences of different policies. Our approach thus better captures settings where a large reform is most informative about how policies map to outcomes. Furthermore, 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.⁴

More generally, our contribution is unique in that none of the papers mentioned above study how valence considerations influence the incumbent’s strategic incentives to engage in policy experimentation. A prominent literature on elections has studied models in which candidates have an exogenous characteristic, such as valence, that is 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

⁴Callander and Hummel (2014) considers forward-looking parties, but assumes exogenous retention probabilities.

can vary significantly depending on whether valence is ex ante known or ex ante uncertain. Thus, our model also contributes to previous works which have studied 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).

Our paper complements Alonso and Câmara (2016), who also study an 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 true 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). 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. Including policy preferences is necessary for the effects of the intensity of the politician’s incentives to emerge in equilibrium, and thus allows us to obtain more comprehensive understanding of the relationship between valence and experimentation.

Furthermore, 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, as a consequence, the informativeness of outcomes. In contrast, in the Bayesian persuasion framework used in Alonso and Câmara (2016) the incumbent can choose *any* signal structure mapping the state to realized outcomes, which captures situations where the officeholder has more flexibility in controlling the information that reaches the voters and how they interpret such information. Finally, Alonso and Câmara (2016) only considers a setting where the incumbent’s relative valence is uncertain and revealed before the election. Instead, we highlight that the relationship between valence and experimentation is fundamentally different depending on whether valence is known or uncertain ex ante. Interestingly, the predictions from our uncertain-valence model still align with theirs. Both models find that higher-valence incumbents implement less informative experiments.⁵

Finally, our paper contributes to a small formal literature that studies how crises impact political and policy outcomes. Most of the work in this tradition 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, and providing a theory of the

⁵This result emerges in Alonso and Câmara’s baseline model, which assumes a log-concave valence distribution. Their prediction is flipped under log-convex distributions.

supply-side effect on policy, one that applies even if the crisis is orthogonal to the policy dimension. Thus, we build on Ashworth, Bueno de Mesquita and Friedenberg (2018) who also study crises as exogenous informative shocks. However, in their model politicians do not engage in policymaking. In turn, Izzo (2022) studies how such crises influence the self-selection of competent candidates into the electoral arena. We extend this approach by analyzing how an exogenous crisis impacts strategic policy experimentation.

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 and a noisy signal of her policy utility (described below). 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.

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 given by:

$$U_t^i(x_t) = -(x_t - \hat{x}_i)^2 + \mathbb{I}_t^i \beta, \quad (1)$$

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 equals 0 otherwise. For simplicity, we assume $\hat{x}_I = -\hat{x}_C > 0$, thus, the incumbent is right-wing and the challenger is symmetrically left-wing. Furthermore, the candidates' bliss points are common knowledge.

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

$$U_t^V(x_t) = -(x_t - \hat{x}_V)^2 + \theta_t, \quad (2)$$

where \hat{x}_V is the voter's ideal point and θ_t is the valence of the period- t officeholder. Thus, consistent with previous models of elections with valence, we model valence as a dimension that is orthogonal to policy.

The location of \hat{x}_V is initially unknown to all players. This policy may take one of two values, $\hat{x}_V \in \{-1, 1\}$, normalized for simplicity. Players have a common prior belief that $Pr(\hat{x}_V = 1) = \gamma \in (0, 1)$. Before making her retention decision, the voter observes a noisy signal of her utility that

is derived from the implemented policy:

$$\sigma = -(x_1 - \hat{x}_V)^2 + \varepsilon, \quad (3)$$

where ε is a shock drawn from the standard normal distribution, with CDF Φ and PDF ϕ .

Valence: We consider two different versions of the model, which vary the information that is known about the valence of the candidates. Across both versions of the model we assume that the incumbent has no private information about his valence. We later relax this assumption in an extension.

The fixed-valence model. In the fixed-valence model, $\theta_I \in [0, 1]$ and $\theta_C \in [0, 1]$ are realized and publicly observed at the beginning of the game, after which they remain fixed throughout. This version captures scenarios where θ_i represents a candidate's immutable and observable characteristics, such as experience, name recognition, or even campaign funds. Alternatively, this can model a setting where the true valence of each politician is unknown to all *and* no further information about valence is revealed during the course of the game. In this case, we interpret θ_i as i 's expected valence.

The uncertain-valence model. In this version of the model, we assume $\theta_I \in \{0, 1\}$ and $\theta_C \in \{0, 1\}$ are ex-ante unknown to all. Valences are drawn independently, and players share a common prior belief that $Pr(\theta_i = 1) = \pi_i \in (0, 1)$. After the incumbent's policy decision, but before the election, θ_I is revealed and publicly observed. Beliefs about the challenger's valence, however, remain fixed throughout the game.⁶ We can interpret this model as a situation in which the incumbent's competence is under scrutiny, and information is anticipated to be released prior to the election. For example, as mentioned above, the country may be undergoing a crisis whose resolution depends on the incumbent's initially unknown ability. Alternatively, it can capture significant uncertainty about whether the voter will perceive the incumbent or challenger as having greater charisma (or likeability, integrity, etc...) when election day arrives.⁷

The two versions of the model capture the two extremes about how much uncertainty the incumbent faces on the valence dimension when setting policy. This simplifies the analysis and clearly emphasizes the differences between the two settings. However, similar qualitative incentives should hold under less stark assumptions.

⁶Assuming only information about the incumbent is revealed simplifies notation and facilitates the application to crises as valence-revealing events. However, similar results would hold if instead information is revealed about the challenger's valence, or the relative valences of the candidates.

⁷The latter interpretation is consistent with the standard valence shock approach in probabilistic models of elections (see, e.g., Groseclose, 2001).

Finally, to facilitate comparison of the results of the two versions of the model, we introduce the following notation:

$$v_i \equiv \mathbb{E}[\theta_i] \text{ for } i \in \{I, C\}.$$

Thus, in the fixed-valence model $v_i = \theta_i$, while in the uncertain-valence model $v_i = \pi_i$. Notice that, in either case, we have $v_i \in [0, 1]$.

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

1. Nature draws $\hat{x}_V \in \{-1, 1\}$.
2. I chooses policy $x_1 \in \mathbb{R}$.
3. The voter observes the policy choice x_1 , a noisy signal of her realized policy utility σ , and then updates her beliefs about \hat{x}_V .
 - In the uncertain-valence model, the voter also observes the realization of θ_I .
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.

Our solution concept is Perfect Bayesian Equilibrium.

We impose the following assumption on office benefit:

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

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. Note that if the incumbent is not too advantaged on the policy dimension, $\gamma \leq \frac{3}{4}$, then this assumption holds even if $\beta = 0$. Thus, Assumption 1 can be interpreted as assuming that the election is highly competitive when the incumbent and challenger are similar ex-ante.

Additionally, to streamline the analysis, we assume that I is never guaranteed to win the election if $\theta_I = 1$, and he is not guaranteed to lose if $\theta_I = 0$. Specifically, we make an assumption on the incumbent's ideal point, which characterizes the degree of polarization between the candidates.

Assumption 2. *Polarization is sufficiently large:* $\hat{x}_I > \max\left\{\frac{v_c}{4}, \frac{1-v_c}{4}\right\}$.

Before concluding this section, we comment on the timing of policymaking and learning in the model. In our model, the voter learns via experience by observing a noisy signal of the policy outcome which realizes prior to the election. However, it is possible that, once a policy is implemented, its effect on voter welfare may take some time to become visible. Our framework incorporates such frictions, which are captured by the shock term ϵ in the voter's signal σ . Furthermore, one element outside the scope of our model, but which would reinforce our mechanism, is that the attention voters (and media) pay to a policy and its consequences is endogenous to the nature of the policy. Here, we can reinterpret our model as one in which more extreme policies result in greater scrutiny from the media and voters, and this generates a more precise signal for voters about the true impact of the policy.

Preliminaries

Before moving to the analysis of equilibrium policymaking in the two models, it is useful to establish some preliminary results that are applicable to both.

The Voter's Retention Rule

We begin by characterizing the voter's retention rule. In the last period, the officeholder has no reelection concerns and thus always implements his ideal policy. Therefore, if I is re-elected then $x_2^* = \hat{x}_I$, otherwise, $x_2^* = \hat{x}_C$. As is typical in models of electoral accountability, the voter faces a selection problem. Here, her problem is two-fold. The voter wants to elect the candidate whose ideal point provides her with the highest expected utility (given her posterior beliefs), but she also prefers the candidate with the greatest valence.

Recall that, in the fixed-valence model, θ_I and θ_C are publicly observed by all players at the beginning of the game. In the uncertain-valence model, the incumbent's valence is initially unknown but is observed prior to the election, while the challenger's valence remains unknown. Using the notation $v_C = \mathbb{E}[\theta_C]$ introduced earlier, we can express a general retention rule for the voter which holds across both versions of the model. In equilibrium, the voter strictly prefers to reelect the incumbent if:

$$\mathbb{E}_{\hat{x}_V} \left[-(\hat{x}_I - \hat{x}_V)^2 | x_1, \sigma \right] + \theta_I > \mathbb{E}_{\hat{x}_V} \left[-(\hat{x}_C - \hat{x}_V)^2 | x_1, \sigma \right] + v_C. \quad (4)$$

If the inequality is reversed then she always elects the challenger. If it holds with equality then the voter is indifferent and we assume she reelects the incumbent with probability $1/2$.⁸ Lemma

⁸This assumption makes the incumbent's problem continuous at $x = 0$ for all v_I . It is only consequential for a measure 0 set of parameters and does not otherwise affect our results.

1 characterizes the retention decision, recalling that $\hat{x}_C = -\hat{x}_I$. All proofs can be found in the Appendix.

Lemma 1. *Define μ as the posterior probability that $\hat{x}_V = 1$. The voter reelects the incumbent if:*

$$\mu > \frac{1}{2} - \frac{\theta_I - v_C}{8\hat{x}_I}.$$

Otherwise, the voter elects the challenger.

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

A key quantity is the valence v_I at which the voter is ex ante indifferent between the incumbent and the challenger, which we denote as \bar{v} . We characterize this value by replacing θ_I with v_I (the incumbent's ex-ante valence) in condition (4), and finding the v_I that solves the condition with equality. Specifically,

$$\bar{v} = v_C - 4\hat{x}_I(2\gamma - 1).$$

Recall that γ is the prior probability that $\hat{x}_V = 1$. Hence, if $v_I > \bar{v}$ then the incumbent is ex-ante preferred by the voter. When instead $v_I < \bar{v}$, the voter ex-ante prefers to replace him. Building on this discussion, we introduce the following terminology, which will be useful for characterizing our results.

Definition 1. *The incumbent is **leading** if $v_I > \bar{v}$ and **trailing** if $v_I < \bar{v}$.*

Policy Experimentation and Voter Learning

As highlighted by Lemma 1, the voter's retention decision depends crucially on her beliefs about the optimal policy. Here, we fully characterize the features of the voter's learning technology. Although the voter directly observes the incumbent's policy choice, x_1 , her inference problem is complicated because she receives only a noisy signal of her realized utility from this policy. In this setting, we show that the amount of voter learning is a function of the implemented policy.

