Policymaking in Times of Crisis

Peter Bils Federica Izzo

January 2021

Abstract

How do crises influence an executive's willingness to implement policy reforms? While existing work focuses on how crises impact voters' *demand* for reform, we instead investigate how they alter politicians' incentives to *supply* policy experimentation, even if the crisis does not shift voters' policy preferences. To study this problem, we develop a model of elections and policy experimentation. In our setting, voters face uncertainty about their optimal policy and politicians' ability to manage a crisis. We show that extreme reforms generate more information for voters about their optimal platform. Consequently, the incumbent has electoral incentives to engage in information control. At the same time, a crisis represents an exogenous test for the incumbent, who must prove competent enough to successfully manage the country during turbulent times. Therefore, a crisis has an (independent) impact on the incumbent's electoral prospects, and this may influence his incentives to engage in risky policy experiments. We find that, in contrast to the conventional wisdom, a crisis induces bolder policy reforms only when the incumbent is sufficiently likely to be competent. If the incumbent is relatively unlikely to be competent, then the crisis instead results in policies that are closer to the status quo. As such, our model qualifies the standard intuition on this matter, and potentially allows us to make sense of the mixed results emerging in the empirical literature.

Introduction

Officeholders often face risks on multiple dimensions. Some of these risks are imposed exogenously by outside events. Pandemics, natural disasters or global economic downturns represent crises that are not generated by the incumbent, but that require his competent intervention to be addressed. Other risks are endogenous and derive from the officeholder's choice. The incumbent may decide to engage in bold reforms or policy experiments, whose outcome is ex-ante uncertain. In this paper, we ask: How does the *exogenous* risk of a crisis influence the incumbent's willingness to take *endogenous* risks in policymaking?

The conventional wisdom in the literature argues that crises may influence policymaking because they alter voters' demands (see, e.g., Drazen and Easterly, 2001; Drazen, 2000; Stewart, McCarty and Bryson, 2020). Our paper complements these accounts by providing a theory of the *supply-side* effect of crises on policymaking, which can arise even absent a demand-side shock to voters' policy preferences. We thus show that even when the crisis dimension is entirely *orthogonal* to the policy one, it may nonetheless influence the incumbent's willingness to supply policy reforms. For example, our theory presents a framework for analyzing how the current COVID-19 crisis may influence policy reforms on largely unrelated issues, such as the environment, immigration, or economic redistribution.

The source of the supply-side effect we uncover is the anticipation that both the crisis and policy dimensions impact the incumbent's electoral prospects. A crisis represents a test for the incumbent, who must prove his competence by successfully managing the country during turbulent times. The incumbent may succeed at this task, improving his electoral prospects, or fail, hurting his chances (Healy and Malhotra, 2013). At the same time, different policy choices entail different levels of electoral risk. Bolder policy reforms involve higher uncertainty over outcomes, and thus higher risk (whether positive or negative) for the officeholder (Majumdar and Mukand, 2004; Dewan and Hortala-Vallve, 2019). Consequently, policymakers' incentives to implement risky policy experiments or avoid gambles depend on their prevailing electoral prospects, which, in turn, may be altered by the crisis. For instance, such considerations were at play during the Obama administration's deliberations over healthcare reform. President Obama and his advisors discussed how the ongoing financial crisis should affect their approach to healthcare reform, given the electoral risks inherent to both the crisis and the choice to engage in bold policy reforms.[1]

Thus, exogenous crises may influence officeholders' incentives to take risky policy choices. To characterize this effect, we analyze a two-period electoral accountability model. In each period, the country is either in a period of business as usual, or experiences a crisis. In either case, the incumbent chooses how much to reform policy on a dimension orthogonal to the crisis. Finally,

[1] See: `www.newyorker.com/magazine/2020/11/02/barack-obama-new-book-excerpt-promised-land-obamacare`.

the voter decides whether to reelect the incumbent or replace him with an ideologically opposed challenger. Consequently, the voter faces a selection problem. All else equal, she prefers the candidate that is most likely to be competent. However, she also wants to select the candidate whose ideological preferences are closer to her true ideal policy.

In our model, two key ingredients complicate this selection problem. First, voters and politicians alike face uncertainty over both the policy dimension and the crisis dimension. Specifically, they are uncertain about which policy reforms are better for voters, and which candidate is better able to successfully manage a crisis. Second, faced with this uncertainty voters look back at their past experiences to inform their electoral choice. Specifically, voters' preferences over policy reforms evolve as they observe past policy outcomes (similar to Fiorina (1981)), and their beliefs over the incumbent's ability update as they observe his performance under a crisis, if one arises (similar to Ashworth, Bueno de Mesquita and Friedenberg (2018)).

We start by fully characterizing how voters learn from past experiences and how this impacts the incumbent's electoral chances. In our setting, learning about policy is endogenous. Specifically, we show that more extreme reforms (i.e., platforms farther from the status quo) produce more informative outcomes and, thus, increase the amount of voter learning on the policy dimension. To see that, suppose the incumbent implements a bold liberal reform, moving the status quo far to the left. If this policy produces a good outcome for the voter, then it is likely that the reform was in the optimal direction. Conversely, because policy outcomes are noisy and thus the voter's learning is imperfect, the consequences of a policy close to the status quo are much less informative. On the other hand, the possibility to learn about the incumbent's ability to deal with a crisis arises exogenously.

Thus, the incumbent may *endogenously* acquire (or consolidate) an advantage if the voter updates in his favor on the policy dimension, or *exogenously* if a crisis arises and he proves able to solve it. Similarly, negative information on either dimension damages the incumbent's electoral chances. This has two immediate implications. First, the incumbent's willingness to implement more or less radical reforms depends on his ideological preferences as well as his incentives to facilitate or prevent voter learning about the optimal policy. Second, the occurrence of a crisis alters these incentives, and thus has a significant impact on the incumbent's equilibrium choice. In the second part of the paper, we characterize this effect.

We begin by characterizing the incumbent's equilibrium choice in the benchmark no-crisis case. In this case, the incumbent's policy choice is conditioned by his *expected* competence. An incumbent with a low reputation for competence is ousted absent new favorable information on the policy dimension. In equilibrium, the incumbent thus engages in bold reforms, distorting policy away from both the status quo and his own ideal point so as to facilitate voter learning on this dimension. In contrast, an incumbent with high expected competence faces opposite incentives. New information can only hurt this leading incumbent's electoral chances. As such, he wants to implement uninformative policies so as to prevent voter learning. In equilibrium, he thus distorts policy away from his own ideal point and towards the status quo. Comparing this benchmark to the incumbent's behavior under a crisis, we show that crises have a threefold effect on policymaking.

First, when the incumbent is ex-ante electorally disadvantaged crises have a *qualitative* effect on policymaking: they alter the nature of the incumbent's incentives to control information. As emphasized above, a trailing incumbent will engage in policy gambles when no crisis arises, so as to facilitate voter learning on the policy dimension. Suppose instead a crisis emerges in the first period. If the incumbent fails to solve the crisis, it persists to the second period. The voter then always prefers to oust the incompetent incumbent, regardless of what she learns on the policy dimension. If instead the incumbent succeeds in solving the crisis, he will acquire a large lead on the competence dimension, and new policy-relevant information can only hurt his retention chances. Thus, even a trailing incumbent will choose to distort policy towards the status quo during times of crisis: the crisis induces this ex-ante disadvantaged incumbent to behave *as if* he was electorally leading.

Second, for leading incumbents, who do not experience a qualitative effect, the crisis nonetheless has a *quantitative* effect on policymaking: it alters the intensity of the incumbent's incentives to control voter learning. If the incumbent fails to solve the crisis, policy information is electorally irrelevant. If the incumbent succeeds, his electoral lead increases (compared to the no-crisis scenario), thereby making his retention prospects less susceptible to what the voter learns on the policy dimension. Thus, the crisis weakens the incumbent's incentives to prevent voter learning, and as such it moves his equilibrium policy closer to his ideological preferences.

Third, crises have an *indirect* effect on policymaking: they alter how the equilibrium policy responds to changes in model primitives. In particular, we study how the policy changes with the ex-ante probability that the incumbent is a competent type. If the country experiences a crisis, then increasing this probability increases the *likelihood* the incumbent acquires an (ex-post) electoral advantage. Consequently, the incumbent has stronger incentives to preserve this advantage by preventing voter learning on the policy dimension. This, in turn, moves the equilibrium policy closer to the status quo. Under normal times, however, the incumbent's expected competence plays a very different role. As his expected competence increases, the *magnitude* of the incumbent's (ex-ante) electoral advantage increases. This impacts the incumbent's incentives in two ways. First, as he becomes very likely or very unlikely to be competent, policy information becomes less electorally relevant. This weakens the incumbent's incentives to control information and distort policy away from his ideological preferences. Second, increasing the incumbent's expected competence decreases the reelection threshold. This makes information control less important. When the incumbent is leading, the two effects go in the same direction. When he is trailing, instead, these go in opposite directions and the net effect is a combination of the two components. We find that if the incumbent is leading or very far behind, then the equilibrium policy in normal times moves away from the status quo as his expected ability increases, in sharp contrast with the comparative statics under a crisis.

Finally, combining the three results described above, we show that an exogenous crisis has a *directional* effect on policies, but this effect may either increase or reduce experimentation. When the incumbent is sufficiently likely to be competent, crises induce bolder policy experiments (which translates into more policy extremism). For incumbents who are unlikely to be a good type, we should instead observe less policy experimentation, i.e., policies are less extreme and closer to the status quo during turbulent times.

We consider two extensions to the baseline model. First, we ask whether officeholders would ever want to generate a crisis in our setting. When the incumbent is ex-ante behind, he needs the voter to learn about either the competence or the policy dimension in order to win reelection. We show that, if the probability the incumbent is a competent type is larger than the probability of the policy-relevant state being favorable to him, then his best bet is to generate the crisis and subject himself to a test on the competence dimension. Our results therefore highlight that, ironically, the incumbent wants to generate an endogenous crisis when he is not too likely (nor unlikely) to be able to solve it. Second, we allow the incumbent to have private information about his competence. We show that our findings hold in this richer setting, even if the incumbent perfectly observes his type.

Although we see our theory as a complement to, rather than a substitute for, demand-driven accounts of how crises impact policymaking, our predictions are novel and potentially allow us to empirically adjudicate between different arguments. The conventional wisdom in the literature maintains that crises should always increase policy experimentation (Tommasi and Velasco, 1996). However, the empirical evidence on this matter is mixed (see the discussion in Prato and Wolton (2018); Mahmalat and Curran (2018)). While some scholars confirm a positive association between crises and reform (Lora and Olivera, 2004; Alesina, Ardagna and Trebbi, 2006; Pitlik and Wirth, 2003), others find crises may have an opposite effect (Campos, Hsiao and Nugent, 2010; Pop-Eleches, 2008; Castanheira, Nicodème and Profeta, 2012; Galasso, 2014; Mian, Sufi and Trebbi, 2014). By identifying conditions under which the effect of crises on reform may go in one or the other direction, our theory may provide a framework to reevaluate these mixed results. Empirical scholars should consider how the impact of a crisis is mediated by the competitiveness of the electoral environment and, especially, the officeholder's (expected) quality.

Literature

First and foremost, our paper contributes to the small formal literature studying how crises impact political and policy outcomes. Most of the existing work in this tradition conceptualizes crises as a shock to the actors' policy tastes (Drazen and Easterly, 2001). That is, a crisis may lead to policy changes or reform because it shifts players' policy preferences (Levy and Razin, 2021; Bils, 2020); it deteriorates the status quo (Fernandez and Rodrik, 1991; Drazen, 1996; Prato and Wolton, 2018); or it increases special interests' willingness to compromise, which reduces gridlock (Williamson, 1994). In short, these works present theories of the demand-side effect of a crisis.