Using Bayes' rule, if the incumbent chooses policy x_1 and this generates signal σ then $\mu(x_1, \sigma)$, the voter's posterior belief that $\hat{x}_V = 1$, is given by:

$$\mu(x_1, \sigma) = \frac{\gamma\phi(\sigma + (x_1 - 1)^2)}{\gamma\phi(\sigma + (x_1 - 1)^2) + (1 - \gamma)\phi(\sigma + (x_1 + 1)^2)}.$$

This highlights two crucial properties of the learning process. First, because the noise distribution satisfies the Monotone Likelihood Ratio Property, $\mu(x_1, \sigma)$ is increasing in σ when $x_1 > 0$, and

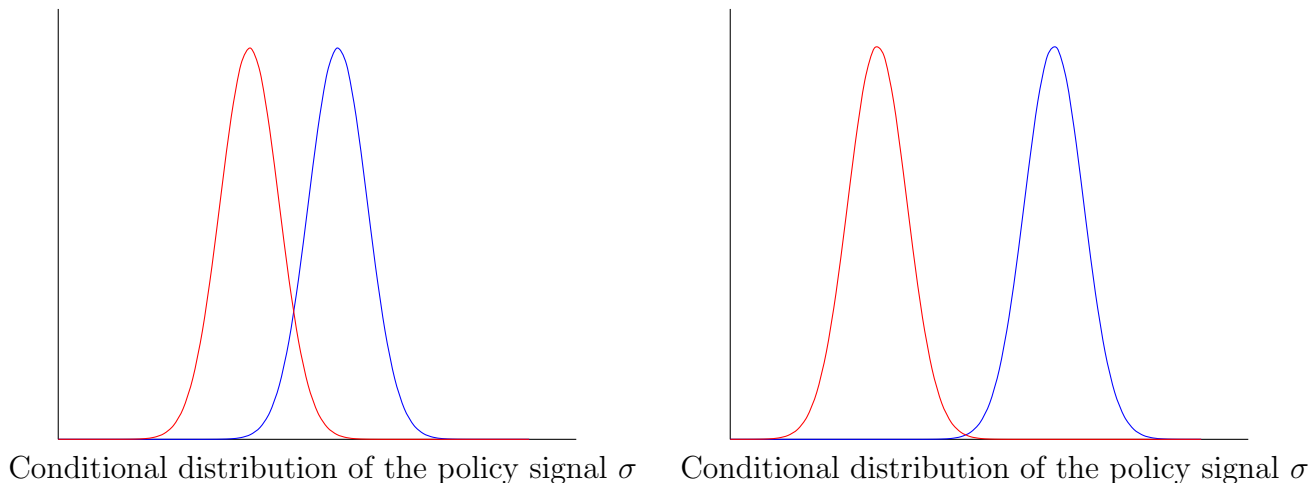
decreasing in σ otherwise. Thus, when a right-wing (left-wing) policy generates a higher signal σ , the voter believes it is more likely that her ideal policy is also right-wing (left-wing).

Second, even fixing the signal σ , the inferences that the voter draws depend on the implemented policy. It is easy to see that $|E[\sigma|\hat{x}_V = 1] - E[\sigma|\hat{x}_V = -1]| = |4x_1|$ increases if x_1 moves away from 0 in either direction. In other words, as x_1 becomes more extreme, the signal distributions conditional on the state \hat{x}_V move farther apart (see Figure 1). As a consequence, the voter is better able to filter out information from noise and draws a more precise inference.

Given these properties, we have that more extreme policies reduce the variance in the posterior distribution. At the extremes, if $x_1 = 0$ then the voter does not update at all from the signal, whereas if $x_1 \rightarrow \infty$ then the signal is perfectly informative. In the Appendix, we formalize this discussion and show that outcomes are more (Blackwell) informative as $|x_1|$ increases.

Substantively, suppose the incumbent implements an extreme left-wing economic policy. Should this platform produce a high signal, it likely aligns with the voter's optimal policy (i.e., matches the sign of \hat{x}_V). Conversely, a high signal from a moderate policy does not necessarily indicate alignment with the voter's interests, it could merely result from random economic fluctuations.

Figure 1: Policy and learning



Note: Figure 1 depicts the effect of moving policy away from 0 on the signal. In each graph, the red curve represents the signal distribution under $\hat{x}_V = -1$, and the blue curve the distribution under $\hat{x}_V = 1$. The left graph fixes a policy $x'_1 > 0$, and the right one a policy $x''_1 > x'_1$.

The Incumbent's Strategic Problem

We now turn to the incumbent's optimal policy choice. The voter's retention decision depends on her posterior beliefs on the policy dimension, as well as I 's valence and C 's expected valence.

As the previous section highlights, the amount of learning on the policy dimension is endogenous, with more extreme policies inducing more learning. Therefore, the incumbent's policy choice is a function of his own ideological preferences, \hat{x}_I , as well as his incentives to either prevent or facilitate voter learning about policy.

Let $P_{\theta_I}(x)$ be the incumbent's probability of winning if he chooses policy x , given valence θ_I . In the fixed-valence model, θ_I is known at the beginning of the game and remains fixed throughout. Thus, the incumbent's probability of winning is $P_{\theta_I}(x)$. Instead, in the uncertain-valence model, θ_I is initially unknown but gets revealed right before the election. Thus, from the perspective of the incumbent the probability of winning is $\pi_I P_1(x) + (1 - \pi_I) P_0(x)$. The following lemma establishes that the optimal policy choices, denoted x_f^* in the fixed-valence model and x_u^* in the uncertain-valence model, solve I 's first-order conditions.

Lemma 2. *In the fixed-valence model any equilibrium policy x_f^* must solve:*

$$2(\hat{x}_I - x) + (\beta + 4\hat{x}_I^2) \frac{\partial P_{\theta_I}}{\partial x} = 0.$$

In the uncertain-valence model any equilibrium policy x_u^ must solve:*

$$2(\hat{x}_I - x) + (\beta + 4\hat{x}_I^2) \left(\pi_I \frac{\partial P_1}{\partial x} + (1 - \pi_I) \frac{\partial P_0}{\partial x} \right) = 0.$$

Throughout, we focus on a selection of equilibrium such that x_f^* and x_u^* are differentiable in v_I for $v_I < \bar{v}$ and for $v_I > \bar{v}$.⁹ The first-order condition highlights that the implemented policy influences the incumbent's probability of being reelected (via the voter learning). Given the symmetry in our setup, any pair of policies x and $-x$ induces the same posterior distribution in expectation, therefore $P_{\theta_I}(x) = P_{\theta_I}(-x)$. This implies that, in equilibrium, a right-wing incumbent never implements a policy $x_1 < 0$. Therefore, if the second term in the first-order condition is always positive, so that more extreme policies increase his probability of winning in this range, then the incumbent implements a policy to the right of his ideological preferences, $x_1^* > \hat{x}_I$. Otherwise, the incumbent distorts his policy choice towards 0, $x_1^* < \hat{x}_I$. Regardless of the direction of the distortion, the intensity of the incumbent's strategic incentives to control information determines the magnitude to which policy moves away from \hat{x}_I .

In what follows, we often use the following terminology:

Definition 2. *The incumbent gambles if $x^* > \hat{x}_I$.*

⁹We note that the equilibrium policy in the fixed-valence case is discontinuous in v_I at $v_I = \bar{v}$, despite including noise in the voter's utility. This is because at $v_I = \bar{v}$ 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 $v_I \neq \bar{v}$. However, the incumbent's problem is well-behaved for $v_I > \bar{v}$ and $v_I \leq \bar{v}$, delivering existence of a differentiable x_u^* on each side of \bar{v} . Furthermore, numerical examples suggest that the equilibrium policy is unique.

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 and more extreme policies are more informative for the voter. The gambling terminology emphasizes that if we observe I implementing a policy more extreme than \hat{x}_I , then it is 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}_I and towards 0, then it is due to electoral incentives to prevent information generation relative to \hat{x}_I .

Policymaking in the Fixed-Valence Model

First, we consider the fixed-valence model, where both $v_I = \theta_I$ and $v_C = \theta_C$ are known at the start of the game to all players. Lemma 2 emphasizes that whether the incumbent chooses a policy more or less extreme than his ideal point depends on his electoral incentives to control information. Thus, our first step in analyzing the equilibrium policy choice is to characterize how I 's probability of winning changes as x moves away from 0 (i.e., the sign of $\frac{\partial P_{\theta_I}}{\partial x}$). Recall that $\bar{v} = v_C - 4\hat{x}_I(2\gamma - 1)$ is the value at which the voter is ex-ante indifferent between incumbent and challenger.

Lemma 3. *Assume $x > 0$.*

- (i) *If $v_I > \max\{v_C, \bar{v}\}$ then I 's probability of winning is decreasing in x .*
- (ii) *If $v_I < \min\{v_C, \bar{v}\}$ then I 's probability of winning is increasing in x .*
- (iii) *If $v_I \in (\min\{v_C, \bar{v}\}, \max\{v_C, \bar{v}\})$ then the probability of winning is non-monotonic in x :*
 - *If $\gamma < \frac{1}{2}$ then I 's probability of winning is single-peaked in x .*
 - *Instead, if $\gamma > \frac{1}{2}$ then I 's probability of winning is decreasing then increasing in x .*

Symmetric results hold for $x < 0$.

Parts (i) and (ii) of Lemma 3 are intuitive. When the incumbent is ex-ante leading, if there is no new information then I always wins the election. If I 's advantage is significant, then generating more policy-relevant information always hurts his reelection chances. Consequently, I 's ex-ante probability of winning is decreasing in the informativeness of the experiment, i.e., as x moves away from 0. In contrast, the incumbent always loses the election absent new information when he is ex-ante trailing. Thus, if he is sufficiently disadvantaged, then generating more information always increases his probability of winning, $\frac{\partial P_{\theta_I}}{\partial x} > 0$.

Point (iii) is more subtle, and it uncovers an unexpected non-monotonicity that arises as a consequence of the voter's valence concerns.

Consider an incumbent who is barely leading: $v_I \in (\bar{v}, v_C)$, which implies $\gamma > \frac{1}{2}$. In this case, I is ex ante leading because he is ahead on the policy dimension, $\gamma > \frac{1}{2}$, but he has lower valence than the challenger, $v_I < v_C$. On the one hand, such an incumbent is always damaged by the revelation of information, as he is guaranteed reelection if the voter learns nothing new. On the other, if information is produced then making the signal more informative is potentially beneficial. This is because a false negative — i.e., a bad outcome that occurs despite policy aligning with the state \hat{x}_V — severely damages I 's electoral prospects when he is close to the indifference threshold. Since the true state is likely to favor the incumbent, $\gamma > \frac{1}{2}$, the probability of a false negative decreases as outcomes become more informative. These two effects push in opposite directions, leading to the probability of winning being “single-dipped” in x for a barely leading incumbent.

If instead the incumbent is barely trailing, then the logic is reversed. Specifically, suppose $v_I \in (v_C, \bar{v})$, which implies $\gamma < \frac{1}{2}$. Hence, I is behind because the policy dimension ex ante favors the challenger, but he is stronger on the valence dimension, $v_I > v_C$. A trailing incumbent needs to reveal information to win reelection, but any learning on the policy dimension is likely to be unfavorable since $\gamma < 1/2$. However, because the incumbent is barely trailing, a favorable signal does not need to be very informative to move him ahead of the challenger. Thus, a barely trailing incumbent's probability of winning is maximized when there is a high probability of generating a false positive, which occurs for intermediate values of $|x|$.

Finally, notice that if valence does not matter ($v_I = v_C$) then these non-monotonic effects do not emerge. Absent heterogeneity in valence, for the incumbent to be barely leading requires γ to be very close to $\frac{1}{2}$. As such, the likelihood that a more informative policy avoids a false negative instead of revealing a true negative is low; hence, the incumbent's probability of winning is always decreasing in x . In contrast, when valence matters I can remain barely leading even if γ is relatively high, as long as $v_I < v_C$. Likewise, for the incumbent to be barely trailing when $v_I = v_C$ also implies γ is close to $1/2$. Therefore, the incentive to try and generate a false positive is muted, and I 's probability of winning is always increasing in x .

If I implemented policy based purely on reelection incentives to control information, Lemma 3 would be sufficient to characterize his optimal choice. In equilibrium, however, the policy choice also depends on I 's policy preferences. Absent reelection concerns, I would implement his ideal point. Proposition 1 characterizes the conditions under which we instead observe the incumbent *gamble* by implementing a policy more extreme than his ideal point.

Proposition 1. *There exists $\hat{v}_f \in (v_C, \bar{v}]$ such that the incumbent gambles if and only if $v_I < \hat{v}_f$. Furthermore, if $\gamma < \frac{1}{2}$ then $\hat{v}_f < \bar{v}$. Otherwise, if $\gamma \geq \frac{1}{2}$ then $\hat{v}_f = \bar{v}$.*

When the incumbent is leading ($v_I > \bar{v}$) his probability of winning is maximized when no new information is revealed. Hence, intuitively, he never gambles and he always chooses a policy more

centrist than \hat{x}_I . However, a trailing incumbent may gamble or not. Consistent with expectations from Lemma 3, incumbents who are significantly trailing, $v_I < \hat{v}_f$, always gamble to generate a highly informative signal. In contrast, if $v_I \in (\hat{v}_f, \bar{v})$ then I moderates and limits the informativeness of the signal relative to \hat{x}_I , despite needing new information to win reelection.

This result is a direct consequence of the non-monotonicity identified in Lemma 3. If the voter's optimal policy is unlikely to align with the incumbent's ideal point ($\gamma < \frac{1}{2}$) and I is only barely trailing, then the probability of winning is maximized at intermediate values of x . Consequently, policymaking by a barely trailing incumbent appears similar to a leading one. We observe I avoiding gambles, as he distorts his policy to a more moderate position.¹⁰ As such, Proposition 1 qualifies insights from work on gambling for resurrection (Downs and Rocke, 1994). The logic of gambling for resurrection suggests that reelection incentives should motivate a trailing incumbent to implement a policy that is riskier than he would choose otherwise. Instead, in our setting, we sometimes observe a trailing incumbent choose a policy less informative than his static optimum.

The previous result characterizes the *direction* in which the incumbent distorts policy away from his ideal point due to career concerns. The *magnitude* of this distortion depends on the intensity of the incumbent's incentives to control information. Our next proposition characterizes how the incumbent's valence influences this intensive margin, and thus the amount of experimentation that emerges in equilibrium.

Proposition 2. *There exists a cut-point $\underline{v} \geq 0$, with $\underline{v} < \bar{v}$, such that, if $v_I \in (\underline{v}, \bar{v})$ then the equilibrium policy in the fixed-valence model, x_f^* , is decreasing in v_I . Otherwise, x_f^* is increasing in v_I .*

Changing v_I has two effects on the incumbent's incentives. First, increasing v_I makes the voter more willing to reelect the incumbent, which implies that I can survive more negative policy information. Second, moving v_I towards \bar{v} creates a more competitive electoral environment, which makes policy information more relevant for the voter's decision. Importantly, these two effects imply that increasing v_I can induce the incumbent to implement either a more moderate or more extreme policy, depending on the degree of his initial advantage (or disadvantage).

If the incumbent is leading, $v_I > \bar{v}$, then these two effects always go in the same direction. Increasing v_I away from \bar{v} increases how much negative policy information the incumbent can reveal while still winning reelection *and* it makes policy information less relevant. Thus, both effects weaken I 's incentives to control information when v_I increases, which pushes x_f^* to be more extreme, moving it towards the incumbent's ideal point.

In contrast, if the incumbent is trailing, $v_I < \bar{v}$, then these two effects compete. Increasing v_I towards \bar{v} increases the salience of the policy dimension (making information control more relevant),

¹⁰We note that when $\gamma > \frac{1}{2}$ the non-monotonicity from Lemma 3 does not cause a barely leading incumbent to gamble because his probability of winning remains maximized at $x = 0$.

but also makes the voter more lenient towards I (making information control less relevant). When v_I is nearer to 0, the former effect can dominate, intensifying the incumbent's incentives to experiment. Conversely, when v_I approaches \bar{v} , the latter effect prevails, dampening incentives to control information. Therefore, as v_I increases, initially x_f^* can move towards more extreme positions; however, as v_I gets closer to the indifference threshold, it shifts the equilibrium policy to more moderate stances closer to \hat{x}_I . Proposition 3 provides conditions under which this non-monotonicity must emerge.

Proposition 3. *If \hat{x}_I is sufficiently small then $\underline{v} > 0$.*


Whether x_f^* is non-monotonic in v_I when the incumbent is trailing depends on the importance of the valence dimension relative to the policy dimension for the voter. This is determined by \hat{x}_I , which captures the degree of polarization between the candidates. When \hat{x}_I is low the candidates deliver similar policy platforms in the second period, thus, electoral outcomes are mostly determined by valence. In this case, when v_I close to 0 the incumbent's probability of winning is always 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_I , and $\underline{v} > 0$. However, when \hat{x}_I is larger we may have $\underline{v} = 0$, in which case x_f^* is strictly decreasing in v_I for all $v_I < \bar{v}_I$.¹¹

Figure 2 pulls together Propositions 1, 2, and 3 to depict how the equilibrium policy changes as a function of v_I . In the fixed-valence model, the incumbent's valence has a non-monotonic effect on policy extremism. In particular, if $\underline{v} > 0$, then the maximum amount of policy experimentation is by incumbents who are trailing, but face only a moderate disadvantage. Furthermore, the least amount of experimentation is by incumbents who are just barely leading.¹² These results highlight the importance of considering ideological preferences as well as reelection motives for understanding incentives to engage in policy experimentation.

¹¹That $\underline{v} = 0$ can emerge for some parameters is straightforward to verify numerically.

¹²Indeed, such incumbents engage in almost no experimentation. In the Appendix, under our assumption on β we show that $\lim_{v_I \rightarrow \bar{v}^+} x_f^* = 0$.

Figure 2: Policymaking in the fixed-valence model



Note: Figure 2 depicts the incumbent's equilibrium policy choice, x_f^* , as a function of his valence, in the fixed-valence model.

Policymaking in the Uncertain-Valence Model

Next we analyze the uncertain-valence model. In this version of the model, the true value of θ_I is revealed to the voter prior to the election but after I chooses x_1 . Here, $v_I = \pi_I$ and it denotes the incumbent's expected valence. Proposition 4 begins by establishing when reelection incentives encourage the incumbent to gamble.


Proposition 4. *There exists \hat{v}_u such that the incumbent gambles if and only if $v_I < \hat{v}_u$. Moreover, generically $\hat{v}_u \neq \bar{v}$.*

Lemma 3 immediately yields that the incumbent's probability of winning conditional on a favorable valence realization, $\theta_I = 1$, is decreasing in $|x|$. Conversely, conditional on $\theta_I = 0$, more extreme policies are electorally beneficial. However, the incumbent must choose policy before the value of θ_I is realized, therefore, his policy choice depends on his probability of being high valence. When v_I is high the incumbent is likely to secure an electoral advantage; thus, he wants to choose a policy more moderate than his ideal point to limit the amount of additional information revealed. In contrast, if v_I is low then the incumbent expects to be significantly behind after his valence is revealed, and thus he gambles to try and generate a very favorable signal on the policy dimension.

However, similar to Proposition 1, whether I is leading or trailing electorally ex ante does *not* determine if he gambles on policy, $\hat{v}_u \neq \bar{v}$. Because θ_I is revealed before the election, how the voter views I 's valence ex ante never plays a role in her retention decision. This further highlights that

valence considerations can significantly alter our understanding of how electoral incentives impact strategic policy gambles.

Figure 3: Policymaking in the uncertain-valence model



Note: Figure 3 depicts equilibrium policy in the uncertain-valence model, x_u^* , as a function of the incumbent's ex-ante valence, v_I .

We conclude the analysis by studying how the exact policy implemented in equilibrium, x_u^* , varies with the incumbent's ex-ante valence.

Proposition 5. *The equilibrium policy in the uncertain-valence model, x_u^* , is decreasing in v_I .*

In sharp contrast to the results of the fixed-valence model, the equilibrium amount of experimentation is *always* decreasing in the incumbent's ex-ante valence. As previously mentioned, in the uncertain-valence model v_I does not play a direct role in the voter's retention decision. This substantially alters how I balances his policy preferences against his incentive to control information compared to the fixed-valence model. Increasing v_I increases the probability that I obtains an ex-post electoral advantage, and thus increases the probability he benefits from limiting the informativeness of the policy signal. However, a higher v_I does not make the voter more lenient towards the incumbent on election day, since the voter's new information supersedes her prior beliefs.¹³ Thus, conditional on the valence realization, a higher v_I does not allow the incumbent to counter the negative downsides from a more informative policy. Consequently, a higher v_I only influences the incumbent's policy by amplifying his incentives to protect an increasingly likely ex-post advantage. In turn, this induces him to implement a more moderate policy.

¹³Necessary 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.

Endogenous Valence Information

Thus far, we have considered two versions of the model: one where valence is commonly known and exogenously fixed, and one where it is uncertain and exogenously revealed. Here, we study a version of the model where $\theta_i \in \{0, 1\}$ for $i \in \{I, C\}$ is ex-ante unknown to all players, with common prior belief $Pr(\theta_i = 1) = \pi_i$, as in the uncertain-valence model. Now, however, I can take an action that reveals his valence. We will refer to this action as a valence test. We assume that if I takes a valence test then θ_I is revealed before the election but after he chooses x_1 , as in the uncertain-valence game. Otherwise, if I chooses not to generate a valence test then θ_I is not revealed, and the voter maintains her prior belief π_I . In other words, I can choose whether to play the fixed-valence game or the uncertain-valence game. For example, the officeholder may purposefully generate a crisis that tests his unknown competence, such as entering a war (as in Downs and Rocke, 1994), and the consequences of the crisis may not materialize before he has to set policy.

We assume the office rent β is sufficiently large that the incumbent's dynamic reelection incentives dominate ideological considerations. This allows us to focus on the electoral incentives to generate valence information, and implies that I 's equilibrium policy is close to the one that maximizes his probability of winning. Finally, we also have that I 's expected valence is $v_I = \pi_I$ in this setting.

Proposition 6. *Assume office benefit is sufficiently large. If $v_I \in [\gamma, \bar{v}]$ then there exist parameter values such that the incumbent generates a valence test in equilibrium. Otherwise, if $v_I \notin [\gamma, \bar{v}]$ then I never generates a valence test.*

In this model, the incumbent can choose to gamble on both policy and valence. Suppose $v_I > \bar{v}$. In this case, the incumbent is ex-ante leading and wins reelection absent any new information. Therefore, I has no incentive to gamble, whether on the policy or the crisis dimension. In contrast, when $v_I < \bar{v}$, the trailing incumbent needs the voter to update positively on at least one of these dimensions to win the election. If $v_I < \gamma$ then gambling on the policy dimension is more likely to succeed than gambling on valence. The incumbent therefore never generates an endogenous valence test, because this maximizes the probability that new policy-relevant information is enough to push him above the retention threshold. If instead $v_I > \gamma$, then the incumbent is more likely to generate a positive valence shock than to generate a signal that convinces the voter \hat{x}_V is right-wing (and thus aligned with \hat{x}_I). Therefore, he sometimes finds it optimal to generate a valence test to improve his retention chances.

Application: Crises as Valence-Revealing Events

As briefly mentioned above, one interpretation of the uncertain-valence model is that information revelation on θ_I occurs as the result of a crisis. If the country is hit by an economic recession, experiences a natural disaster, or suffers a military attack, then the incumbent's management of the crisis will reveal information about his ability. In the previous section, we explored the officeholder's incentives to generate such a crisis. However, even absent these incentives, the country may nonetheless be hit by an *exogenous* crisis that acts as a valence-revealing shock. In this section, we ask: How do the incumbent's incentives to experiment on the ideological dimension change in times of crisis?

As described above, within our framework, the fixed-valence model 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 valence $v_I = \pi_I$. Conversely, the uncertain-valence model represents 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.


Proposition 7 compares policymaking in the two versions of the model, and characterizes conditions under which a crisis induces more or less experimentation.

Proposition 7. *If $v_I > \bar{v}$ then the incumbent enacts a more extreme policy during times of crisis than during normal times. If $v_I < \underline{v}$ then the incumbent implements a more moderate policy during times of crisis than during normal times. If $v_I \in (\underline{v}, \bar{v})$ then the crisis can lead to more or less extreme policies.*

The crisis can alter the extent of I 's existing incentives to control information, that is, it can have a *quantitative* effect on policymaking. Furthermore, the crisis can have a *qualitative* effect on policymaking and incentivize the incumbent to switch from gambling to not, or vice versa. Proposition 7 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. This conditional effect is due to the crisis having a quantitative impact on policymaking incentives when the incumbent is very far ahead or behind, and a qualitative impact when elections are more competitive.