In contrast, we model a crisis as a test of the officeholder's ability, and provide a theory of the supply-side effect, that will apply even to crises that are orthogonal to the policy dimension. We build on Ashworth, Bueno de Mesquita and Friedenberg (2018), who also study crises as exogenous informative shocks that alter the inferences voters draw upon observing governance outcomes. However, in their model politicians take no strategic action. In turn, Izzo (2020) studies how such crises influence the self-selection of competent candidates into the electoral arena. Here, we complement these approaches by analyzing how crises of this nature impact policymakers' strategic choices.

Our paper also connects to work that studies how elections impact an officeholder's incentives to initiate a crisis, such as an international conflict (e.g., Downs and Rocke, 1994; Smith, 1996; Tarar, 2006; Judd, 2017). In particular, we note that our intuition and results relate to the concept of gambling for resurrection, introduced by Downs and Rocke (1994). The authors study officeholders' incentives to take risky choices, as we do, focusing on the decision whether to initiate an international conflict. They show that incumbents who are ex-ante are electorally disadvantaged will want to gamble for resurrection: engage in inefficient wars in hopes of generating a good outcome and thus winning reelection. Our analysis qualifies this result.[2] We confirm a similar intuition, but also show that disadvantaged incumbents may shy away from a crisis when they can experiment on the policy dimension. Thus, we find that disadvantaged incumbents will choose to generate an endogenous crisis only when the likelihood that the voters' ideal policy is aligned with their own is sufficiently low.

In our world, voters are uncertain about which policy is best for them and learn via experience. Officeholders thus have incentives to strategically use policy to control the *amount* of voter learning, by either engaging in risky policy experiments or playing it safe. In this perspective, our work contributes to the literature on policy experimentation within the context of multi-armed bandit problems (e.g., Strumpf, 2002; Volden, Ting and Carpenter, 2008; Strulovici, 2010; Hirsch, 2016; Dewan and Hortala-Vallve, 2019; Gieczewski and Kosterina, 2020).

[2] As an aside, we note that Dewan and Hortala-Vallve (2019) instead obtain the opposite results: the incumbent gambles when he is very likely to be competent. This is because, in their model, incumbent and challenger have the same expected ability, therefore neither is ex-ante advantaged or disadvantaged.

Within this tradition, our paper is the first to study how crises influence politicians' incentives to take risks in policymaking. Beyond this different substantive focus, our work provides a technical contribution to this literature. Existing works typically consider a binary policy space, with a risky option and a safe one.[3] As such, these works can only analyze a decision-maker's choice *whether* to experiment. In contrast, we consider a framework to study policy experimentation with a continuous space. As such, we can analyze the intensity of the policymaker's dynamic incentives to take risks, and thus characterize the *amount* of policy experimentation emerging in equilibrium. Further, by modelling policy experiments as movements along the left-right spectrum, our framework also allows us to make directional predictions about equilibrium policy.

Callander (2011) also considers 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. As a consequence, bold policy experiments 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 policy changes that facilitate learning. Further, the focus of his 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 builds on the intuition that reelection-facing politicians will have *dynamic* incentives to control information.[4]

Finally, our model contributes to a literature on elections in which candidates have exogenous characteristics that are orthogonal to the policy dimension (e.g., Ansolabehere and Snyder, 2000; Groseclose, 2001; Bernhardt, Câmara and Squintani, 2011; Krasa and Polborn, 2012). In our paper, voters face uncertainty about the incumbent's exogenous characteristic, i.e., his ability to manage a crisis. We show that the resolution of this uncertainty can have a significant impact on policymaking. Moreover, our paper's focus on policy experimentation is unique in this literature.

Model

Players and actions: We consider a two period model of electoral accountability. There is an incumbent (I), challenger (C), and representative voter (V). Players face uncertainty about a policy-relevant state of the world ($\omega \in \{-\alpha, \alpha\}$) and about each politician's competence in the event of a crisis ($\theta_j \in \{0, 1\}$).

In the first period, the country is either in a crisis, or experiences a period of business as usual. In either case, the officeholder chooses a policy $x_1 \in \mathbb{R}$, which leads to a policy outcome o_1^p as described

[3]Strumpf (2002) considers an extension with two experimental policies. Hirsch (2016) considers a binary policy space where one option is not inherently more risky than the other, but a correct policy succeeds only if a bureaucrat exerts sufficient effort in its implementation.

[4]Callander and Hummel (2014) consider forward-looking parties, but assume exogenous retention probabilities.

below. If there is a crisis, it is either successfully managed, $o_1^c = 1$, or not, $o_1^c = 0$, depending solely on the incumbent's competence. Importantly, o_1^c realizes *after* the officeholder chooses which policy to implement in the first period. The voter observes the policy choice, the realized policy outcome, the realization of the crisis and whether it was successfully managed, and decides whether to reelect I or to elect C.

The game then proceeds to the second period. As in the first period, the country either experiences a crisis or a period of business as usual. The probability or the severity of a crisis in the second period (potentially) depends on the first-period crisis outcome. Finally, the officeholder chooses policy $x_2 \in \mathbb{R}$. If a crisis arises, it is successfully managed, or not, depending on the officeholder's competence.

Payoffs and information: The voter cares about outcomes on both the policy and the crisis dimensions. Overall, the voter's *total* per-period utility is:

$$U_t^v = U_t^p + U_t^c,$$

where U_t^p denotes her policy utility and U_t^c her crisis utility for period $t \in \{1, 2\}$.

The voter's *policy* utility in period t is equal to the policy outcome, that is, $U_T^p = o_t^p$. This outcome evolves from one period to the next as a function of policy change, an unknown state of the world, and an idiosyncratic shock:

$$o_t^p = o_{t-1}^p + \omega(x_t - x_{t-1}) + \epsilon_t,$$

where x_t is the implemented policy in period t, x_{t-1} is the previous period's policy choice, o_{t-1}^p is the previous period's policy outcome, and ϵ_t is an idiosyncratic shock. We assume that ϵ_t is i.i.d. in each period from the standard normal distribution Φ, with mean 0, variance 1, and pdf ϕ.

Uncertainty about the voter's optimal policy is captured by the state of the world $\omega \in \{-\alpha, \alpha\}$.[5] Specifically, ω represents the optimal direction of policy change. A positive ω indicates that moving the policy to the right of the status quo is optimal for the voter. Conversely, a negative ω implies that more progressive reforms would improve the voter's (expected) welfare. The voter and the politicians are uncertain about the true value of ω. However, prior to making her electoral decision, the voter observes the realization of the first-period policy outcome, and draws inferences abut the state of the world. Therefore, we interpret the voter's uncertainty as referring to the consequences of the various policy choices. Alternatively, we may think about a voter who sees policy as an experience good, so that she does not know how different policy changes will impact her welfare

[5]The assumption that the two possible values of ω are symmetric around zero is not necessary for our qualitative results.

until such policies are implemented. In order to focus on policy experimentation, we assume away any asymmetry of information: players share a common prior belief that $p(\omega = \alpha) = \gamma \in (0, 1)$.

The voter's *crisis* utility depends on whether there is a crisis and the officeholder's competence. Specifically, in period t her crisis utility is

$$U_t^c = -(1 - o_t^c)\mathbb{I}_c K$$

Where $\mathbb{I}_c = 1$ if a crisis arises in period t and $\mathbb{I}_c = 0$ otherwise. The realization of o_1^c depends on the incumbent's competence. Specifically, each politician $j \in \{I, C\}$ is either the good type, $\theta_j = 1$, or bad type, $\theta_j = 0$. We then have that $o_t^c = \theta$, where θ is the period t's officeholder's type. We assume that j's type is unknown to all players, who share a common prior belief that $p(\theta_j = 1) = \pi_j$. In a separate section below, we analyze an extension of the model where the incumbent has private information about his competence (specifically, he observes a private and potentially fully informative signal of his ability). We show that all our qualitative results survive under this richer setting.

Finally, politicians are motivated by both ideology and winning office. Politician j's per-period utility is

$$U_t^j = \mathbb{I}_j \beta - (x_j - x_t)^2,$$

where $\mathbb{I}_j = 1$ if politician j is in office at time t and $\mathbb{I}_j = 0$ otherwise. For simplicity, we assume $x_I = -x_C > 0$.

Players' dynamic payoffs are given by the sum of per-period utility and we assume no discounting.

Crisis persistence: At the moment of making her retention decision, the voter does not know whether a crisis will materialize in the next period. We allow the second-period probability of a crisis to depend on the first-period crisis outcome. Denote the first-period crisis outcome by $\chi \in \{N, S, F\}$. That is, if a crisis did not arise then $\chi = N$. If there was a crisis and it was successfully managed then $\chi = S$. Finally, if the incumbent failed to solve the crisis then $\chi = F$. Then, the probability of a crisis in period 2 is given by $\xi_\chi p_b$, where p_b is the baseline probability of a crisis, $\xi_N = \xi_S = 1$, and $\xi_F \geq 1$. Therefore, ξ_F indicates the increased probability of a crisis tomorrow due to the incumbent failing to solve the crisis today. Alternatively, we can view ξ_F as the increased severity of the crisis in the second period if it is not solved in the first period. For example, $\xi_F > 1/p_b$ can be interpreted as the crisis persisting into the second period *and* worsening due to incompetent management in the first period. Under either interpretation, the voter's *expected* second-period crisis utility is

$$-\xi_\chi p_b K(1-\pi_j),$$

where π_j is the probability that the second-period officeholder is competent. Thus, we refer to $\xi_\chi p_b K$ as the weighted cost of a (mismanaged) crisis in the second period.

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

1. Nature draws the state of the world ω, I's competence θ_I, and C's competence θ_C.
2. The country is either in crisis or not.
3. I chooses policy $x_1 \in \mathbb{R}$.
4. The voter observes the policy choice, x_1, the policy outcome, o_1^p, the crisis outcome, o_1^c, and updates her beliefs about ω and θ_I.
5. The voter makes her retention decision.
6. If a crisis did not occur, or if I successfully managed the crisis, then a new crisis occurs with probability p_b. Otherwise, if there was a crisis and I failed to solve it, then the weighted cost of a crisis in the second period is $\xi_F p_b K$.
7. The second period officeholder chooses policy $x_2 \in \mathbb{R}$.
8. Utilities are realized and the game ends.

Before proceeding to the analysis, let us emphasize that in our setting crises are entirely orthogonal to the policy dimension. That is, crises have no impact on voters' policy preferences, or on their beliefs over the policy dimension. This also implies that the policy choice has no impact on the probability that the crisis is resolved. For example, environmental or redistributive policies do not determine a country's performance in an international conflict or the likelihood of successfully controlling the COVID-19 crisis. While these elements could be incorporated in our model, our goal in this paper is to provide a theory of the *supply-side* effect of crises on policymaking. For this purpose, we choose to shut down potential demand-side channels.

Voter Learning

As highlighted in the introduction, the voter learning about the incumbent's ability and the optimal policy (reform) is at the core of this model. Thus, before delving into equilibrium analysis, it is worth pausing to analyze how these learning processes work in our setting.

Exogenous Learning on the Incumbent's Competence

First, we consider how the voter updates her beliefs about the incumbent's ability to deal with a crisis. We denote $\mu_\theta(\chi)$ the voter's posterior belief that the incumbent is competent, as a function of the first-period crisis outcome. Recall that $\chi \in \{N, S, F\}$ indicates the first-period crisis occurrence and outcome.

In our setting, a crisis represents a test that is *exogenously* imposed on the incumbent. If there is no crisis, then the voter does not have an opportunity to learn about the incumbent's competence, and simply retains her prior belief, $\mu_\theta(N) = \pi_I$. On the other hand, if there is a crisis then the outcome is entirely dependent on the incumbent's ability. If the crisis is successfully managed, then the voter learns that the incumbent is competent, $\mu_\theta(S) = 1$. Otherwise, if the incumbent fails to manage the crisis, then the voter learns she is not competent, $\mu_\theta(F) = 0$.[6]

Endogenous Learning on the Voter's Ideal Policy

Next, we analyze voter learning on the policy dimension. In our setting, the voter does not know if reforms to the left or the right of the status quo are optimal for her, i.e., if $\omega = \alpha$ or $\omega = -\alpha$. Faced with this uncertainty, she looks back at how much she liked (or disliked) the implemented policy, and updates her beliefs. Her inference problem, however, is complicated by the fact that her realized utility (i.e., the policy outcome) is also a function of the idiosyncratic shock ϵ_1. Thus, the voter's policy-payoff realization is a *noisy* signal of the state of the world.