To understand these two effects, consider an incumbent who is electorally leading ex-ante. During normal times, such an incumbent distorts policy away from \hat{x}_I and towards a more moderate position, to prevent information generation and protect his electoral advantage. When this incumbent is hit by a crisis it changes his policymaking calculus. He may score a success on the crisis dimension and further increase his electoral lead, which makes voter learning on the policy dimension less electorally relevant. Alternatively, he may fail to solve the crisis and lose his initial

Figure 4: Crisis vs. no crisis



Note: Figure 4 compares equilibrium policy in normal times, x_f^* (depicted in red) against the equilibrium policy in times of crisis, x_u^* (depicted in blue), as a function of the incumbent's expected competence, v_I .

advantage. In this case, voter learning on policy becomes electorally valuable, as a failure turns a leading incumbent into a heavily trailing one. Although both forces push a leading incumbent to choose a more extreme policy, which force dominates determines whether the crisis has a quantitative or qualitative effect on policymaking.

For a leading incumbent who is very likely to be competent, the crisis has a quantitative effect on policy. In this case, the incumbent anticipates he will likely solve the crisis. Thus, he is still incentivized to choose a policy more moderate than his ideal point to prevent voter learning. However, such an incumbent also knows that he can afford more risk in policymaking because a success on the crisis dimension makes his retention probability less elastic to information on the policy dimension. Therefore, the crisis weakens his incentives to control information, and causes him to enact a policy more extreme (closer to \hat{x}_I) than during normal times.

Similarly, the crisis has a quantitative effect when the incumbent is very far behind, although the effect goes in the opposite direction. During normal times, severely trailing incumbents always gamble and choose an extreme policy that facilitates voter learning. Suppose instead the country is hit by a crisis. The incumbent continues to gamble in hopes of securing a policy success because v_I is very low and he anticipates not solving the crisis. However, as above, the crisis weakens the incumbent's strategic incentives to experiment because failure on the crisis makes policy outcomes less electorally relevant. Thus, the incumbent implements a more moderate policy, one closer to \hat{x}_I , during a crisis than in normal times.

Finally, suppose the election is ex-ante close. Here, the crisis can have a qualitative effect on

policymaking: it alters the nature of the incumbent’s incentives to control information. Depending on the parameter values, this qualitative effect can mean that crises induce more or less experimentation.

Suppose that $v_I \in [\hat{v}_u, \hat{v}_f]$. In this case, the qualitative effect produces more extreme policies. Such an incumbent is ex-ante electorally trailing and gambles on policy during normal times. However, it is not too unlikely that he can prove he is competent if a crisis hits. As such, in a crisis, this marginally trailing incumbent finds it optimal to behave as if he was electorally leading, and avoids gambling. Here, the qualitative effect implies that the crisis induces less experimentation. A symmetric reasoning applies when $v_I \in [\hat{v}_f, \hat{v}_u]$. This incumbent avoids policy gambles in a period of business as usual. However, in a crisis he anticipates that he is likely to fail and worsen his electoral prospects. In turn, this induces him to gamble on policy in hopes of scoring a success on the ideological dimension.¹⁴ In the Appendix, we provide numerical examples demonstrating that both cases, $\hat{v}_f < \hat{v}_u$ or $\hat{v}_f > \hat{v}_u$, can occur, depending on the parameters.

So far, we have described policy experiments in ideological terms, where choices further from the center are more informative. An alternative interpretation is to assume that 0 represents the status quo policy. Under this interpretation, our results suggest that more radical reforms generate more information, compared to policies that stay closer to the status quo. In this perspective, our model complements existing work that studies how crises impact policy reform.

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

¹⁴We note a potential ambiguity. There may exist $v_I \in [\underline{v}, \bar{v}]$ such that the incumbent experiences a quantitative rather than qualitative effect. In this case, the crisis may have a non-monotonic effect on reform. However, if the equilibrium policies intersect for at most one value of $v_I < \bar{v}$ then this type of non-monotonicity cannot arise. In particular, there is a unique cutoff in v_I above which the crisis always leads to more reform, and below which it leads to less reform. An example is depicted in Figure 4, where the crisis and no-crisis policies never intersect at interior values. Moreover, numerical simulations support the claim that the policies can only intersect at most once. However, due to the nuanced competing effects of v_I on the equilibrium policies it is difficult to obtain a general analytical characterization of this region.

implies that researchers may even recover a zero effect when averaging across different values of v_I . 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 7 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.

Extension: Asymmetric Information

Before concluding, we relax the assumption of symmetric uncertainty and 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 \hat{x}_V . 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.

At the beginning of the game nature draws the type of each politician, $\theta_i \in \{0, 1\}$, according to the commonly known distributions $Pr(\theta_i = 1) = \pi_i \in (0, 1)$. Next, the incumbent observes a private signal $s_I \in \{0, 1\}$, where $Pr(s_I = 1 | \theta_I = 1) = Pr(s_I = 0 | \theta_I = 0) \in (1/2, 1)$. The game then proceeds as before. In the uncertain-valence model, the incumbent’s type is fully revealed prior to the election. Instead, in the fixed-valence model, the voter receives no exogenous information about θ_I . As in the previous section, we maintain the definition of $v_I = \pi_I$ as the prior probability that the incumbent is a high type. Our solution concept is perfect Bayesian equilibrium, henceforth “equilibrium”.

After observing the signal, the incumbent updates his beliefs about his own valence according to Bayes’ rule. Let ψ_{s_I} be the incumbent’s (interim) posterior belief that $\theta_I = 1$ conditional on the realization of his private signal. Thus, $0 < \psi_0 < \pi_I < \psi_1 < 1$.

We show that, across the two models, information asymmetries do not change our qualitative results from the baseline setup, in particular, the comparative statics on v_I .

The Fixed-valence Model

We define $x_f^b(v_I)$ as the equilibrium policy in the baseline fixed-valence model with symmetric uncertainty. In the asymmetric information setting, we let $x_f^a(s_I)$ denote the equilibrium policy choice of the incumbent after observing the signal s_I . 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 $x_f^b(v_I)$ following either signal $s_I \in \{0, 1\}$. That is, both types of the incumbent choose the equilibrium policy from the baseline model without asymmetric information.

Lemma 4. *There always exists an equilibrium where the incumbent chooses $x_f^b(v_I)$ following either signal, $x_f^a(0) = x_f^a(1) = x_f^b(v_I)$.*

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 $s_I = 0$. 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 5. *In any equilibrium, both types of the incumbent are always 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 $s_I = 0$ 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 $\theta_I = \psi_0$. In contrast, in the pooling equilibrium described earlier both types receive the payoff from the complete information game with $\theta_I = \pi_I > \psi_0$. From here, our next result follows from a standard envelope argument.

Proposition 8. *Among all equilibria, the equilibrium where $x_f^a(0) = x_f^a(1) = x_f^b(v_I)$ maximizes the incumbent's expected utility under each signal.*

Proposition 8 shows 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_I . Thus, the unique equilibrium policy choice surviving this refinement is the same as the complete information fixed-valence model.

The Uncertain-valence Model

Here, we define $x_u^b(v_I)$ as the equilibrium policy in the baseline uncertain-valence model with symmetric information. We let $x_u^a(s_I)$ denote the equilibrium policy choice of the incumbent after observing the signal s_I in the asymmetric information setting with uncertain valence.

Lemma 6. *In every equilibrium, the incumbent chooses different policies following each signal, $x_u^a(0) \neq x_u^a(1)$. Furthermore, $x_u^a(s_I = 0) = x_u^b(v_I = \psi_0)$ and $x_u^a(s_I = 1) = x_u^b(v_I = \psi_1)$.*


The incumbent's 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_I} . As a consequence, the *expected* equilibrium policy is decreasing in $v_I = \pi_I$. As in the symmetric uncertainty baseline, officeholders with higher valence in expectation enact more moderate policies, all else equal (as depicted in Figure 5).

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 signalled by the incumbent's actions (formally, the public signal revealed before the election is more informative than the incumbent's private signal).

Conclusion

When it comes to deciding which candidate to support, two elements crucially influence voters' choice. First, their past experiences with policy outcomes. Second, their perception of the relative

Figure 5: Policymaking with asymmetric information



Note: Figure 5 depicts the equilibrium policy following the $s_\theta = 0$ signal (upper dashed blue line), the one following the $s_\theta = 1$ signal (lower dashed blue line), and the expected policy (solid blue line).

valences of each candidate. While the former is endogenous to the officeholder’s choice, the latter is often exogenously determined. Here, we study the interaction between these two elements and ask: How does the incumbent’s valence influence his incentives to experiment with risky policies in order to control voters’ learning?

Our key results are as follows. First, we show that the voter’s preferences over valence may fundamentally alter the *nature* of the incumbent’s strategic incentives. In turn, this may induce electorally trailing incumbents to behave as if they were leading and avoid policy gambles, and vice versa. Second, we show that the effect of the incumbent’s valence on the *intensity* of his incentives to control information varies significantly depending on whether we consider valence as a fixed characteristic that is known ex-ante (e.g., the candidates’ experience), or as a shock that realizes right before the election (e.g., the voters’ perception of the candidates’ relative charisma or competence). While in the fixed-valence model the effect is non-monotonic (and experimentation can be maximized at intermediate values), in the uncertain-valence setting the amount of experimentation decreases monotonically as the incumbent’s expected valence increases. Finally, conceptualizing crises as valence-revealing events, we compare the two versions of the model (with fixed and uncertain valence) to study their effect on policy experimentation. We find that crises prompt bolder policy experiments when the incumbent’s expected valence is high, but lead to more cautious reforms from those with low valence.

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.
- Aragonès, Enriqueta and Dimitrios Xefteris. 2017. "Voters' private valuation of candidates' quality." *Journal of Public Economics* 156:121–130.
- 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.
- Bils, Peter. 2023. "Overreacting and Posturing: How Accountability and Ideology Shape Executive Policies." *Quarterly Journal of Political Science* 18(2):153–182.
- 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.
- 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.
- Hörner, Johannes and Andrzej Skrzypacz. 2017. "Learning, experimentation and information design." *Advances in Economics and Econometrics* 1:63–98.
- Izzo, Federica. 2022. "Who runs when?" *Working Paper* .
URL: https://www.federicaizzo.com/pdf/DWG_JOP.pdf
- Izzo, Federica. 2023. "Ideology for the Future." *American Political Science Review* 117(3):1089–1104.
- 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.
- 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.
- 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.
- 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.
- Yu, Tinghua. 2022. "Accountability and learning with motivated agents." *Journal of Theoretical Politics* 34(2):313–329.

A Proofs for Baseline Model

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_t 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 Friedenber (2017) applies, and delivers that outcomes are more Blackwell informative as x moves away from 0 in either direction. \square

Lemma 1. *Define μ as the posterior probability that $\hat{x}_V = 1$. The voter reelects the incumbent if:*

$$\mu > \frac{1}{2} - \frac{\theta_I - v_C}{8\hat{x}_I}.$$

Otherwise, the voter elects the challenger.

Proof. The voter's expected utility from re-electing the incumbent is greater than her utility from electing the challenger if

$$\begin{aligned} & -\mu(\hat{x}_I - 1)^2 - (1 - \mu)(\hat{x}_I + 1)^2 + \theta_I \\ & \geq -\mu(\hat{x}_C - 1)^2 - (1 - \mu)(\hat{x}_C + 1)^2 + v_C. \end{aligned}$$

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

$$\mu \geq \frac{1}{2} - \frac{\theta_I - v_C}{8\hat{x}_I}$$

\square

Definition 3. *Let $\bar{\mu}_{\theta_I} = \frac{1}{2} - \frac{\theta_I - v_C}{8\hat{x}_I}$ and $\lambda_{\theta_I} = \ln\left(\frac{(1-\gamma)\bar{\mu}_{\theta_I}}{\gamma(1-\bar{\mu}_{\theta_I})}\right)$.*

Lemma (A2). *Given valence θ_I , the probability of reelection for the incumbent is*

$$P_{\theta_I}(x_1) = \gamma\left(1 - \Phi\left(\frac{\lambda_{\theta_I}}{4|x_1|} - 2|x_1|\right)\right) + (1 - \gamma)\left(1 - \Phi\left(\frac{\lambda_{\theta_I}}{4|x_1|} + 2|x_1|\right)\right). \quad (5)$$

Proof. By Bayes rule we have

$$\mu(x_1, \sigma) = \frac{\gamma\phi(\sigma + (x_1 - 1)^2)}{\gamma\phi(\sigma + (x_1 - 1)^2) + (1 - \gamma)\phi(\sigma + (x_1 + 1)^2)}.$$

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

$$Pr \left(\frac{\gamma\phi(\sigma + (1 - x_1)^2)}{\gamma\phi(\sigma + (1 - x_1)^2) + (1 - \gamma)\phi(\sigma + (-1 - x_1)^2)} > \bar{\mu}_{\theta_I} \right), \quad (6)$$

where ϕ is the PDF of the standard normal distribution.

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

$$\begin{aligned} & \gamma \cdot Pr \left(\frac{\gamma\phi(-(\hat{x}_I - 1)^2 + \epsilon + (x_1 - 1)^2)}{\gamma\phi(-(\hat{x}_I - 1)^2 + \epsilon + (x_1 - 1)^2) + (1 - \gamma)\phi(-(\hat{x}_I + 1)^2 + \epsilon + (x_1 + 1)^2)} > \bar{\mu}_{\theta_I} \right) \\ & + (1 - \gamma) \cdot Pr \left(\frac{\gamma\phi(-(\hat{x}_I + 1)^2 + \epsilon + (x_1 - 1)^2)}{\gamma\phi(-(\hat{x}_I + 1)^2 + \epsilon + (x_1 - 1)^2) + (1 - \gamma)\phi(-(\hat{x}_I + 1)^2 + \epsilon + (x_1 + 1)^2)} > \bar{\mu}_{\theta_I} \right). \end{aligned} \quad (7)$$

Equation 7 further reduces to

$$\gamma \cdot Pr \left(\frac{\gamma\phi(\epsilon)}{\gamma\phi(\epsilon) + (1 - \gamma)\phi(4x_1 + \epsilon)} > \bar{\mu}_{\theta_I} \right) + (1 - \gamma) \cdot Pr \left(\frac{\gamma\phi(-4x_1 + \epsilon)}{\gamma\phi(-4x_1 + \epsilon) + (1 - \gamma)\phi(\epsilon)} > \bar{\mu}_{\theta_I} \right), \quad (8)$$

and we can rewrite this probability as

$$\gamma \cdot Pr \left(e^{-\frac{\epsilon^2}{2} + \frac{(4x_1 + \epsilon)^2}{2}} > \frac{\bar{\mu}_{\theta_I}(1 - \gamma)}{\gamma(1 - \bar{\mu}_{\theta_I})} \right) + (1 - \gamma) \cdot Pr \left(e^{-\frac{(-4x_1 + \epsilon)^2}{2} + \frac{\epsilon^2}{2}} > \frac{\bar{\mu}_{\theta_I}(1 - \gamma)}{\gamma(1 - \bar{\mu}_{\theta_I})} \right). \quad (9)$$

Suppose that $x_1 > 0$. After rearranging and applying a logarithmic transformation, the above obliquely reduces to

$$\gamma \cdot Pr(\epsilon > \frac{\lambda_{\theta_I}}{4x_1} - 2x_1) + (1 - \gamma) \cdot Pr(\epsilon > \frac{\lambda_{\theta_I}}{4x_1} + 2x_1), \quad (10)$$

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

Definition 4. Let $\Delta_{\theta_I}^- = \frac{\lambda_{\theta_I}}{4x} - 2x$ and $\Delta_{\theta_I}^+ = \frac{\lambda_{\theta_I}}{4x} + 2x$.

Lemma (A3).

1. $\frac{\partial \Delta_{\theta_I}^+}{\partial x} = \frac{-\lambda_{\theta_I}}{4x^2} + 2 = -\frac{1}{x}\Delta_{\theta_I}^-$, and
2. $\frac{\partial \Delta_{\theta_I}^-}{\partial x} = \frac{-\lambda_{\theta_I}}{4x^2} - 2 = -\frac{1}{x}\Delta_{\theta_I}^+$.

Proof. Follows immediately by differentiating. □

Definition 5. Let $\Gamma_{\theta_I}(x) = \gamma\Delta_{\theta_I}^+\phi(\Delta_{\theta_I}^-) + (1 - \gamma)\Delta_{\theta_I}^-\phi(\Delta_{\theta_I}^+)$ and $\Omega_{\theta_I}(x) = \gamma\phi(\Delta_{\theta_I}^-) + (1 - \gamma)\phi(\Delta_{\theta_I}^+)$.

Lemma (A4). We have $\frac{\partial P_{\theta_I}}{\partial x} = \frac{1}{x}\Gamma_{\theta_I}(x)$ and $\frac{\partial^2 P_{\theta_I}}{\partial x^2} = -\frac{1}{x^2}\Gamma_{\theta_I}(x) + \frac{1}{x}\frac{\partial \Gamma_{\theta_I}}{\partial x}$, where $\frac{\partial \Gamma_{\theta_I}}{\partial x} = \frac{1}{x}\left[\Delta_{\theta_I}^- \Delta_{\theta_I}^+ \Gamma_{\theta_I}(x) - \Omega_{\theta_I}(x)\right]$.

Proof. Taking the derivative of $P_{\theta_I}(x)$ with respect to x yields:

$$\begin{aligned}\frac{\partial P_{\theta_I}}{\partial x} &= \gamma \phi\left(\frac{\lambda_{\theta_I}}{4x} - 2x\right)\left[\frac{\lambda_{\theta_I}}{4x^2} + 2\right] + (1 - \gamma)\phi\left(\frac{\lambda_{\theta_I}}{4x} + 2x\right)\left[\frac{\lambda_{\theta_I}}{4x^2} - 2\right] \\ &= \frac{1}{x}\Gamma_{\theta_I}(x).\end{aligned}\tag{11}$$

Next, we differentiate Γ_{θ_I} and obtain:

$$\begin{aligned}\frac{\partial \Gamma_{\theta_I}}{\partial x} &= \gamma \frac{\partial \Delta_{\theta_I}^+}{\partial x} \phi(\Delta_{\theta_I}^-) - \gamma \Delta_{\theta_I}^+ \frac{\partial \Delta_{\theta_I}^-}{\partial x} \phi(\Delta_{\theta_I}^-) + (1 - \gamma) \frac{\partial \Delta_{\theta_I}^-}{\partial x} \phi(\Delta_{\theta_I}^+) - (1 - \gamma) \Delta_{\theta_I}^- \frac{\partial \Delta_{\theta_I}^+}{\partial x} \phi(\Delta_{\theta_I}^+) \\ &= \frac{1}{x} \gamma \phi(\Delta_{\theta_I}^-) (\Delta_{\theta_I}^- (\Delta_{\theta_I}^+)^2 - 1) + \frac{1}{x} (1 - \gamma) \phi(\Delta_{\theta_I}^+) (\Delta_{\theta_I}^+ (\Delta_{\theta_I}^-)^2 - 1) \\ &= \frac{1}{x} \left[\Delta_{\theta_I}^- \Delta_{\theta_I}^+ \Gamma_{\theta_I}(x) - \Omega_{\theta_I}(x) \right].\end{aligned}$$

The second derivative of $P_{\theta_I}(x)$ then follows straightforwardly by differentiating. \square

Lemma 2. In the fixed-valence model any equilibrium policy x_f^* must solve:

$$2(\hat{x}_I - x) + (\beta + 4\hat{x}_I^2) \frac{\partial P_{\theta_I}}{\partial x} = 0.$$

In the uncertain-valence model any equilibrium policy x_u^* must solve:

$$2(\hat{x}_I - x) + (\beta + 4\hat{x}_I^2) \left(\pi_I \frac{\partial P_1}{\partial x} + (1 - \pi_I) \frac{\partial P_0}{\partial x} \right) = 0.$$

Proof. We show that any equilibrium policy x_f^* in the fixed-valence model 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, $\beta + 4\hat{x}_I^2 < \infty$ implies $x_f^* < \infty$. Hence, recalling that $x_f^* \geq 0$, a maximizer exists.

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

A similar argument yields the result for x_u^* in the uncertain-valence model. \square

Lemma 3. Assume $x > 0$.

(i) If $v_I > \max\{v_C, \bar{v}\}$ then I 's probability of winning is decreasing in x .

(ii) If $v_I < \min\{v_C, \bar{v}\}$ then I 's probability of winning is increasing in x .

(iii) If $v_I \in (\min\{v_C, \bar{v}\}, \max\{v_C, \bar{v}\})$ then the probability of winning is non-monotonic in x :

– If $\gamma < \frac{1}{2}$ then I 's probability of winning is single-peaked in x .

– Instead, if $\gamma > \frac{1}{2}$ then I 's probability of winning is decreasing then increasing in x .

Symmetric results hold for $x < 0$.

Proof. We break our analysis into cases depending on the value of $\theta_I = v_I$. Recall that $\frac{\partial P_{\theta_I}}{\partial x}$ is negative if and only if equation (11) is negative.

Case 1: $\theta_I > \max\{v_C, \bar{v}\}$. If $\theta_I > \bar{v}$ then $\lambda_{\theta_I} < 0$. Hence, (11) is always negative if $\frac{\lambda_{\theta_I}}{2x^2} + 1 < 0$, which holds for all $x \in [0, \sqrt{\frac{-\lambda_{\theta_I}}{2}}]$.

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

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

Which we rewrite as:

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

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

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

which holds if and only if:

$$\frac{\bar{\mu}_{\theta_I}(1 - \gamma)}{\gamma(1 - \bar{\mu}_{\theta_I})} < \frac{1 - \gamma}{\gamma} \frac{-\frac{\lambda_{\theta_I}}{8x^2} + 1}{\frac{\lambda_{\theta_I}}{8x^2} + 1}. \quad (15)$$

This condition further simplifies to:

$$2\bar{\mu}_{\theta_I} + \frac{\lambda_{\theta_I}}{8x^2} < 1. \quad (16)$$

Recall that $\lambda_{\theta_I} < 0$ by $\theta_I > \bar{v}$. Thus, a sufficient condition for (16) to always hold is that $\bar{\mu}_{\theta_I} < \frac{1}{2}$. Expanding, $\bar{\mu}_{\theta_I} = \frac{1}{2} - \frac{\theta_I - v_C}{8\hat{x}_I} < \frac{1}{2}$ by assumption that $\theta_I > v_C$, which completes the argument.

Case 2: $\theta_I < \min\{v_C, \bar{v}\}$. By $\theta_I < \bar{v}$ we have $\lambda_{\theta_I} > 0$. Thus, the same rearrangement of (11) as in the previous case yields that $\frac{\partial P_{\theta_I}}{\partial x} > 0$ if and only if:

$$2\bar{\mu}_{\theta_I} + \frac{\lambda_{\theta_I}}{2x^2} > 1. \quad (17)$$

By assumption $\theta_I < v_C$, hence, $\bar{\mu}_{\theta_I} > \frac{1}{2}$. Together with $\lambda_{\theta_I} > 0$ this yields that 17 always holds.

Case 3: $\theta_I \in (\min\{v_C, \bar{v}\}, \max\{v_C, \bar{v}\})$. By definition of \bar{v} , $v_C < \bar{v}$ if and only if $v_C < v_C - 4\hat{x}_I(2\gamma - 1) \Leftrightarrow \gamma < \frac{1}{2}$. To prove the result, we further break the argument into two more cases, depending on the prior belief γ .

1. $\gamma < \frac{1}{2}$: In this case, $v_C < \bar{v}$, hence, $\theta_I \in (v_C, \bar{v})$. By $\theta_I < \bar{v}$ we have $\lambda_{\theta_I} > 0$. Therefore, $\frac{\partial P_{\theta_I}}{\partial x} < 0$ 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, $\lim_{x \rightarrow 0} \text{LHS (16)} = \infty > 1$, because $\lambda_{\theta_I} > 0$. Second $\lim_{x \rightarrow \infty} \text{LHS (16)} = 2\bar{\mu}_{\theta_I} < 1$, where the inequality holds because $\bar{\mu}_{\theta_I} < \frac{1}{2}$ by $\theta_I > v_C$. To complete the argument, notice that LHS (16) is clearly strictly decreasing in x because $\lambda_{\theta_I} > 0$. Furthermore, this implies that I 's probability of winning is maximized at the unique $x > 0$ that solves (16) at equality.
2. $\gamma > \frac{1}{2}$: In this case, $v_C > \bar{v}$, hence, $\theta_I \in (\bar{v}, v_C)$. Recall that $\theta_I < \bar{v}$ implies $\lambda_{\theta_I} < 0$. Thus, the same argument as in Case 1 yields that (11) is negative for all $x \in [0, \sqrt{\frac{-\lambda_{\theta_I}}{2}}]$. As before, if $x > \sqrt{\frac{-\lambda_{\theta_I}}{2}}$ then (11) is negative if and only if (16) holds. By $\lambda_{\theta_I} < 0$ LHS (16) is strictly increasing in x . Furthermore, $\lim_{x \rightarrow \infty} \text{LHS (16)} = 2\bar{\mu}_{\theta_I} > 1$ because $\theta_I < v_C$ implies $\bar{\mu}_{\theta_I} > 1/2$. Thus, there must be 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.