Denote the voter's posterior belief that $\omega = \alpha$, i.e., right-wing reforms are optimal for her, as $\mu_\omega(x_1, o_1^p)$. By Bayes' rule we have the following:

Lemma 1. *Assume the incumbent chooses policy x_t and this yields utility $U_t^p = o_t^p$ to the voter. The voter's posterior belief over ω is given by:*

$$\mu_\omega(x_t, o_t^p) = \frac{\gamma\phi\big(o_t^p - o_{t-1}^p - \alpha(x_t - x_{t-1})\big)}{\gamma\phi\big(o_t^p - o_{t-1}^p - \alpha(x_t - x_{t-1})\big) + (1-\gamma)\phi\big(o_t^p - o_{t-1}^p + \alpha(x_t - x_{t-1})\big)}$$

Upon observing policy outcome o_t^p, the voter considers the likelihood of such a realization emerging under one or the other state of the world. Formally, this is the likelihood that the noise ϵ_t has realization $o_t^p - E[o_t^p|\omega = \alpha]$ or $o_t^p - E[o_t^p|\omega = -\alpha]$.

Lemma 1 highlights two crucial property of the voter learning process. First, the higher the utility provided by a policy, the more likely the voter is to believe such policy moved the status quo

[6]The assumption that a crisis fully reveals the incumbent's competence is imposed for simplicity. Similar results hold if the crisis outcome is a noisy signal of politician competence, as long as a success or failure is sufficiently informative about competence.

in the optimal direction. Because the noise distribution satisfies the Monotonic Likelihood Ratio Property, μ_ω is increasing in o_1^p when $x_t > x_{t-1}$, and decreasing in o_1^p otherwise.

Second, even fixing the policy outcome o_1^p, the inferences that the voter draws depend on the implemented policy. It is easy to see that $|2\alpha(x_t - x_{t-1})|$ increases as x_t moves away from the status quo in either direction. In other words, as x_t moves away from the status quo, the utility distributions conditional on the state ω move farther apart. As a consequence, the voter is better able to filter out information from noise. Substantively, suppose that the incumbent implements a bold liberal reform, moving the status quo far to the left. If this policy produces a good outcome for the voter, then it is likely that the reform was in the optimal direction. Conversely, because policy outcomes are noisy and thus the voter's learning is imperfect, the consequences of a policy close to the status quo are much less informative. Thus, the *amount* of voter learning increases under bolder policy reforms.

Figure 1: Policy and learning

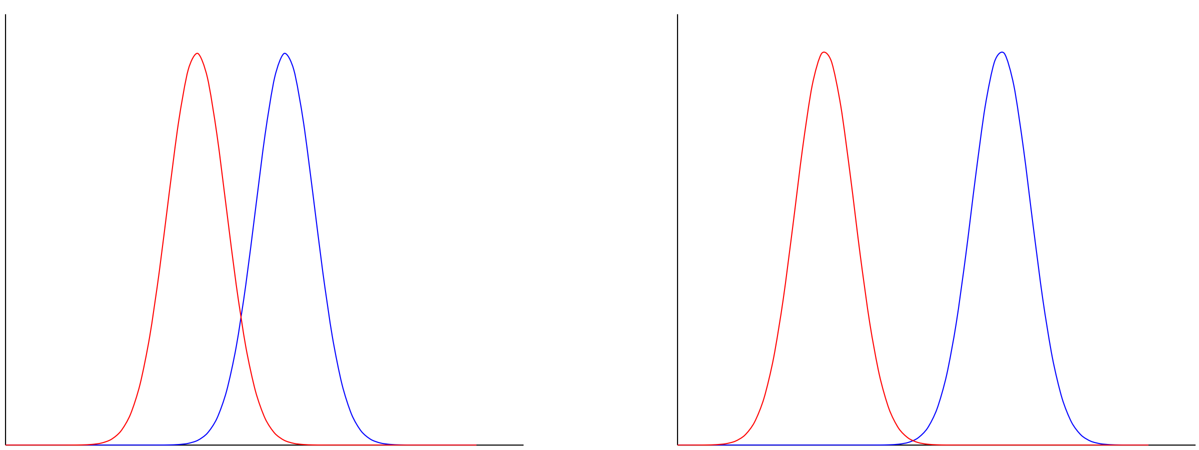

Note: Figure 1 depicts the effect of moving policy away from the status quo on the signal. In each graph, the red curve represents the policy outcome distribution under $\omega = -\alpha$, and the blue curve the distribution under $\omega = \alpha$. The left graph fixes a policy $x_t' > x_{t-1}$, and the right one a policy $x_t'' > x_t'$.

We adopt the following definition of informativeness:

Definition 1. *One signal (i.e., a mapping from the unobserved state to policy outcomes) is more Blackwell informative than another if the distribution of posteriors under the first signal is a mean-preserving spread of the distribution of posteriors under the second signal.*

With this definition, Lemma 2 formalizes the discussed relationship between policy reforms and informativeness:

Lemma 2. *Outcomes are more Blackwell-informative as x_t moves away from the status quo in either direction.*

This section emphasized that, while the possibility to learn on the competence dimension arises exogenously, the amount of learning on the policy dimension is endogenous to the officeholder's choice. In other words, while a crisis *imposes* a risk on the incumbent, he is instead able to *choose* how much risk to take in policymaking. Our goal in this paper is to study how these two processes interact, i.e., how does the *exogenous* risk of a crisis influence incumbent's incentives to take *endogenous* risks in policymaking?

Equilibrium Analysis

Having analyzed voter learning, we now characterize equilibrium behavior. As usual, we proceed by backwards induction, starting from the second-period officeholder's policy choice.

Absent any electoral incentives, whoever is in office in the second (and last) period simply implements his ideologically preferred policy, regardless of whether a crisis emerges or not. Thus, if I is re-elected, then $x_2^* = x_I$. Otherwise, $x_2^* = x_C$.

When making her retention decision, the voter therefore faces a selection problem. Here, her problem is two-fold. She wants to avoid the potential cost of a mismanaged future crisis, and therefore prefers to elect the candidate who is most likely to be competent. At the same time, anticipating that whoever is elected follows his own ideological preferences in policymaking, she wants to select the candidate whose ideal policy provides her with the highest expected utility. Thus, in equilibrium she reelects the incumbent if and only if:

$$\mathbb{E}_\omega[U_2^p(x_I)|x_1, o_1^p] + \mathbb{E}_{\theta_I}[U_2^c|\chi] \geq \mathbb{E}_\omega[U_2^p(x_C)|x_1, o_1^p] + \mathbb{E}_{\theta_C}[U_2^c]$$

The voter's retention decision is thus a function of her posterior beliefs on both the state of the world, $\mu_\omega(x_1, U_1^p)$, and the incumbent's ability, $\mu_\theta(\chi)$. Our next lemma characterizes this decision.

Lemma 3. *Assume the crisis outcome is χ. If $\mu_\omega \geq \overline{\gamma}_\chi$ then the voter reelects the right-wing incumbent, otherwise, the voter elects the challenger, where*

$$\overline{\gamma}_\chi = \frac{1}{2} - \frac{\xi_\chi p_b K(\mu_\theta(\chi) - \pi_C)}{2\alpha x_I}.$$

The intuition for this reelection rule is straightforward. Suppose $\omega = \alpha$. In this case, the voter's expected second-period policy utility is $o_1^p + \alpha(x_2 - x_1)$. Thus, (regardless of the location of the status quo or the first-period policy outcome) right-wing reforms provide her with higher utility than left-wing ones. Likewise, the opposite holds if $\omega = -\alpha$. Therefore, the voter reelects the right-wing incumbent if and only if her posterior belief that $\omega = \alpha$ (i.e., μ_ω) is sufficiently high.

Further, notice that the reelection threshold $\bar{\gamma}_\chi$ is decreasing in μ_θ. In other words, the higher the voter's posterior on the incumbent's competence, the more lenient she becomes on the ideology dimension (so that she is willing to reelect the right-wing incumbent even if she believes policies to the left of the status quo more likely to be optimal for her). Moreover, this effect is amplified by the weighted cost of a crisis $\xi_\chi p_b K$: the higher the risk and/or expected severity of future crises, the more the voter cares about competence when making her retention decision.

Moving forward, we make the following assumption:

Assumption 1. $\xi > \frac{1-\pi_C}{\pi_C}$ *and* $p_b \in \left(\frac{\alpha x_I}{\xi K \pi_C}, \min\{\frac{\alpha x_I}{K\pi_C}, \frac{\alpha x_I}{K(1-\pi_C)}\}\right)$.

The lower bounds on ξ and p_b imply that when an incumbent fails to solve a crisis, and thus reveals himself as a bad type, the likelihood that the crisis continues to the next period and/or its severity are sufficiently large that the voter prefers to replace him with the challenger (whose probability of being competent is larger than zero). In particular, this holds regardless of the voter's (posterior) beliefs over the location of the optimal policy. Substantively, a *failed* crisis today increases the salience of the competence dimension, and makes ideology electorally irrelevant (i.e., $\bar{\gamma}_F > 1$). In contrast, because the baseline probability of a crisis in the second period (p_b) is low, proving himself as a good type by solving a crisis is not enough for the incumbent to always get reelected (i.e., $\bar{\gamma}_S > 0$).

Most of our results continue to hold if the probability of winning after a failed crisis is interior. In our analysis below, we highlight when this is not the case and briefly discuss how the relevant results change under alternative assumptions.

Equilibrium Policy

Moving backwards, we now consider the incumbent's optimal policy choice. In what follows, we assume that $\gamma > \frac{1}{2}$. This simplifies notation and facilitates the presentation of our results, but it does not impact our qualitative conclusions. In Appendix B, we re-state and prove all our results allowing for $\gamma < \frac{1}{2}$. For simplicity, we set the status quo policy at the beginning of the game to 0 (i.e., $x_0 = 0$). Notice that this also implies that the ideal policies of the incumbent and challenger are on opposite sides of the status quo. We often refer to policies farther from the status quo as more extreme and polices closer to the status quo as more moderate.

Lemma 3 highlighted that the voter's retention decision is a function of her posterior belief on the policy dimension. From Lemmas 1 and 2, we know that the distribution of this posterior belief depends on the implemented x_1: the farther this policy is from the status quo, the larger is the amount of voter learning in equilibrium. Thus, the incumbent's policy choice will be a function of his own ideological preferences, x_I, *and* his incentives to either prevent or facilitate voter learning on the policy dimension.

Let $P_\chi(x)$ be the incumbent's probability of winning if he chooses policy x and the first-period crisis outcome is χ. Then, the incumbent's expected utility for policy x in normal times is given by

$$U_n(x) = -(x - x_I)^2 - (x_I - x_C)^2 \Big[1 - P_N(x)\Big] + \beta P_N(x).$$

Instead, his expected utility for policy x during times of crisis is

$$U_c(x) = -(x - x_I)^2 - (x_I - x_C)^2 \Big[1 - \pi_I P_S(x) - (1 - \pi_I) P_F(x)\Big] + \beta \Big[\pi_I P_S(x) + (1 - \pi_I) P_F(x)\Big].$$

With these in hand, we can now characterize the incumbent's equilibrium policy choice. In what follows, we maintain the following assumption:

Assumption 2. *Office rents are sufficiently large, $\beta > \widehat{\beta}$.*

Substantively, this guarantees that dynamic electoral incentives are sufficiently strong that they can dominate over the incumbent's static ideological considerations. The $\widehat{\beta}$ is characterized in the Appendix (in the proof of Lemma A5).[7]

Lemma 4. *Let x_c^* the equilibrium policy under a crisis and x_n^* denote the equilibrium policy in normal times. In times of crisis the equilibrium policy x_c^* solves*

$$2(x_I - x) + (\beta + 4x_I^2)\Big(\pi_I \frac{\partial P_S}{\partial x} + (1 - \pi_I)\frac{\partial P_F}{\partial x}\Big) = 0.$$

In normal times the equilibrium policy x_n^* solves

$$2(x_I - x) + (\beta + 4x_I^2)\frac{\partial P_N}{\partial x} = 0.$$

As discussed above, the first order condition highlights that the implemented policy directly influences the incumbent's probability of being reelected (via the voter learning). Given the symmetry in our setup, for any pair of policies x and $-x$ we have that $P_\chi(x) = P_\chi(-x)$. This implies that, in equilibrium, a right-wing incumbent will never implement a policy $x < 0$. Therefore, if the second term in the FOC is negative, so that more extreme reforms increase his probability of winning in this range, then the incumbent implements a policy to the right of his ideological preferences, $x^* > x_I$. Otherwise, the incumbent distorts his policy choice away from his static optimum and towards the status quo, $x^* < x_I$.