□

Proposition 1. *There exists $\hat{v}_f \in (v_C, \bar{v}]$ such that the incumbent gambles if and only if $v_I < \hat{v}_f$. Furthermore, if $\gamma < \frac{1}{2}$ then $\hat{v}_f < \bar{v}$. Otherwise, if $\gamma \geq \frac{1}{2}$ then $\hat{v}_f = \bar{v}$.*

Proof. Note that in the fixed-valence model $v_I = \theta_I$, thus, we will use the θ_I notation in proving the result. We break the argument in two cases, depending on whether γ is greater than 1/2.

Case 1: $\gamma < \frac{1}{2}$. Recall, from the proof of Lemma 3, that $\gamma < \frac{1}{2}$ implies $v_C < \bar{v}$.

First, we show that if $\theta_I > \bar{v}$ then the incumbent does not gamble. In this case, $v_I > \bar{v} > v_C$ and, therefore, by Lemma 3 I 's probability of winning is decreasing in x . Consequently, for any policy $x > \hat{x}_I$ deviating to \hat{x}_I gives I higher policy utility and a greater probability of winning. Thus, if $v_I > \bar{v}$ then $x_f^* \leq \hat{x}_I$ and the incumbent never gambles in equilibrium.

Second, consider $v_I < v_C < \bar{v}$. By Lemma 3 the incumbent's probability of winning is increasing in x . Therefore, for any $x < \hat{x}_I$ I deviating to \hat{x}_I gives I higher policy utility and a greater probability of winning. Thus, if $v_I < v_C$ then $x_f^* \geq \hat{x}_I$, and the incumbent always gambles in equilibrium.

Finally, assume $\theta_I \in (v_C, \bar{v})$. From the proof of Lemma 3, we have that the incumbent's probability of winning is maximized at the positive solution to:

$$2\bar{\mu}_{\theta_I} + \frac{\lambda_{\theta_I}}{2x^2} = 1, \quad (18)$$

which we denote as $x_{\theta_I}^w$. If $x_{\theta_I}^w < \hat{x}_I$ then the incumbent's probability of winning is decreasing in x for $x > \hat{x}_I$. Thus, $x_f^* \leq \hat{x}_I$. On the other hand, if $x_{\theta_I}^w > \hat{x}_I$ then the incumbent's probability of winning is increasing in x for $x < \hat{x}_I$. Thus, $x_f^* \geq \hat{x}_I$. Solving (18) explicitly for $x_{\theta_I}^w$ yields:

$$x_{\theta_I}^w = \sqrt{\frac{\lambda_{\theta_I}}{2 - 4\bar{\mu}_{\theta_I}}}.$$

Recall that $\lambda_{\theta_I} > 0$ for $\theta_I < \bar{v}$ and that $\bar{\mu}_{\theta_I} = 1/2$ when evaluated at $\theta_I = \theta_C$. Thus, $\lim_{\theta_I \rightarrow \theta_C} x_{\theta_I}^w = \infty$. Additionally, $\lim_{\theta_I \rightarrow \bar{v}} x_{\theta_I}^w = 0$, because $\lambda_{\theta_I} = 0$ at $\theta_I = \bar{v}$.

Finally, differentiating we have $\frac{\partial \bar{\mu}_{\theta_I}}{\partial \theta_I} = -\frac{1}{8\hat{x}_I} < 0$ and $\frac{\partial \lambda_{\theta_I}}{\partial \theta_I} = \frac{1}{\bar{\mu}_{\theta_I}(1-\bar{\mu}_{\theta_I})} \frac{\partial \bar{\mu}_{\theta_I}}{\partial \theta_I} < 0$. Therefore, $\frac{\partial x_{\theta_I}^w}{\partial \theta_I} < 0$. This implies there is a unique cutoff $v_f^g \in (v_C, \bar{v})$ such that $x_f^* \geq \hat{x}_I$ if and only if $\theta_I = v_I \leq \hat{v}_f$, as required. Specifically, v_f^g is given by the θ_I that solves

$$\hat{x}_I = \sqrt{\frac{\lambda_{\theta_I}}{2 - 4\bar{\mu}_{\theta_I}}}.$$

Case 2: $\gamma > \frac{1}{2}$. In this case, $v_C > \bar{v}$. First, assume $\theta_I > v_C > \bar{v}$. By Lemma 3, the probability of winning is decreasing in x . Therefore, $x_f^* \leq \hat{x}_I$, and the incumbent never gambles if $\theta_I > v_C$. Second, if $\theta_I < \bar{v} < v_C$ then by Lemma 3 the incumbent's probability of winning is increasing in x . Therefore, $x_f^* \geq \hat{x}_I$ and the incumbent always gambles when $\theta_I = v_I < \bar{v}$, as required.

Finally, consider $\theta_I \in (\bar{v}, v_C)$. To finish proving the result we must show that I does not gamble in this case, $x_f^* < \hat{x}_I$. From the proof of Lemma 3 there exists a cutoff $x_{\theta_I}^w$ such that I 's probability of winning is increasing in x for $x > x_{\theta_I}^w$ and decreasing in x for $x < x_{\theta_I}^w$.

We first prove that the optimal policy cannot be above $x_{\theta_I}^w$. Because the voter's signal becomes

perfectly informative about \hat{x}_V when $x \rightarrow \infty$, we have $\lim_{x \rightarrow \infty} P_{\theta_I}(x) = \gamma$. Thus, the incumbent's expected utility from choosing any $x \geq x_{\theta_I}^w$, is bound above by $-4\hat{x}_I^2(1 - \gamma) + \gamma\beta$, which is the payoff from getting policy $x = \hat{x}_I$ and winning with probability γ . On the other hand, because $\theta_I > \bar{v}$, 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: $-\hat{x}_I^2 + \beta$. Consequently, a sufficient condition to ensure $x_f^* < \hat{x}_{\theta_I}^w$ is that

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

which holds by Assumption 1.

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

Lemma (A5). *The following holds: $\lim_{\theta_I \rightarrow \bar{v}_I^+} x_f^* = 0$.*

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

Step 1. We start by establishing that the $\theta_I = \bar{v}$ incumbent's equilibrium payoff must be lower than his payoff if could get policy $x = 0$ and win the election with probability 1. To do so, we show that the \bar{v} politician's probability of winning is bound above by $\max\{\gamma, 1/2\}$. Note, from our assumption that the voter randomizes when indifferent, inspecting equation 5 we have that at $\theta_I = \bar{v}$ the probability of winning must go to $\frac{1}{2}$ as x goes to 0.¹⁵ Furthermore, from the proof of Lemma 3, at $\theta_I = \bar{v}$ 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 $\gamma > 1/2$. As $x \rightarrow \infty$ the voter's signal is perfectly informative. Thus, I 's probability of winning is γ . Therefore, I 's probability of winning must be strictly increasing in x . Second, consider $\gamma < 1/2$. In this case, an analogous argument yields that I 's probability of winning is strictly decreasing in x . Thus, I 's probability of winning is maximized at $x = 0$. At $\theta_I = \bar{v}$ the voter is indifferent between I and C ex ante. Hence, if $x = 0$ then the voter remains indifferent between I and C , as there is no learning when $x = 0$, and by assumption the voter reelects I with probability $1/2$, as required.

¹⁵Without this tie-breaking assumption, for example if the voter always reelects when indifferent, we still obtain similar results but the \bar{v} type actually implements $x_f^* = 0$.

Step 2. Let $U_{\theta_I}(x) = -(x - \hat{x}_I)^2 + P_{\theta_I}(x)\beta - (1 - P_{\theta_I}(x))4\hat{x}_I^2$ be the expected utility to an incumbent with valence θ_I from choosing policy x . Additionally, define $U^0 = \beta - \hat{x}_I^2$ as I 's payoff from $x = 0$ and winning with probability 1. The above argument yields that for $\theta_I = \bar{v}$ we have $U_{\bar{v}}(x) \leq \max\{\frac{1}{2}, \gamma\}\beta - 4\hat{x}_I^2(1 - \max\{\frac{1}{2}, \gamma\})$, which is the utility from choosing \hat{x}_I and winning with the highest possible probability. Thus, a sufficient condition for $U^0 > U_{\theta_I}(x)$ to hold for all x is that:

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

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

Step 3. We show that for any $x > 0$ there exists δ_x such that if $|\theta_I - \bar{v}| < \delta_x$ then $U^0 - U_{\theta_I}(x) > 0$. Consider the following manipulation:

$$\begin{aligned} U^0 - U_{\theta_I}(x) &= U^0 - U_{\theta_I}(x) + U_{\bar{v}}(x) - U_{\bar{v}}(x) \\ &= U^0 - U_{\theta_I}(x) + U_{\bar{v}}(x) - U_{\bar{v}}(x) \\ &= U^0 - U_{\bar{v}}(x) - (U_{\theta_I}(x) - U_{\bar{v}}(x)), \end{aligned}$$

Therefore, $\lim_{\theta_I \rightarrow \bar{v}} U^0 - U_{\theta_I}(x) = \lim_{\theta_I \rightarrow \bar{v}} U^0 - U_{\bar{v}}(x) - (U_{\theta_I}(x) - U_{\bar{v}}(x))$. If $x > 0$ then $U_{\theta_I}(x)$ is continuous in θ_I . Hence,

$$\lim_{\theta_I \rightarrow \bar{v}} U^0 - U_{\bar{v}}(x) - (U_{\theta_I}(x) - U_{\bar{v}}(x)) = U^0 - U_{\bar{v}}(x).$$

By step 1 of the proof $U^0 - U_{\bar{v}}(x) > 0$ for any $x > 0$. For $\theta_I > \bar{v}$ we have $U_{\theta_I}(0) = U^0$. Thus, if $\lim_{\theta_I \rightarrow \bar{v}^+} x_f^* > 0$, then there exists θ_I arbitrarily close to \bar{v} such that $U^0 > U_{\theta_I}(x_f^*)$, a contradiction. \square

Proposition 2. *There exists a cut-point $\underline{v} \geq 0$, with $\underline{v} < \bar{v}$, such that, if $v_I \in (\underline{v}, \bar{v})$ then the equilibrium policy in the fixed-valence model, x_f^* , is decreasing in v_I . Otherwise, x_f^* is increasing in v_I .*

Proof. By Lemma 2 any equilibrium policy must solve the first-order condition. Applying the implicit function theorem we have

$$\frac{\partial x_f^*}{\partial v_I} = \frac{\partial x_f^*}{\partial \theta_I} = -(\beta + 4\hat{x}_I^2) \frac{\frac{\partial^2 P_{\theta_I}}{\partial \theta_I \partial x_1}}{-2 + \frac{\partial^2 P_{\theta_I}}{\partial x^2}}.$$

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

$$\begin{aligned} \frac{\partial^2 P_{\theta_I}}{\partial \theta_I \partial x_1} &< 0. \\ \Leftrightarrow \frac{1}{2x^2} \frac{\partial \lambda_{\theta_I}}{\partial \theta_I} &\left(\gamma \phi(\Delta_{\theta_I}^-) + (1 - \gamma) \phi(\Delta_{\theta_I}^+) + \gamma \phi'(\Delta_{\theta_I}^-) \Delta_{\theta_I}^+ + (1 - \gamma) \phi'(\Delta_{\theta_I}^+) \Delta_{\theta_I}^- \right) < 0, \end{aligned}$$

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

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

Therefore, $\frac{\partial x_f^*}{\partial \theta_I} < 0$ if and only if $1 - \Delta_{\theta_I}^- \Delta_{\theta_I}^+ > 0$, which can be rewritten as

$$4(x_f^*)^2(1 + (x_f^*)^2) > \lambda_{\theta_I}^2. \quad (19)$$

First, we show that if $\theta_I > \bar{v}$ then inequality (19) fails, and thus $\frac{\partial x_f^*}{\partial \theta_I} > 0$ for all $\theta_I > \bar{v}$. Lemma A5 shows that $\lim_{\theta_I \rightarrow \bar{v}} x_f^* = 0$. Therefore, x_f^* must be increasing in θ_I for $\theta_I > \bar{v}$ sufficiently close to \bar{v} . Hence, (19) fails for all θ_I sufficiently close to \bar{v} . Suppose that eventually (19) holds at some $\theta_I > \bar{v}$. Since x_f^* and λ_{θ_I} are continuous in θ_I there exists a point $v' > \bar{v}$ such that (19) fails for $\theta_I < v'$, and holds with equality at $\theta_I = v'$.