Thus, a crucial first step in our analysis is characterizing the conditions under which the incumbent's probability of winning is increasing as x moves away from the status quo, i.e., bolder policy

[7] We note that, if $\gamma < \frac{3}{4}$, so the incumbent is not too advantaged on the policy dimension, then Assumption 2 always holds, $\widehat{\beta} = 0$.

reforms are electorally beneficial, and when instead it is decreasing, i.e., policies closer to the status quo are electorally optimal. Let $\tilde{\pi}$ be the value of π_I at which, under normal times, the voter is ex-ante indifferent between incumbent and challenger. Specifically, $\tilde{\pi} = \pi_C + \frac{\alpha x_I}{p_b K}(1-2\gamma)$. We make the following assumption to guarantee that $\tilde{\pi} \in (0, 1)$:

Assumption 3. $\gamma < \frac{1}{2} + \frac{\pi_C}{2}\frac{p_b K}{\alpha x_I}$.

We can now describe how the incumbent's probability of winning changes in x.

Lemma 5. *Suppose $x > 0$.*

1. *Under a crisis, I's ex-ante probability of winning is decreasing in x.*
2. *In normal times, the following holds:*
    - *If $\pi_I > \pi_C$ then I's probability of winning is decreasing in x;*
    - *If $\pi_I < \tilde{\pi}$ then I's probability of winning is increasing in x;*
    - *If $\pi_I \in (\tilde{\pi}, \pi_C)$, then I's probability of winning is decreasing then increasing in x.*

Symmetric results hold for $x < 0$.

Suppose a crisis arises in the first period. If the incumbent fails to solve it, then he is always ousted. Hence, changing x does not alter his probability of winning in the event of a failure. In contrast, solving the crisis secures an electoral advantage for the incumbent. In this case, new policy-relevant information can therefore only hurt his reelection chances. In particular, the voter reelects a successful incumbent absent new information on the policy dimension. Thus, in a crisis, the incumbent's ex-ante probability of winning is decreasing in the amount of voter learning, i.e., decreasing in x.

Instead, suppose a crisis does not materialize. For a logic analogous to the one above, the probability of winning is always decreasing in x when the incumbent is sufficiently far ahead ex-ante, $\pi_I > \pi_C$. The opposite reasoning holds when he is sufficiently far behind, $\pi_I < \tilde{\pi}$. Since the incumbent needs to generate information on the policy dimension to win, the probability of winning is increasing in x. Now assume that $\pi_I \in (\tilde{\pi}, \pi_C)$, so the election is more competitive. Here, the incumbent is less competent than the challenger in expectation, but nonetheless ex-ante ahead electorally (due to the voter's prior on the policy dimension γ). If no information at all is generated, the ex-ante advantaged incumbent wins for sure. Thus, his probability of winning is initially decreasing as x increases away from the status quo. However, since $\gamma > \frac{1}{2}$, new policy-relevant information is more likely to reveal to the voter that she prefers the incumbent to the challenges. As x continues moving away from zero, this second effect dominates, and $\frac{\partial P_N}{\partial x}$ takes a positive sign.

How Crises Impact Policymaking

We can now characterize how the crisis influences the incumbent's equilibrium policy choice. Specifically, does the crisis alter the direction in which the incumbent distorts policy away from his static optimum (towards or away from the status quo)? Does it alter the magnitude of this distortion? What is the overall effect on policymaking? To answer these question, we introduce the following definition:

Definition 2. *We say the incumbent gambles if* $x^* > x_I$.

Thus, the incumbent gambles if he chooses a policy that is further from the status quo than his ideal point. We use this terminology to reflect the fact that more extreme reforms, by generating more policy-relevant information, entail more uncertainty for the incumbent.

The next result follows directly from Lemma 5:

Lemma 6. *If there is a crisis, then the incumbent never gambles in equilibrium,* $x_c^* \leq x_I$.

When a crisis arises in the first period, the incumbent is always ousted in the event he fails to solve it. Otherwise, if he successfully solves it, we know from Lemma 5 that his probability of winning is decreasing in x. The incumbent does not know for sure if he will be able to solve a crisis or not, but he can anticipate that policy information is either electorally irrelevant, or hurts his retention chances. This implies that he has incentives to prevent voter learning and, thus, always implements a policy that is closer to the status quo than his static optimum.[8]

A straightforward implication is that the incumbent's incentives for moderation get stronger as office rents increase.

Corollary 1. *Suppose there is a crisis. The equilibrium policy reform is always decreasing in office benefit,* $\frac{\partial x_c^*}{\partial \beta} < 0$.

Next, we consider the case in which no crisis arises.

Lemma 7. *Suppose there is no crisis. If* $\pi_I < \tilde{\pi}$ *then the incumbent always gambles in equilibrium. Otherwise, if* $\pi_I > \tilde{\pi}$ *then the incumbent never gambles.*

Recall that at $\pi_I = \tilde{\pi}$ the voter is ex-ante indifferent between retaining the incumbent or replacing him with the challenger. When $\pi_I > \tilde{\pi}$, the incumbent is electorally leading, and this creates incentives to prevent voter learning about policy in order to preserve his advantage. Thus,

[8] If we relax Assumption 1, so that even an incumbent who fails to solve a crisis is retained with positive probability (ex-ante), then we would have that under a crisis the incumbent gambles if and only π_I is sufficiently low. This is because information hurts the incumbent conditional on success, but helps him conditional on failure. The net effect in expectation thus depends on the probability of the incumbent successfully managing the crisis, i.e., π_I.

Figure 2: Policymaking in times of crisis

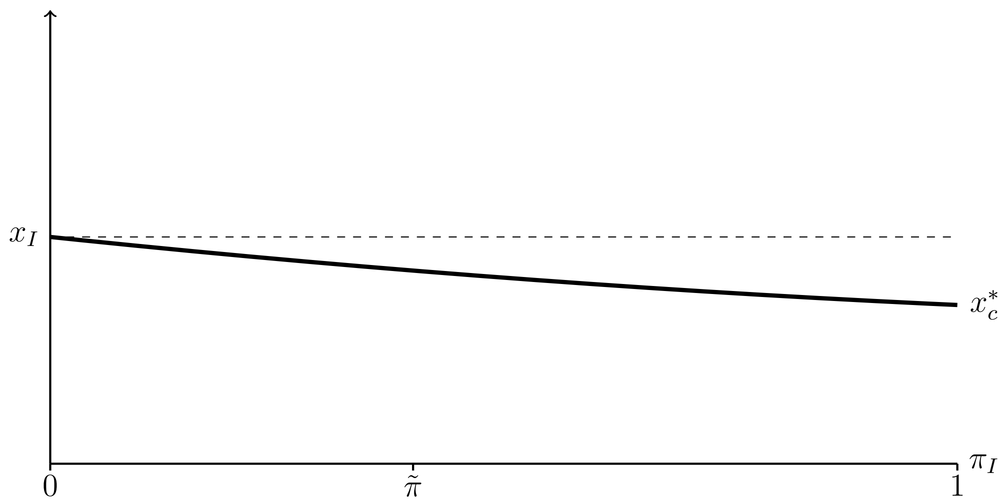

Note: Figure 2 depicts the incumbent's equilibrium policy choice in times of crisis, x_c^*, as a function of the probability he is competent.

he implements a reform that is closer to the status quo than his ideal point. In contrast, if π_I is lower than $\tilde{\pi}$, then the incumbent needs to generate a significant shock to the voter's beliefs to have any hope of winning reelection. Therefore, he wants to experiment with extreme policies farther from the status quo.

Figure 3: Policymaking in normal times

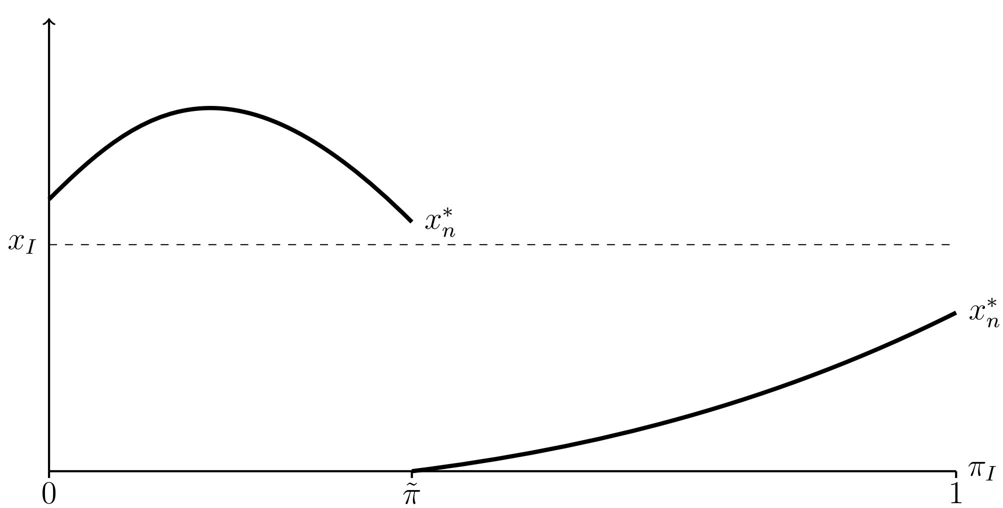

Note: Figure 3 depicts the incumbent's equilibrium policy choice, x_n^*, as a function of the probability he is competent, when there is not a crisis.

As in times of crisis, the incumbent's incentives to control information get stronger as his reelection motives intensify:

Corollary 2. *Suppose there is no crisis. If $\pi_I < \tilde{\pi}$, then the equilibrium policy choice is increasing in office benefit, $\frac{\partial x_n^*}{\partial \beta} > 0$. Otherwise, if $\pi_I > \tilde{\pi}$, then the equilibrium policy choice is decreasing in office benefit, $\frac{\partial x_n^*}{\partial \beta} < 0$.*

Comparing the equilibrium policy under crisis to policy during normal times, we show that the crisis has three effects on reform.

First, from Lemmas 6 and 7 we can immediately see that a crisis can impact the incumbent's decision to gamble or not in equilibrium. We refer to this as the *qualitative* effect of the crisis.

Proposition 1. *If $\pi_I < \tilde{\pi}$, then the crisis has a qualitative effect on policymaking: I gambles in normal times, but does not gamble under a crisis, $x_c^* < x_I < x_n^*$. If $\pi_I > \tilde{\pi}$, then the crisis has no qualitative effect, $x_c^*, x_n^* < x_I$.*

As discussed earlier, if there is a crisis then the incumbent is always incentivized to act *as if* he is ahead, since he is always removed after failing the crisis. Therefore, the crisis does not change the incumbent's decision to not gamble when $\pi_I > \tilde{\pi}$. On the other hand, for $\pi_I < \tilde{\pi}$, the trailing incumbent wants to facilitate voter learning during normal times, but wants to prevent information generation if he experiences a crisis. Here, the crisis qualitatively impacts the incumbent's incentives to control information and, consequently, his policy choice. Exogenous crises thus induce low (expected) quality incumbents to distort policy towards the status quo rather than engage in extreme policy experiments.