Since (19) fails for $\theta_I < v'$ it must be that LHS (19) is increasing faster in θ_I than RHS (19) at $\theta_I = v'$. To prove that (19) fails for all $\theta_I > \bar{v}$ we show that the existence of a such a v' where LHS (19) is increasing faster than RHS (19) yields a contradiction. Differentiating both sides:

$$\begin{aligned} \frac{\partial LHS(19)}{\partial \theta_I} &= 8 \frac{\partial x_f^*}{\partial \theta_I} x_f^* \left[1 + 2(x_f^*)^2 \right] \\ \frac{\partial RHS(19)}{\partial \theta_I} &= \frac{\partial \lambda_{\theta_I}}{\partial \theta_I} 2\lambda_{\theta_I}. \end{aligned}$$

Since $\frac{\partial \lambda_{\theta_I}}{\partial \theta_I} < 0$ and $\lambda_{\theta_I} < 0$ for all $\theta_I > \bar{v}$, this yields $\frac{\partial RHS(19)}{\partial \theta_I} > 0$. However, by construction, if (19) holds with equality at $\theta_I = v'$ then $\frac{\partial x_f^*}{\partial \theta_I} |_{\theta_I=v'} = 0$, contradicting that $\frac{\partial LHS(19)}{\partial \theta_I} > \frac{\partial RHS(19)}{\partial \theta_I}$ at $\theta_I = v'$.

Second, suppose $\theta_I \leq \bar{v}$. We show there exists $\underline{v} < \bar{v}$ such that $\frac{\partial x_f^*}{\partial v_I} = \frac{\partial x_f^*}{\partial \theta_I} < 0$ if and only if $\theta_I \in (\underline{v}, \bar{v})$. By construction, $\lambda_{\theta_I} = 0$ at $\theta_I = \bar{v}$. Thus, inequality (19) must hold because $x_f^* > 0$. Since x_f^* is continuous in θ_I in this range, it must also hold for all θ_I sufficiently close to \bar{v} . That is, there must exist \underline{v} such that (19) holds for all $\theta_I \in (\underline{v}, \bar{v}]$. We now prove that (19) can only hold

over this interval. Specifically, we show that once (19) holds at some v' it must also hold for all $\theta_I \in (v', \bar{v})$.

Suppose (19) holds for some $\theta_I < v'$ and fails for some $\theta_I > v'$. Since both sides of (19) are continuous in θ_I for $\theta_I < \bar{v}$ there must exist some v' such that (19) holds with equality at $\theta_I = v'$. Recall that $\lambda_{\theta_I} > 0$ for $\theta_I < \bar{v}$ and $\frac{\partial \lambda_{\theta_I}}{\partial \theta_I} < 0$. Thus, $\lambda_{\theta_I}^2$ is decreasing in θ_I for $\theta_I < \bar{v}$. Additionally for $\theta_I < v'$ we have that $\frac{\partial x_f^*}{\partial \theta_I} < 0$. Hence, for $\theta_I < v'$ both sides of (19) are decreasing in θ_I . Since (19) holds for $\theta_I < v'$, for (19) to hold with equality at $\theta_I = v'$ it must be that $\frac{\partial LHS(19)}{\partial \theta_I}|_{\theta_I=\bar{v}} < \frac{\partial RHS(19)}{\partial \theta_I}|_{\theta_I=\bar{v}}$. By the same argument in the previous case, $\frac{\partial LHS(19)}{\partial \theta_I}|_{\theta_I=\bar{v}} = 0$. In contrast, $\frac{\partial RHS(19)}{\partial \theta_I} < 0$ for all $\theta_I < \bar{v}$, a contradiction. Thus, such a v' cannot exist, as required. \square

Proposition 3. *If \hat{x}_I is sufficiently small then $\underline{v} > 0$.*

Proof. When $\hat{x}_I = \frac{v_C}{4}$ the voter kicks out the incumbent if $\theta_I = 0$, regardless of the signal on the policy dimension. Thus, $x_f^* = \hat{x}_I$ when $\theta_I = 0$. However, for $\theta_I > 0$ the incumbent has a strictly positive probability of winning. Thus, by Lemma 3, we have $\frac{\partial P_{\theta_I}}{\partial x} > 0$ for θ_I sufficiently small. Inspecting I 's first-order condition, clearly $x_f^* = \hat{x}_I$ cannot be optimal if $\frac{\partial P_{\theta_I}}{\partial x} > 0$. Therefore, for θ_I sufficiently close to 0 we have $x_f^* > \hat{x}_I$, which implies that x_f^* is increasing and $\theta_I = v_I$ and hence $\underline{v} > 0$ at $\hat{x}_I = \frac{v_C}{4}$. As \underline{v} is continuous in \hat{x}_I , this yields $\underline{v} > 0$ for all \hat{x}_I sufficiently close to $\frac{v_C}{4}$. \square

Proposition 4. *There exists \hat{v}_u such that the incumbent gambles if and only if $v_I < \hat{v}_u$. Moreover, generically $\hat{v}_u \neq \bar{v}$.*

Proof. Recall that $v_I = \pi_I$ in the uncertain-valence model. At $\pi_I = 0$ the incumbent's probability of winning is strictly decreasing in x by Lemma 3, thus, $x_u^* > \hat{x}_I$. At $\pi_I = 1$ the incumbent's probability of winning is strictly increasing in x by Lemma 3, thus, $x_u^* < \hat{x}_I$. Finally, Proposition 5 (proved below) shows that x_u^* is decreasing in π_I , which yields the existence of a unique cutoff \hat{v}_u .

To see that $\hat{v}_u \neq \bar{v}$ for almost all parameters we show that \hat{v}_u is decreasing in v_C at $\hat{v}_u = \bar{v}$, whereas \bar{v} is clearly always strictly increasing in v_C . Notice that \hat{v}_u solves $x_u^* = \hat{x}_I$. Therefore:

$$\frac{\partial \hat{v}_u}{\partial v_C} = -\frac{\partial x_u^*/\partial v_C}{\partial x_u^*/\partial v_I}.$$

By Proposition 5 below we have $\frac{\partial x_u^*}{\partial v_I} < 0$. Thus, if $\frac{\partial x_u^*}{\partial v_C} < 0$ then $\frac{\partial \hat{v}_u}{\partial v_C} < 0$. Applying the implicit function theorem to I 's first-order condition yields:

$$\frac{\partial x_u^*}{\partial v_C} = -\frac{(\beta + 4\hat{x}_I^2) \frac{\partial^2 P_{\theta_I}}{\partial x \partial v_C}}{-2 + (\beta + 4\hat{x}_I^2) \frac{\partial^2 P_{\theta_I}}{\partial x^2}}. \quad (20)$$

At $x = x_u^*$ the denominator of 20 must be negative. Therefore, $\frac{\partial x_u^*}{\partial v_C}$ is negative if and only if

$\frac{\partial^2 P_{\theta_I}}{\partial x \partial v_C} < 0$. Recall that $\frac{\partial P_{\theta_I}}{\partial x} = \frac{1}{x} \Gamma_{\theta_I}(x)$. Differentiating with respect to v_C we obtain

$$\frac{\partial^2 P_{\theta_I}}{\partial x \partial v_C} = \gamma \frac{\partial \Delta_{\theta_I}^+}{\partial v_C} \phi(\Delta_{\theta_I}^-) + \gamma \Delta_{\theta_I}^+ \frac{\partial \Delta_{\theta_I}^-}{\partial v_C} \phi'(\Delta_{\theta_I}^-) + (1 - \gamma) \frac{\partial \Delta_{\theta_I}^-}{\partial v_C} \phi(\Delta_{\theta_I}^+) + (1 - \gamma) \Delta_{\theta_I}^- \frac{\partial \Delta_{\theta_I}^+}{\partial v_C} \phi'(\Delta_{\theta_I}^+). \quad (21)$$

At $\theta_I = \bar{v}$: $\Delta_{\theta_I}^- = -2$, $\Delta_{\theta_I}^+ = 2$, $\bar{\mu} = \gamma$, $\frac{\partial \Delta_{\theta_I}^+}{\partial v_C} = \frac{\partial \Delta_{\theta_I}^-}{\partial v_C} = -\frac{1}{4x^2} \frac{1}{8\hat{x}_I \gamma (1 - \gamma)}$. Thus at $\theta_I = \bar{v}$ (21) reduces to:

$$\frac{\partial^2 P_{\theta_I}}{\partial x \partial v_C} = \frac{1}{4x^2} \frac{1}{8\hat{x}_I \gamma (1 - \gamma)} \left(-\gamma \phi(-2) - \gamma 2\phi'(-2) - (1 - \gamma)\phi(2) + (1 - \gamma)2\phi'(2) \right).$$

Recall that ϕ is the standard normal distribution, therefore, $\phi'(-2) > 0$ and $\phi'(2) < 0$. Hence, $\frac{\partial^2 P_{\theta_I}}{\partial x \partial v_C} \Big|_{\theta_I = \bar{v}} < 0$, completing the proof. \square

Proposition 5. *The equilibrium policy in the uncertain-valence model, x_u^* , is decreasing in v_I .*

Proof. From Proposition 2 any equilibrium policy must solve the first-order condition. Applying the implicit function theorem yields:

$$\frac{\partial x_u^*}{\partial v_I} = \frac{\partial x_u^*}{\partial \pi_I} = -(\beta + 4\hat{x}_I^2) \frac{\frac{\partial P_1}{\partial x_1} - \frac{\partial P_0}{\partial x_1}}{-2 + \pi_I \frac{\partial^2 P_1}{\partial x^2} + (1 - \pi_I) \frac{\partial^2 P_0}{\partial x^2}}.$$

Therefore, $\frac{\partial x_u^*}{\partial \pi_I} \leq 0$ if and only if $\frac{\partial P_1}{\partial x_1} - \frac{\partial P_0}{\partial x_1} \leq 0$. We have $1 > \max\{v_C, \tilde{v}_C\}$ and $0 < \min\{v_C, \tilde{v}_C\}$, and so Lemma 3 yields $\frac{\partial P_1}{\partial x_1} \leq 0$ and $\frac{\partial P_0}{\partial x_1} \geq 0$. Therefore, $\frac{\partial x_u^*}{\partial \pi_I} \leq 0$, as required. \square

Proposition 6. *Assume office benefit is sufficiently large. If $v_I \in [\gamma, \bar{v}]$ then there exist parameter values such that the incumbent generates a valence test in equilibrium. Otherwise, if $v_I \notin [\gamma, \bar{v}]$ then I never generates a valence test.*

Proof. We first establish that when β is sufficiently high the incumbent chooses to have a valence test or not based only on which action yields a higher probability of winning.