Next, we consider incumbents that do not experience a qualitative effect, i.e., for whom $\pi_I > \tilde{\pi}$. For these policymakers, we identify a *quantitative* effect of exogenous crises: the intensity of their incentives to control information, and thus the amount by which they are willing to move away from their ideological preferences, changes under a crisis.

Proposition 2. *If $\pi_I > \tilde{\pi}$ then the crisis always dampens incentives to control information, i.e., $|x_c^* - x_I| < |x_n^* - x_I|$.*

When the incumbent is ex-ante leading, $\pi > \tilde{\pi}$, he has incentives to prevent voter learning on policy both in normal times and in a crisis. However, if the incumbent succeeds in solving the crisis he obtains an even greater electoral advantage due to the voter learning he is competent and, in turn, becoming more lenient on the policy dimension. If he fails, policy information will be electorally irrelevant. Overall, this implies that the crisis *dampens* incentives to control information. The incumbent is thus relatively more willing to take more risks on the policy dimension during times of crisis, and moves the equilibrium policy closer to the his static optimum. Substantively, a leading incumbent engages in bolder policy reforms under a crisis.

Figure 4: Crisis vs. no crisis

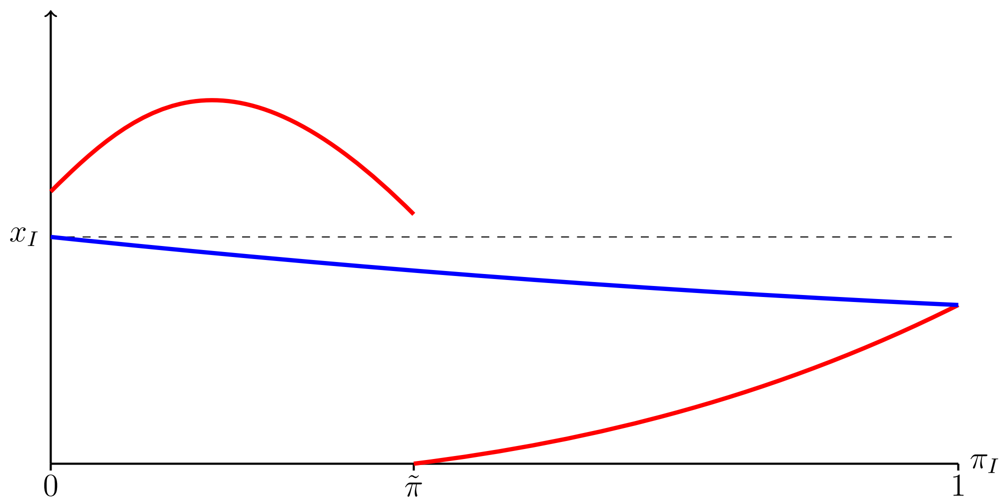

Note: Figure 4 compares equilibrium policy in normal times, x_n^* (depicted in red) against the equilibrium policy in times of crisis, x_c^* (depicted in blue).

Next, we show that crises may have an *indirect* effect on policy making, by influencing how the incumbent's optimal choice responds to changes in his expected ability.

Proposition 3.

1. *The equilibrium crisis policy, x_c^*, is decreasing in π_I.*

2. *There exists a cut-point $\underline{\pi}_I < \tilde{\pi}$, such that: if $\pi_I \in (\underline{\pi}_I, \tilde{\pi})$ then the equilibrium policy under no crisis, x_n^*, is decreasing in π_I. Otherwise, x_n^* is increasing in π_I.*

First, suppose the country experiences a crisis. As the incumbent becomes ex ante more likely to be competent the *probability* that he will acquire an ex-post electoral advantage increases. In turn, this strengthens his incentives to preserve this advantage by preventing learning on the policy dimension. The equilibrium crisis policy thus decreases (moves closer to the status quo) as π_I increases.

During normal times, however, the incumbent's expected competence plays a very different role. As π_I increases, the voter becomes more lenient towards the incumbent on the policy dimension (i.e., the reelection threshold $\bar{\gamma}$ decreases). Thus, π_I changes the *magnitude* of the incumbent's electoral advantage (or disadvantage). Here, increasing π_I has a qualitatively distinct impact on trailing incumbents and leading ones. First, consider a trailing incumbent, i.e. $\pi_I < \tilde{\pi}$. Increasing π_I has two competing effects. The incumbent's initial disadvantage shrinks, therefore the amount of (favorable) policy information he needs to generate in order to win reelection decreases. This weakens his incentives to experiment with extreme policies. At the same time, however, as the incumbent's disadvantage on the competence dimension decreases, policy information is more likely to be electorally relevant. This strengthens his incentives to control information and thus move to the extreme away from his bliss point.

Our analysis shows that the first effect dominates when π_I is close to $\tilde{\pi}$, and the second dominates when π_I is close to 0. Thus, x_n^* is decreasing in π_I when elections are very competitive ($\pi_I \in (\underline{\pi}_I, \tilde{\pi})$), and increasing in π_I when the incumbent is very far behind ($\pi_I < \underline{\pi}_I$).

In contrast, when the incumbent is leading ($\pi_I > \tilde{\pi}$), the two effects always go in the same direction. As π_I increases, the amount of (negative) policy information the incumbent can generate while still winning reelection increases. This makes the incumbent less risk averse. Simultaneously, policy information becomes less electorally relevant because the voter becomes more and more lenient. This weakens incentives to control information. Both effects therefore push x_n^* to the extreme, away from the status quo and towards the incumbent's static optimum.

Thus, when the incumbent is ex-ante leading or very far behind, the equilibrium policy is increasing in his expected ability during normal times, but decreasing during times of crisis.

Finally, combining all of the above, we can characterize the crisis overall directional effect on equilibrium policy:

Proposition 4. *If $\pi_I > \tilde{\pi}$, then $x_c^* > x_n^*$. Otherwise, if $\pi_I < \tilde{\pi}$, then $x_c^* < x_n^*$.*

This proposition highlights that crises induce more policy extremism (and thus experimentation) when the incumbent is sufficiently likely to be competent, and induce moderation otherwise. This result thus qualifies the conventional wisdom that crises always represent an opportunity for reforms and policy experiments. We show that, when we take the incumbent's electoral incentives into account, the impact of a crisis can go in the opposite direction.

Further, our analysis highlights that exogenous crises also have a crucial impact on policy volatility. Formally, denote x_c^{max} as the most extreme policy sustainable in equilibrium in times of crisis, for any π_I, and x_c^{min} the least extreme. Similarly define x_n^{max} and x_n^{min}. This yields:

Corollary 3. $x_n^{max} - x_n^{min} > x_c^{max} - x_c^{min}$.

As is obvious from an analysis of the indirect effect, the range of policies that the incumbent may implement in equilibrium (under different values of the primitives) is larger in times of no crisis.

Our results offer several avenues for future research. First, they highlight the importance of the political environment for understanding how crises impact reform. Although earlier theories imply that crises should always increase policy experimentation (e.g., Tommasi and Velasco, 1996), findings in the empirical literature are mixed. Some scholars do find that crises and reform are positively related (Lora and Olivera, 2004; Alesina, Ardagna and Trebbi, 2006; Pitlik and Wirth, 2003), however, others find crises may lead to less reform (Campos, Hsiao and Nugent, 2010; Pop-Eleches, 2008; Castanheira, Nicodème and Profeta, 2012; Galasso, 2014; Mian, Sufi and Trebbi, 2014). Propositions 1 and 2 suggest that whether crises induce more or less reform is dependent on whether the incumbent already has an electoral advantage. Second, our prediction that policy volatility is lower during times of crisis (Corollary 3) provides additional observable implications that are specific to our model. Finally, our theory is unique in that it applies even to crises that are completely orthogonal to the policy dimension.

Endogenous Crises

In our baseline model, the incumbent's only tool to improve his reelection prospects is to control information on the policy dimension, by either gambling with bold experiments or implementing uninformative platforms closer to the status quo. Within this setting, we have considered how officeholders react to crises that *exogenously* impose a risk on them.

Suppose instead that, in addition to controlling policy, the incumbent could also decide whether to take risks on the competence dimension. In other words, suppose the incumbent could *endogenously* generate a crisis, so as to provide voters with a test of his ability. When (if ever) would he choose to do so?

In what follows, we address this question. We will be working under the assumption that the office rents β are sufficiently large that the incumbent's dynamic electoral incentives dominate ideological considerations, so that the equilibrium policy converges to the one that maximizes his probability of winning. Under this assumption, we have:

Proposition 5. *Suppose office benefit is sufficiently large. If $\pi_I \in [\gamma, \tilde{\pi}]$, then endogenous crises emerge in equilibrium. Otherwise, if $\pi_I \notin [\gamma, \tilde{\pi}]$, then endogenous crises never emerge.*

The intuition for this result is straightforward. Suppose that $\pi_I > \tilde{\pi}$. In this case, the incumbent is ex-ante leading and, absent any new information, will be reelected. This incumbent therefore never has any reason to gamble, whether on the policy or the crisis dimension. In contrast, when $\pi_I < \tilde{\pi}$, the trailing incumbent needs the voter to learn on at least one of these dimensions in order to win the upcoming election. Here, the incumbent has two options: he can gamble on the competence dimension by generating a crisis, but play it safe on policy, or he can engage in bold policy experiments, but avoid the risk of exposing himself as an incompetent. The optimal choice depends on the relative probability of the two gambles succeeding. Recall that γ is the probability that the voter's optimal direction of reform is aligned with the incumbent's ideology. If $\pi_I < \gamma$, then gambling on the policy dimension is more likely to succeed. The incumbent therefore never chooses to generate a crisis, in order to maximize the probability that new policy-relevant information is enough to push him above the retention threshold. Instead, when $\pi_I > \gamma$ the incumbent is more likely to succeed in solving a crisis than he is to convince the voter that her ideal policy is a right-wing one. In this range, he therefore finds it optimal to induce a crisis in order to improve his retention chances. Our results therefore highlight that, ironically, the incumbent will choose to generate an endogenous crisis only when he is not too likely (nor unlikely) to be able to solve it.

Asymmetric Information

In this section, we allow the incumbent to have private information about his ability. Introducing such asymmetry of information has a crucial implication in our setting: the incumbent's policy choice may influence the voter's beliefs about his ability. Our goal in this section is to show that the qualitative results from Propositions 1-4 in the main body survive even in this richer information setting.

Formally, we assume that at the beginning of the game the incumbent observes a private signal of his own ability, given by $s_\theta \in \{0, 1\}$. Let $p(s_\theta = 1|\theta = 1) = p(s_\theta = 0|\theta = 0) \in (1/2, 1]$. After observing the signal, the incumbent updates his beliefs about his own ability and the game proceeds as in the main model. Denote ψ_s as the incumbent's (interim) posterior that he is a good type, conditional on the realization of his private signal, where $0 \leq \psi_0 < \psi_1 \leq 1$. Notice that we allow the incumbent's signal to be either partially or fully informative about his ability.

In what follows, we denote $x_n^b(\pi_I)$ as the equilibrium policy under normal times in the baseline symmetric uncertainty model, and $x_n^a(s_\theta)$ the equilibrium policy choice of the incumbent after observing the signal s_θ in the asymmetric information setting. Similarly, let $x_c^b(\pi_I)$ and $x_c^a(s_\theta)$ denote the equilibrium policies in times of crisis. Finally, let $\mu_\theta(x_1)$ be the voter's interim beliefs about the incumbent's ability after observing his policy choice.

First, we show that the equilibrium is always separating in times of crisis.

Lemma 8. *Suppose there is a crisis. In any perfect Bayesian equilibrium, we always have that the incumbent chooses different policies following each signal, $x_c^a(0) \neq x_c^a(1)$. Furthermore, $x_c^a(0) = x_c^b(\pi_I = \psi_0)$ and $x_c^a(1) = x_c^b(\pi_I = \psi_1)$.*

If a crisis emerges in the first period, the incumbent's type is always revealed via the crisis outcome. As a consequence, the voter's interim posterior $\mu_\theta(x_1)$ is electorally irrelevant.[9] Therefore, the incumbent's policy choice influences his reelection chances only via experimentation and the

[9]Here we are assuming that, if the voter's interim posterior and the posterior conditional on crisis outcome are incompatible, the voter only updates conditioning on o_1^c. Specifically, if the voter forms interim posterior that the incumbent is competent but then observes a failed crisis, we assume that she reaches final beliefs that $\theta_I = 0$ with certainty.

voter learning on the policy dimension. The incumbent's strategic problem is therefore identical to the one he faces in the baseline model. Thus, in equilibrium he always acts *as if* there was no asymmetry of information between him and the voter, and implements the dynamically optimal policy given his interim posterior ψ_s.

A corollary follows straightforwardly:

Corollary 4. *The expected policy in times of crisis is decreasing in* π_I.

As in the symmetric uncertainty baseline, more competent (in expectation) officeholders implement more moderate reforms, everything else being equal.

Next, we show that under normal times the incumbent can never do better than in the symmetric uncertainty model. First, we verify that there always exists an equilibrium where both types of the incumbent pool on $x_n^b(\pi_I)$, the optimal policy in the baseline model without asymmetric information.

Lemma 9. *There always exists an equilibrium where, regardless of his private signal, the incumbent adopts* $x_n^b(\pi_I)$.

The intuition is as follows. When a crisis does not materialize, fixing the voter's interim posterior $\mu_\theta(x_1)$, the incumbent's dynamically optimal policy is not a function of his own beliefs over θ_I. Thus, the optimal policy does not depend on the incumbent's private signal. Notice that this immediately implies that the usual beliefs refinements (intuitive criterion, D1, etc.) do not have bite here. Next, suppose that, following a deviation off the equilibrium path, the voter forms beliefs that $s_\theta = 0$. It is straightforward to see that the conjectured equilibrium always exists under this assumption.

Next, we show that this equilibrium is the one that yields both types their highest expected utility. As a first step, we establish an indifference-based separation result:

Lemma 10. *In any separating equilibrium, both types are indifferent between policies on the equilibrium path.*

As discussed 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 the case that the incumbent is always indifferent between the policies on the equilibrium path.

Further, notice that, as it is always the case in such signalling games, in any separating equilibrium an incumbent who observes $s_\theta = 0$ must be locating at his dynamically optimal policy from the baseline model.

Our result then follows from a standard envelope theorem argument:

Proposition 6. *Suppose no crisis emerges in the first period. Among all perfect Bayesian equilibria, the one where the incumbent adopts* $x_n^a(0) = x_n^a(1) = x_n^b(\pi_I)$ *maximizes the expected utility of both types of the incumbent.*

Thus, there is no equilibrium in which the incumbent can do better than the one where he ignores his private information, even if he learns he is the competent type. An important implication follows immediately. If we focus on the equilibrium that provides the incumbent with the highest expected utility, asymmetric information has no impact on the equilibrium policy during normal times. This, together with Corollary 4 showing that under a crisis the policy is decreasing in π_I, implies that our qualitative results from Propositions 1-4 continue to hold. Figure 5 illustrates this robustness. Here, the red curve is exactly as in the baseline, as it plots the incumbent's policy from the welfare-maximizing equilibrium under normal times, $x_n^a(0) = x_n^a(1) = x_n^b(\pi_I)$. The blue dashed curves are the equilibrium policies in times of crisis, under both possible realizations of the private signal s_θ. Finally, the blue solid line is the ex-ante expected policy (where the expectation is over s_θ). As in the baseline, we can see the qualitative, quantitative, and indirect effects of a crisis clearly at play in our setting.

Figure 5: Policymaking with asymmetric information

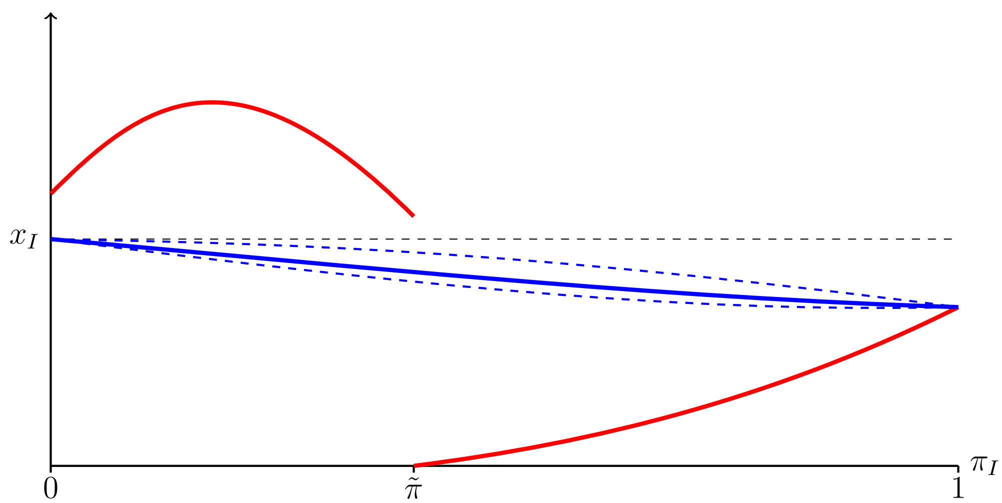

Note: Figure 5 compares equilibrium policy in normal times (red line) against the equilibrium policy in times of crisis following the $s_\theta = 0$ signal (upper dashed blue line), following the $s_\theta = 1$ signal (lower dashed blue line), and the expected crisis policy (solid blue line).

Conclusion

How do crises influence an executive's willingness to implement policy reforms? While existing work focuses on how crises impact voters' *demand* for reform, this paper investigated how they alter politicians' incentives to *supply* policy experimentation, even if the crisis does not shift voters' policy preferences. To study this problem, we developed a model of elections and policy experimentation, where voters face uncertainty about their optimal policy and politicians' ability to manage a crisis. We showed that extreme reforms generate more information for voters about their optimal platform. Consequently, the incumbent has electoral incentives to engage in information control. At the same time, a crisis represents an exogenous test for the incumbent, who must prove competent enough to successfully manage the country during turbulent times. Therefore, a crisis has an (independent) impact on the incumbent's electoral prospects, and this may influence his incentives to engage in risky policy experiments.

Our results showed that crises may either have a qualitative effect on policymaking, inducing officeholders who would would otherwise engage in policy gambles to play it safe, or a quantitative one, altering the intensity but not the nature of the incumbent's incentives to control voter learning. Further, crises may alter how policies respond to changes in the features of the electoral environment. Combining these effects we found that, in contrast to the conventional wisdom, a crisis induces bolder policy reforms only when the incumbent is sufficiently likely to be competent. If the incumbent is relatively unlikely to be competent, then the crisis instead results in policies that are closer to the status quo.

Appendix

Lemma 2. *Outcomes are more Blackwell-informative as x_t moves away from the status quo in either direction.*

Proof. Notice that, in our setting, the noise distribution satisfies the MLRP property. Further, 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

$$\alpha(z - x_{t-1}) + \alpha(z - x_{t-1}) > \alpha(y - x_{t-1}) + \alpha(y - x_{t-1})$$

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

Lemma 3. *Assume the crisis outcome is χ. If $\mu_\omega \geq \overline{\gamma}_\chi$ then the voter reelects the right-wing incumbent, otherwise, the voter elects the challenger, where*

$$\overline{\gamma}_\chi = \frac{1}{2} - \frac{\xi_\chi p_b K(\mu_\theta(\chi) - \pi_C)}{2\alpha x_I}.$$

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

$$\begin{aligned} o_1^p + \mu_\omega \alpha(x_I - x_1) - (1 - \mu_\omega)\alpha(x_I - x_1) - (1 - \mu_\theta(\chi))\xi_\chi p_2 K > \\ o_1^p + \mu_\omega \alpha(x_C - x_1) - (1 - \mu_\omega)\alpha(x_C - x_1) - (1 - \pi_C)\xi_\chi p_2 K. \end{aligned}$$

Substituting $x_C = -x_I$, the above reduces to

$$(2\mu_\omega - 1)\alpha 2x_I > \xi_\chi p_2 K(\pi_C - \mu_\theta(\chi)),$$

and

$$\mu_\omega > \frac{1}{2} - \frac{\xi_\chi p_2 K(\mu_\theta(\chi) - \pi_C)}{2\alpha x_I}$$

□

Lemma (A1). *If the crisis outcome is $\chi \in \{N, S\}$, then the probability of reelection for the incumbent is*

$$P_\chi(x_1) = \gamma\Big(1 - \Phi(\frac{\lambda_\chi}{2\alpha|x_1|} - \alpha|x_1|)\Big) + (1-\gamma)\Big(1 - \Phi(\frac{\lambda_\chi}{2\alpha|x_1|} + \alpha|x_1|)\Big),$$

where $\lambda_\chi = \ln\left(\frac{(1-\gamma)\overline{\gamma}_\chi}{\gamma(1-\overline{\gamma}_\chi)}\right)$. If $\chi = F$ then the probability of reelection for the incumbent is $P_F(x) = 0$.

Proof. By Bayes rule we have

$$\mu_\omega = \frac{\gamma\phi\big(o_t^p - o_{t-1}^p - \alpha(x_t - x_{t-1})\big)}{\gamma\phi\big(o_t^p - o_{t-1}^p - \alpha(x_t - x_{t-1})\big) + (1-\gamma)\phi\big(o_t^p - o_{t-1}^p + \alpha(x_t - x_{t-1})\big)}$$

Thus, using Lemma 3, the incumbent's probability of being re-elected is

$$p\left(\frac{\gamma\phi\big(o_t^p - o_{t-1}^p - \alpha(x_t - x_{t-1})\big)}{\gamma\phi\big(o_t^p - o_{t-1}^p - \alpha(x_t - x_{t-1})\big) + (1-\gamma)\phi\big(o_t^p - o_{t-1}^p + \alpha(x_t - x_{t-1})\big)} > \overline{\gamma}_\chi\right),$$

where $\phi(\cdot)$ is the PDF of the standard normal.

From the incumbent's perspective, U_t is probabilistic, therefore the above can be rewritten as

$$\begin{aligned}
&\gamma[p\left(\frac{\gamma\phi\Big(o_{t-1}^p+\alpha(x_t-x_{t-1})+\epsilon_t-o_{t-1}^p-\alpha(x_t-x_{t-1})\Big)}{\gamma\phi\Big(o_{t-1}^p+\alpha(x_t-x_{t-1})+\epsilon_t-o_{t-1}^p-\alpha(x_t-x_{t-1})\Big)+(1-\gamma)\phi\Big(o_{t-1}^p+\alpha(x_t-x_{t-1})+\epsilon_t-o_{t-1}^p+\alpha(x_t-x_{t-1})\Big)} > \overline{\gamma}_\chi\right)] \\
&+(1-\gamma)[p\left(\frac{\gamma\phi\Big(o_{t-1}^p-\alpha(x_t-x_{t-1})+\epsilon_t-o_{t-1}^p-\alpha(x_t-x_{t-1})\Big)}{\gamma\phi\Big(o_{t-1}^p-\alpha(x_t-x_{t-1})+\epsilon_t-o_{t-1}^p-\alpha(x_t-x_{t-1})\Big)+(1-\gamma)\phi\Big(o_{t-1}^p-\alpha(x_t-x_{t-1})+\epsilon_t-o_{t-1}^p+\alpha(x_t-x_{t-1})\Big)} > \overline{\gamma}_\chi\right)].
\end{aligned} \tag{1}$$

Recall that we assume $x_0 = 0$. The above reduces to;

$$\begin{aligned}
&\gamma[p\left(\frac{\gamma\phi(\epsilon)}{\gamma\phi(\epsilon) + (1-\gamma)\phi(2\alpha x_1 + \epsilon)} > \overline{\gamma}_\chi\right)] \\
&+(1-\gamma)[p\left(\frac{\gamma\phi(-2\alpha x_1 + \epsilon)}{\gamma\phi(-2\alpha x_1 + \epsilon) + (1-\gamma)\phi(\epsilon)} > \overline{\gamma}_\chi\right)].
\end{aligned} \tag{2}$$