Recall that in normal times the incumbent's equilibrium policy solves

$$\frac{-2(x - \hat{x}_I)}{\beta + 4\hat{x}_I^2} + \frac{\partial P_{\theta_I}}{\partial x} = 0.$$

As $\beta \rightarrow \infty$ the LHS of the FOC goes to $\frac{\partial P_{\theta_I}}{\partial x}$. Since the incumbent's problem is continuous in β it must be that x_f^* approaches the policy that maximizes I 's winning probability. Denote this policy as $x_p^* = \operatorname{argmax}_x P_{\theta_I}(x)$. Moreover, $\lim_{\beta \rightarrow \infty} U_f(x_f^*) = (\beta + 2\hat{x}_I) P_{\theta_I}(x_p^*)$. To see this, consider the

ratio

$$\begin{aligned} \frac{U_f(x_f^*)}{(\beta + 2\hat{x}_I)P_{\theta_I}(x_p^*)} &= \frac{-(x_f^* - \hat{x}_I)^2 + (\beta + 2\hat{x}_I)P_{\theta_I}(x_n^*)}{(\beta + 2\hat{x}_I)P_{\theta_I}(x_p^*)} \\ &= \frac{P_{\theta_I}(x_f^*)}{P_{\theta_I}(x_p^*)} - \frac{|x_f^* - \hat{x}_I|}{(\beta + 2\hat{x}_I)P_{\theta_I}(x_p^*)}. \end{aligned}$$

Thus, $\lim_{\beta \rightarrow \infty} \frac{U_f(x_f^*)}{(\beta + 2\hat{x}_I)P_{\theta_I}(x_p^*)} = 1$. A similar argument yields that $\lim_{\beta \rightarrow \infty} \frac{U_u(x_u^*)}{(\beta + 2\hat{x}_I)\pi_I P_S(x)} = 1$, by noting that $x = 0$ maximizes $(\beta + 4\hat{x}_I^2)\pi_I P_1(x)$. Consequently, for β sufficiently high, $U_f(x_f^*) < U_u(x_u^*)$ if and only if $(\beta + 4\hat{x}_I^2)P_{\theta_I}(x_p^*) < (\beta + 4\hat{x}_I^2)\pi_I P_S(0)$, which reduces to $P_{\theta_I}(x_p^*) < \pi_I$.

We now prove the result by considering different values of v_I .

First, we show that if $\pi_I > \bar{v}$ then the incumbent never wants to generate a valence test for β sufficiently high. If $\pi_I > \bar{v}$ then $P_{\theta_I}(x_p^*) = P_{\theta_I}(0) = 1 > \pi_I$, and so the incumbent's equilibrium payoff is higher with no valence test.

Next, suppose $\pi_I < \gamma$. In this case, $\max_x P_{\theta_I}(x) \geq \gamma > \pi_I$, and so the incumbent's equilibrium policy is higher without a valence test. Moreover, this implies that if $\gamma > \bar{v}$ then the incumbent is never better under a valence test.

Now suppose $\pi_I \in (\gamma, \bar{v})$. We show there exists parameters such that the incumbent's equilibrium payoff is higher under a valence test. Specifically, assume $\gamma > 1/2$. This implies $\bar{v} < v_C = \pi_C$. Let $\gamma < \bar{v}$. If $\pi_I \in [\gamma, \bar{v}]$ then $\max_x P_{\theta_I}(x) = \gamma$, which is strictly less than π_I . Thus, if $\gamma > 1/2$, then for all $\pi_I \in [\gamma, \bar{v}]$ the incumbent's equilibrium payoff is higher in a valence test. Moreover, our previous arguments imply that these are the only values of π_I for which the incumbent does better in a valence test when $\gamma > 1/2$.

To finish showing the result, next assume $\gamma < \min\{\bar{v}, \pi_C\}$. Similar to before, we have that if $\pi_I \in (\gamma, \min\{\bar{v}, \pi_C\})$ then $\max_x P_{\theta_I}(x) = \gamma < \pi_I$, and so the incumbent's equilibrium payoff is higher with the valence test. We also note that $\gamma < \pi_C$ also implies $1/2 < \bar{v}$. Therefore, for all $\pi_I \in (1/2, \bar{v}]$ we have $\max_x P_{\theta_I}(x) \leq 1/2 < \pi_I$. Thus, all $\pi_I \in (1/2, \bar{v}]$ prefer a valence test over no test. Furthermore, note that these two arguments together imply that if $\gamma < 1/2 < \pi_C$ then all $\pi_I \in [\gamma, \bar{v}]$ prefer a valence test, and all $\pi_I \notin [\gamma, \bar{v}]$ do better when there is no valence test. \square

B Proofs for Crisis Application

Proposition 7. *If $v_I > \bar{v}$ then the incumbent enacts a more extreme policy during times of crisis than during normal times. If $v_I < \underline{v}$ then the incumbent implements a more moderate policy during times of crisis than during normal times. If $v_I \in (\underline{v}, \bar{v})$ then the crisis can lead to more or less extreme policies.*

Proof. First, suppose $v_I > \bar{v}_I$. By Proposition 2 x_f^* is increasing in v_I , whereas x_u^* is decreasing in v_I by Proposition 5. Moreover, $x_f^* = x_u^*$ at $v_I = 1$, thus $x_f^* < x_u^*$, as required. Second, let $v_I < \underline{v}$. Again, by Proposition 2 x_f^* is increasing in v_I , while x_u^* is decreasing in v_I by Proposition 5. Furthermore, $x_f^* = x_u^*$ at $v_I = 0$, hence $x_f^* > x_u^*$ for $v_I < \underline{v}$. \square

We now present numerical examples to demonstrate that the crisis can lead to more or less reform when $v_I \in (\underline{v}, \bar{v})$. Figure 6 shows an example where $\hat{v}_u < \hat{v}_f$, and hence the crisis leads to less extreme policies for $v_I \in (\hat{v}_u, \hat{v}_f)$. In this example, $\hat{v}_u = 0.071$ and $\hat{v}_f \approx 0.631$. Figure 7 instead provides an example where $\hat{v}_u > \hat{v}_f$, and thus the crisis generates more extreme policies for $v_I \in (\hat{v}_f, \hat{v}_u)$. In this example, $\hat{v}_u \approx 0.776$ and $\hat{v}_f \approx 0.438$.




Figure 6: Numerical example where $\hat{v}_u < \hat{v}_f$. The figure compares equilibrium policies in a crisis and in normal times, as a function of v_I . Parameters: $\gamma = .55$, $v_C = .7$, $\beta = .1$, and $\hat{x}_I = .2$.




Figure 7: Numerical example where $\hat{v}_u > \hat{v}_f$. The figure compares equilibrium policies in a crisis and in normal times, as a function of v_I . Parameters: $\gamma = .55$, $v_C = .5$, $\beta = .1$, and $\hat{x}_I = .176$.

Non-orthogonal Crises

Consistent with earlier models, we have assumed that the valence dimension is orthogonal to the policymaking dimension. For the crisis application, this orthogonality allows us to isolate the incumbent’s incentives to supply policy experimentation. However, policymakers may also implement interventions on dimensions related to an ongoing crisis, even if they are not directly aimed at solving it. For example, Pitlik and Wirth (2003) studies whether economic crises lead countries to adopt market-oriented reforms. In this section we show that our conditional effect of crises on policymaking can hold even if the crisis and policy dimensions are not fully orthogonal. In particular, the supply-side incentives underlying Proposition 7 can continue to hold even if the crisis changes voters’ policy demands.

To capture this demand-side channel in our model we assume that the crisis alters the players’ prior belief about the voter’s optimal policy \hat{x}_V . Specifically, under a crisis $Pr(\hat{x}_V = 1 | \text{crisis}) = \gamma_c \neq \gamma = Pr(\hat{x}_V = 1 | \text{no crisis})$. Here, a crisis affects the incumbent’s strategic considerations because it tests his ability *and* changes how favorable voters are ex-ante towards his policy position. Although this shifts the incumbent’s optimal policy choice, his policymaking incentives are similar to the baseline model.

For simplicity, we consider cases in which the crisis has a large impact on the voter’s policy demands. Additionally, this highlights that even a significant demand-side shock does not necessarily alter our insights into supply-side incentives,

Proposition (A1). *Assume γ_c is sufficiently large or sufficiently small. If v_I is sufficiently large, then the incumbent enacts a more extreme policy during times of crisis than during normal times. Otherwise, the incumbent implements a more moderate policy during times of crisis than during*

normal times.

Proof. At $\gamma_c = 1$, for almost all π_I the voter either strictly prefers to reelect the incumbent or the challenger, because the outcome on the policy dimension does not shift her prior belief. Consequently, the incumbent has no incentive to choose any $x \neq \hat{x}_I$. For γ sufficiently high x_c^* is continuous in γ . Thus, for each π_I for any $\epsilon > 0$ we can find a γ sufficiently close to 1 such that $|x_u^* - \hat{x}_I| < \epsilon$, which yields the directional prediction. A similar argument yields the result for γ_c close to 0. \square

Consider a crisis that makes the voter sufficiently convinced that her ideal policy is aligned with the incumbent, i.e., γ_c close to 1. This decreases the incumbent's incentives to control information because under such a strong prior his retention chances are very inelastic to the realized policy outcome. As a consequence, the equilibrium policy in times of crisis moves closer to the incumbent's static optimum. Graphically, this flattens the blue curve in Figure 4 towards the incumbent's ideal point. However, the policy remains decreasing in the incumbent's expected ability for the same logic described earlier. Substantively, the demand-side channel implies that the crisis further dampens the incumbent's strategic incentives to control information. Thus, the crisis continues to create more reform by incumbents who are electorally leading and less reform by those who are behind. A similar argument holds if the crisis instead convinces the voter she is likely to be ideologically aligned with the challenger.

C Proofs for Asymmetric Uncertainty Extension

Lemma 4. *There always exists an equilibrium where the incumbent chooses $x_f^b(v_I)$ following either signal, $x_f^a(0) = x_f^a(1) = x_f^b(v_I)$.*

Proof. Suppose that $x_f^a(0) = x_f^a(1) = x_f^b(v_I)$, which implies $\mu_{\theta_I}(x_f^b(v_I)) = \pi_I$. For any x_1 off the path of play assume that $\mu_{\theta_I}(x_1) = \pi_I$. Then the expected utility to the $s_I \in \{0, 1\}$ type from any policy x is:

$$-(\hat{x}_I - x)^2 + P_{\pi_I}(x) \cdot \beta - (1 - P_{\pi_I}(x)) \cdot 4\hat{x}_I^2. \quad (22)$$

By definition $x_f^b(v_I)$ maximizes equation (22), thus, neither type of I has a profitable deviation. \square

Lemma (A6). *Consider the fixed-valence model. Among all pooling equilibria, the one where both types of the incumbent choose $x_f^*(v_I)$ 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 $\mu_{\theta_I}(x_1) = \pi_I$. Thus, in any such equilibrium the payoff of type s_I is given by:

$$-(\hat{x}_I - x_1)^2 + P_{\pi_I}(x_1) \cdot \beta - (1 - P_{\pi_I}(x_1)) \cdot 4\hat{x}_I^2. \quad (23)$$

By construction, $x_f^*(v_I)$ maximizes (23), as required. \square

Lemma 5. *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 $s_I = 1$ type strictly prefers policy x' over policy x'' . Recall that both types have the same expected utility for any policy x_1 . Thus, the $s_I = 0$ 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. \square

Proposition 8. *Among all equilibria, the equilibrium where $x_f^a(0) = x_f^a(1) = x_f^b(v_I)$ maximizes the incumbent's expected utility under each signal.*

Proof. By Lemma (A6) 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 $\pi_I \geq \mu_{\theta_I}(x') \equiv \mu'_{\theta_I}$. By Lemma 5 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:

$$\max_x -(\hat{x}_I - x)^2 + P_{\mu'_{\theta_I}}(x) \cdot \beta - (1 - P_{\mu'_{\theta_I}}(x)) \cdot 4\hat{x}_I^2. \quad (24)$$

In contrast, in the equilibrium where $x_f^a(0) = x_f^a(1) = x_f^b(v_I)$, the equilibrium payoff of both types is given by:

$$\max_x -(\hat{x}_I - x)^2 + P_{\pi_I}(x) \cdot \beta - (1 - P_{\pi_I}(x)) \cdot 4\hat{x}_I^2. \quad (25)$$

As $\pi_I \geq \mu'_{\theta_I}$, the envelope theorem delivers that (25) is greater than (24), completing the argument. \square

Lemma 6. *In every equilibrium, the incumbent chooses different policies following each signal, $x_u^a(0) \neq x_u^a(1)$. Furthermore, $x_u^a(s_I = 0) = x_u^b(v_I = \psi_0)$ and $x_u^a(s_I = 1) = x_u^b(v_I = \psi_1)$.*

Proof. In the uncertain-valence model 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_I} . \square