Standardizing $\phi(\cdot)$, we can rewrite this probability as

$$\gamma[p\left(\frac{\gamma e^{-\frac{\epsilon^2}{2}}}{\gamma e^{-\frac{\epsilon^2}{2}}+(1-\gamma)e^{-\frac{(2\alpha x_1+\epsilon)^2}{2}}}>\overline{\gamma}_\chi\right)]$$
$$+(1-\gamma)[p\left(\frac{\gamma e^{-\frac{(-2\alpha x_1+\epsilon)^2}{2}}}{\gamma e^{-\frac{(-2\alpha x_1+\epsilon)^2}{2}}+(1-\gamma)e^{-\frac{\epsilon^2}{2}}}>\overline{\gamma}_\chi\right)], \tag{3}$$

which further reduces to

$$\gamma[p(\gamma(1-\overline{\gamma}_\chi)e^{-\frac{\epsilon^2}{2}}>\overline{\gamma}_\chi(1-\gamma)e^{-\frac{(2\alpha x_1+\epsilon)^2}{2}})]$$
$$+(1-\gamma)[p(\gamma(1-\overline{\gamma}_\chi)e^{-\frac{-(-2\alpha x_1+\epsilon)^2}{2}}>\overline{\gamma}_\chi(1-\gamma)e^{-\frac{\epsilon^2}{2}})]. \tag{4}$$

After some rearrangement, we obtain

$$\gamma[p\left(e^{-\frac{\epsilon^2}{2}+\frac{(2\alpha x_1+\epsilon)^2}{2}}>\frac{\overline{\gamma}_\chi(1-\gamma)}{\gamma(1-\overline{\gamma}_\chi)}\right)]+(1-\gamma)[p\left(e^{-\frac{-(-2\alpha x_1+\epsilon)^2}{2}+\frac{\epsilon^2}{2}}>\frac{\overline{\gamma}_\chi(1-\gamma)}{\gamma(1-\overline{\gamma}_\chi)}\right)]. \tag{5}$$

If $x>0$, then after rearranging and apply a logarithmic transformation, the above obligingly reduces to

$$\gamma[p(\epsilon>\frac{\lambda_\chi}{2\alpha x_1}-\alpha x^1)]+(1-\gamma)[p(\epsilon>\frac{\lambda_\chi}{2\alpha x_1}+\alpha x^1)], \tag{6}$$

Finally, one more standardization yields

$$\gamma(1-\Phi(\frac{\lambda_\chi}{2\alpha x^1}-\alpha x^1))+(1-\gamma)(1-\Phi(\frac{\lambda_\chi}{2\alpha x^1}+\alpha x^1)), \tag{7}$$

where $\Phi(\cdot)$ is the CDF of the standard normal.

Suppose instead $x<0$, then the probability of winning is

$$\gamma\Phi\Big(\frac{-\lambda_\chi}{-2\alpha x_1}-\alpha x_1\Big)+(1-\gamma)\Phi\Big(\frac{-\lambda_\chi}{-2\alpha x_1}+\alpha x_1\Big)$$

Which can be rewritten as

$$\gamma\Big(1-\Phi\Big(\frac{\lambda_\chi}{-2\alpha x_1}+\alpha x_1\Big)\Big)+(1-\gamma)\Big(1-\Phi\Big(\frac{\lambda_\chi}{-2\alpha x_1}-\alpha x_1\Big)\Big)$$

Thus, for a generic x_1, we can express the probability of winning as

$$P_\chi(x_1) = \gamma\Big(1 - \Phi(\frac{\lambda_\chi}{2\alpha|x_1|} - \alpha|x_1|)\Big) + (1-\gamma)\Big(1 - \Phi(\frac{\lambda_\chi}{2\alpha|x_1|} + \alpha|x_1|)\Big),$$

That $P_F(x) = 0$ follows from the assumption that $p_b K > \frac{4\alpha x_I}{\xi \pi_C}$, which implies $\overline{\gamma}_F > 1$. Since $\mu_\omega < 1$ for all realizations of U_1, by Lemma 3 the voter always removes the incumbent following $\chi = F$. □

Before presenting the next lemma, we define $\Delta^-_\chi = \frac{\lambda_\chi}{2\alpha x} - \alpha x$ and $\Delta^+_\chi = \frac{\lambda_\chi}{2\alpha x} + \alpha x$. Note,

1. $\frac{\partial \Delta^+_\chi}{\partial x} = \frac{-\lambda_\chi}{2\alpha x^2} + \alpha = -\frac{1}{x}\Delta^-_\chi$, and
2. $\frac{\partial \Delta^-_\chi}{\partial x} = \frac{-\lambda_\chi}{2\alpha x^2} - \alpha = -\frac{1}{x}\Delta^+_\chi$

Furthermore, define $\Gamma_\chi(x) = \gamma\Delta^+_\chi\phi(\Delta^-) + (1-\gamma)\Delta^-_\chi\phi(\Delta^+_\chi)$ and define $\Omega_\chi(x) = \gamma\phi(\Delta^-_\chi) + (1-\gamma)\phi(\Delta^+_\chi)$.

Lemma (A2). *For $\chi \in \{N, S\}$, we have $\frac{\partial P_\chi}{\partial x} = \frac{1}{x}\Gamma_\chi(x)$ and $\frac{\partial^2 P_\chi}{\partial x^2} = -\frac{1}{x^2}\Gamma_\chi(x) + \frac{1}{x}\frac{\partial \Gamma_\chi}{\partial x}$, where $\frac{\partial \Gamma_\chi}{\partial x} = \frac{1}{x}\Big[\Delta^-_\chi\Delta^+_\chi\Gamma_\chi(x) - \Omega_\chi(x)\Big]$.*

Proof. Taking the derivative of $P_\chi(x)$ with respect to x we get

$$\begin{aligned} \frac{\partial P_\chi}{\partial x} &= \gamma\phi(\frac{\lambda_\chi}{2\alpha x} - \alpha x)[\frac{\lambda_\chi}{2\alpha x^2} + \alpha] + (1-\gamma)\phi(\frac{\lambda_\chi}{2\alpha x} + \alpha x)[\frac{\lambda_\chi}{2\alpha x^2} - \alpha] \\ &= \frac{1}{x}\Gamma_\chi(x). \end{aligned} \tag{8}$$

Next, differentiating Γ_χ, we obtain

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

The second derivative of $P_\chi(x)$ then follows. □

For a given π_I, let $\overline{x}_{\pi I}$ solve $4\alpha^2 x^2(1 + \alpha^2 x^2) = \lambda_N^2$.

Next, define x' as the positive solution to

$$2\overline{\gamma}_N + \frac{\lambda_N}{2\alpha^2(x')^2} = 1. \tag{9}$$

Let $\pi^\dagger$ be the value of π_I for which equation 9 is satisfied at $x = x_I$. Lemmas 5 and 7 below show that if $\pi_I > \pi^\dagger$, the policy maximizing the incumbent's probability of winning under no crisis is to the left of x_I. Otherwise, it is to the right of x_I. We use this to prove the following results.

Lemma (A3). *Assume there is no crisis and $\pi_I < \pi_I^\dagger$. If $x \in [0, \overline{x}_{\pi_I}]$ then $P_N(x)$ can only change concavity once. If $x > \overline{x}_{\pi_I}$ then $P_N(x)$ is concave in x.*

Proof. We have $\frac{\partial P_N}{\partial x} = \frac{1}{x}\Gamma(x)$. Differentiating Γ yields

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

P_N is concave in x if and only if

$$\begin{aligned}
&\frac{\partial^2 P_N}{\partial x^2} < 0\\
&\Leftrightarrow -\frac{1}{x^2}\Gamma(x) + \frac{1}{x}\frac{\partial \Gamma}{\partial x} < 0\\
&\Leftrightarrow \frac{1}{x}\frac{\partial \Gamma}{\partial x} < \frac{1}{x^2}\Gamma(x)\\
&\Leftrightarrow \frac{1}{x^2}\Big[\Delta^-\Delta^+\Gamma(x) - \Omega(x)\Big] < \frac{1}{x^2}\Gamma(x)\\
&\Leftrightarrow 0 < (1 - \Delta^-\Delta^+)\Gamma(x) + \Omega(x)
\end{aligned}$$

We have $\Omega(x) \geq 0$. Moreover, by $\pi_I < \pi^\dagger$, we have $\Gamma(x) > 0$ for any x which are possible solutions to the first-order condition. Thus, a sufficient condition for this inequality to hold is that $1 - \Delta^-\Delta^+ \geq 0$. By definition, this holds for all $x \geq \overline{x}_{\pi_I}$, completing the second part of the lemma.

Next, assume $x < \overline{x}_\pi$, so $1 - \Delta^-\Delta^+ < 0$. For P_N to be convex requires

$$0 > (1 - \Delta^-\Delta^+)\Gamma(x) + \Omega(x) \tag{10}$$

Assume that (10) holds. We show this implies that the RHS of (10) is increasing in x. Differentiating yields

$$\frac{\partial RHS(10)}{\partial x} = \frac{1}{x}[(\Delta^+)^2 + (\Delta^-)^2]\Gamma(x) + \frac{1}{x}[1 - \Delta^-\Delta^+]\Big[\Delta^-\Delta^+\Gamma(x) - \Omega(x)\Big]$$

We must have $\Delta^-\Delta^+\Gamma(x) - \Omega(x) < 0$, otherwise, (10) does not hold. Thus, if P_N is convex in x then $\frac{\partial RHS(10)}{\partial x} > 0$, which implies that once (10) no longer holds at some y it cannot again hold for $x > y$. □

Lemma (A4). *Assume there is no crisis and $\pi_I > \pi_I^\dagger$. If $x \in [0, \overline{x}_{\pi_I}]$ then $P_N(x)$ is concave. If $x > \overline{x}_{\pi_I}$ then $P_N(x)$ can only change concavity once.*

Proof. From the previous lemma we have that P_N is concave if and only if

$$0 < (1 - \Delta^-\Delta^+)\Gamma(x) + \Omega(x)$$

Since $\pi_I > \pi_I^\dagger$ we must have that $\Gamma(x) < 0$ for all x that are candidate solutions to the incumbent's problem. Additionally, $\Omega(x) > 0$. Therefore, a sufficient condition for this inequality to hold is $1 - \Delta^-\Delta^+ \leq 0$, which holds if and only if $x \leq \overline{x}_{\pi_I}$.

Next, assume $x > \overline{x}_{\pi_I}$, so $1 - \Delta^-\Delta^+ > 0$. Again, for P_N to be convex requires

$$0 > (1 - \Delta^-\Delta^+)\Gamma(x) + \Omega(x) \tag{11}$$

Assume that (11) holds. We show this implies that the RHS of (11) is decreasing in x. Differentiating yields

$$\frac{\partial RHS(11)}{\partial x} = \frac{1}{x}[(\Delta^+)^2 + (\Delta^-)^2]\Gamma(x) + \frac{1}{x}[1 - \Delta^-\Delta^+]\Big[\Delta^-\Delta^+\Gamma(x) - \Omega(x)\Big]$$

Since $\pi_I > \pi^\dagger$ we have $\Gamma(x) < 0$, thus, the first term in $\frac{\partial RHS(11)}{\partial x}$ is negative. Now consider the second term of $\frac{\partial RHS(11)}{\partial x}$. For (11) to hold we must have $\Delta^-\Delta^+\Gamma(x) - \Omega(x) > \Gamma(x)$. As such,

$$\Delta^-\Delta^+\Gamma(x) - \Omega(x) < \Delta^-\Delta^+(\Delta^-\Delta^+\Gamma(x) - \Omega(x)) - \Omega(x)$$

The RHS of this inequality rearranges to $(\Delta^-\Delta^+)^2\Gamma(x) - \Omega(x)(1 - \Delta^-\Delta^+) < 0$. Therefore, the second term of $\frac{\partial RHS(11)}{\partial x}$ is also negative. Consequently, if $P_N(x)$ is convex at $\hat{x}$, for $x' > \overline{x}_{\pi_I}$, then it is convex for all $x > \hat{x}$. □

Corollary (A1). *Suppose that $x_I > 0$. Then, $x^* > 0$. Otherwise, $x^* < 0$.*

Proof. By Lemma A1 we have that, for any $x_1 < 0$, a deviation to $-x_1 > 0$ leaves the probability of winning unchanged. This deviation is thus always profitable for a right-wing incumbent. Symmetric argument holds for a left-wing one. The strict inequality follows from Proposition 1. □

Lemma 4. *Let x_c^* the equilibrium policy under a crisis and x_n^* denote the equilibrium policy in normal times. In times of crisis the equilibrium policy x_c^* solves*

$$2(x_I - x) + (\beta + 4x_I^2)\Big(\pi_I \frac{\partial P_S}{\partial x} + (1 - \pi_I)\frac{\partial P_F}{\partial x}\Big) = 0.$$

In normal times the equilibrium policy x_n^* solves

$$2(x_I - x) + (\beta + 4x_I^2)\frac{\partial P_N}{\partial x} = 0.$$

Proof. Assume there is no crisis. By Lemma 5 we must have $x_n^* \in [0, x_I]$ for $\pi_I > \pi^\dagger$ and $x_n^* \in [x_I, \infty)$. We show that the equilibrium policy must be interior in each case and, thus, solves the first order condition. Note that the objective function is continuously differentiable for $x > 0$. Suppose $\pi_I > \pi^\dagger$. We have that as $x \to 0$, the first order condition goes to $x_I > 0$, since the normal PDF goes to 0 faster than any polynomial goes to ∞, sending $\frac{\partial P}{\partial x} \to 0$. Thus, the objective function is increasing as x increases away from x. For $\pi_I < \pi^\dagger$, clearly we must have $x_n^* < \infty$. Finally, if $x = x_I$ then the first order condition becomes

$$(\beta + 4x_I^2)\frac{\partial P}{\partial x}|_{x=x_I}.$$

By Lemma 5 we have $\frac{\partial P}{\partial x}|_{x=x_I}$ for all $\pi_I \neq \pi^\dagger$, and so $x_n^* \neq x_I$.

In a crisis we must have $x_c^* \leq$, since $\frac{\partial P_S}{\partial x} < 0$. The previous argument for $\pi \geq \pi^\dagger$ implies $x_c^* > 0$. Finally, at $x = x_I$ the first order condition becomes $\pi_I \frac{\partial P_S}{\partial x} < 0$ for all $\pi_I > 0$.

Next, we show that the equilibrium policy is unique for $\pi_I \leq \pi^\dagger$. Assume there is not a crisis. We have that U is concave in x if and only if

$$\frac{\partial^2 U}{\partial x^2} = (\beta + 2x_I)\frac{\partial^2 P_N}{\partial x^2} < 0$$

We show that I's optimal policy must lie in an interval over which his expected utility is concave.

If $\pi_I < \pi^\dagger$, then U is concave in x if and only if P_N is concave. By Lemma A3 either U is concave for all x, in which case we are done, or there is some $\tilde{x}$ such that U is convex over $[0, \tilde{x}]$

and concave for $x \geq \tilde{x}$. However, by $\pi_I < \pi^\dagger$ we must have $x^* \geq x_I > 0$. Thus, the optimal policy must lie in $[\tilde{x}, \infty)$, and U is concave over this interval. $\square$

Lemma 5. *Suppose $x > 0$.*

1. *Under a crisis, I's ex-ante probability of winning is decreasing in x.*
2. *In normal times, the following holds:*
   - *If $\pi_I > \pi_C$ then I's probability of winning is decreasing in x;*
   - *If $\pi_I < \tilde{\pi}$ then I's probability of winning is increasing in x;*
   - *If $\pi_I \in (\tilde{\pi}, \pi_C)$, then I's probability of winning is decreasing then increasing in x.*

Symmetric results hold for $x < 0$.

Proof. First, consider $\lambda_\chi < 0$. This implies either $\chi = S$ or $\chi = N$ and $\pi_I > \tilde{\pi}$. In this case, (8) is always negative if $\frac{\lambda_\chi}{2\alpha x^2} + \alpha < 0$, which holds for all $x \in [0, \sqrt{\frac{-\lambda_\chi}{2\alpha^2}}]$.

Now, assume $x > \sqrt{\frac{-\lambda_\chi}{2\alpha^2}}$. From (8) we have that $\frac{\partial P_\chi}{\partial x}$ is negative iff

$$\frac{\phi(\frac{\lambda_\chi}{2\alpha x} - \alpha x)}{\phi(\frac{\lambda_\chi}{2\alpha x} + \alpha x)} < \frac{1-\gamma}{\gamma}\left(\frac{-\frac{\lambda_\chi}{2\alpha x^2} + \alpha}{\frac{\lambda_\chi}{2\alpha x^2} + \alpha}\right). \tag{12}$$

Which we rewrite as

$$e^{-\frac{1}{2}(\frac{\lambda_\chi}{2\alpha x} - \alpha x)^2 + \frac{1}{2}(\frac{\lambda_\chi}{2\alpha x} + \alpha x)^2} < \frac{1-\gamma}{\gamma}\left(\frac{-\frac{\lambda_\chi}{2\alpha x^2} + \alpha}{\frac{\lambda_\chi}{2\alpha x^2} + \alpha}\right). \tag{13}$$

Logarithmic transformation on both sides. The above reduces to

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

which holds if and only if

$$\frac{\overline{\gamma}_\chi(1-\gamma)}{\gamma(1-\overline{\gamma}_\chi)} < \frac{1-\gamma}{\gamma}\frac{\frac{-\lambda_\chi}{2\alpha x^2} + \alpha}{\frac{\lambda_\chi}{2\alpha x^2} + \alpha}. \tag{15}$$

This obligingly reduces to

$$2\overline{\gamma}_\chi + \frac{\lambda_\chi}{2\alpha^2 x^2} < 1. \tag{16}$$

The above is satisfied at any x if $\overline{\gamma}_\chi < \frac{1}{2}$, which always holds when $\mu_\theta(\chi) > \pi_C$. If $\chi = S$ then $\mu_\theta(\chi) = 1 > \pi_C$ and (16) always holds. If $\chi = N$, then $\mu_\theta(\chi) = \pi_I$, thus (16) always holds when $\pi_I > \pi_C$. Recall that $\pi_C < \tilde{\pi}$ iff $\gamma < \frac{1}{2}$. Thus, when $\gamma < \frac{1}{2}$, $\pi_I > \tilde{\pi}$ implies $\pi_I > \pi_C$, and (16) is always satisfied. Finally, suppose $\chi = N$, $\gamma > \frac{1}{2}$ and $\pi_I \in (\tilde{\pi}, \pi_C)$. Then, there exists a unique $x' > \sqrt{\frac{-\lambda_N}{2\alpha^2}}$ s.t. (16) holds for $x < x'$, and fails otherwise. x' satisfies (9).

Next, consider $\lambda_\chi > 0$. This implies either $\chi = F$ or $\chi = N$ and $\pi_I < \tilde{\pi}$. (16) always fails when $\gamma_\chi > \frac{1}{2}$, which always holds when $\mu_\theta(\chi) < \pi_C$. If $\chi = F$, then $\mu_\theta = 0 < \pi_C$, and (16) never holds. If $\chi = N$, then $\mu_\theta(\chi) = \pi_I$ thus (16) fails when $\pi_I < \pi_C$. Recall that $\pi_C < \tilde{\pi}$ iff $\gamma < \frac{1}{2}$. Thus, when $\gamma > \frac{1}{2}$, $\pi_I < \tilde{\pi}$ implies $\pi_I < \pi_C$, and (16) is never satisfied. Finally, suppose $\chi = N$, $\gamma < \frac{1}{2}$ and $\pi_I \in (\pi_C, \tilde{\pi})$. Then, it is easy to see that (16) holds for $x > x'$, and fails otherwise, where x' satisfies (9). □

Lemma 6. *If there is a crisis, then the incumbent never gambles in equilibrium,* $x_c^* \leq x_I$.

Proof. Follows straightforwardly from Corollary A1, Lemma 5 point 1 and the FOC. □

Lemma (A5). *Define* $\widehat{\beta} = \frac{x^2(4\gamma-3)}{1-\gamma}$. *Define* $\tilde{\pi}_I = \pi_C + \frac{\alpha x_I}{p_b K}\Big(1 - 2\gamma\Big)$. *Assume* $\tilde{\pi}_I \in (0,1)$, *which holds if and only if*

$$\gamma \in \Big(\frac{1}{2} - \frac{1-\pi_C}{2}\frac{p_b K}{\alpha x_I}, \frac{1}{2} + \frac{\pi_C}{2}\frac{p_b K}{\alpha x_I}\Big).$$

Then, we have

1. $\lim_{\pi_I \to \tilde{\pi}^+} x_n^* = 0$.
2. *The optimal policy choice* x_n^* *is discontinuous in* π_I *at* $\pi_I = \tilde{\pi}$.
3. *The optimal policy choice* x_n^* *is differentiable in* π_I *for* $\pi_I \leq \tilde{\pi}$.
4. *The optimal policy choice* x_n^* *is differentiable in* π_I *for* $\pi_I > \tilde{\pi}$

Proof. First, we show that the $\tilde{\pi}$ politician's expected utility from any x is lower than her expected utility if she could choose $x = 0$ and win with probability 1. First, note that the $\tilde{\pi}$ politician's probability of winning is bound above by γ, since his probability of winning is either strictly increasing in x ($\gamma > 1/2$) or strictly decreasing ($\gamma < 1/2$). Thus, for $\tilde{\pi}$ we have $U_{\tilde{\pi}_I}(x) \leq \gamma\beta - 4x_I^2(1-\gamma)$. If the $\tilde{\pi}$ incumbent chooses $x = 0$ and wins with probability 1 her expected utility is $\beta - x_I^2$. Define $\tilde{U}_0 = \beta - x_I^2$. Thus, if $\beta - x_I^2 > \gamma\beta - 4x_I^2(1-\gamma)$ then $\tilde{U}_0 > U(x)$ for any policy x. This holds if and only if β is sufficiently large, which holds by Assumption that $\beta > \widehat{\beta}$. For the $\tilde{\pi}$ incumbent and any $x > 0$ define $\Delta_x = \tilde{U}_0 - U(x)$. Next, we show that for any $x^\dagger > 0$ there exists δ_x such that if $|\pi_I - \tilde{\pi}| < \delta_x$ then $U_{\pi_I}(x = 0) - U_{\pi_I}(x^\dagger) > 0$. We have

$$\begin{aligned}
&U_{\pi_I}(x = 0) - U_{\pi_I}(x^\dagger) \\
&= U_{\pi_I}(x = 0) - U_{\pi_I}(x^\dagger) + U_{\tilde{\pi}_I}(x^\dagger) - U_{\tilde{\pi}_I}(x^\dagger) \\
&= U_{\pi_I}(x = 0) - U_{\tilde{\pi}_I}(x^\dagger) - |U_{\pi_I}(x^\dagger) - U_{\tilde{\pi}_I}(x^\dagger)| \\
&= \Delta_x - |U_{\pi_I}(x^\dagger) - U_{\tilde{\pi}_I}(x^\dagger)|,
\end{aligned}$$

where the final line follow from $U_{\pi_I}(x = 0) = \tilde{U}_0$ and $U_{\pi_I}(x^\dagger) > U_{\tilde{\pi}_I}(x^\dagger)$. Because $U(x)$ is continuous in (x, π_I) for $x > 0$, we have that

$$\lim_{\pi_I \to \tilde{\pi}} \Delta'_x - |U_{\pi_I}(x^\dagger) - U_{\tilde{\pi}_I}(x^\dagger)| = \Delta'_x > 0.$$

Since this holds for any $x^\dagger > 0$, we must have that $\lim_{\pi_I \to \tilde{\pi}} x^*_n = 0$.

Next, notice that for $\pi_I = \tilde{\pi}$, the equilibrium policy is bounded away from 0 (which we can show by verifying that the incumbent's utility is increasing in x at $x = 0$ since the normal pdf goes to 0 faster than any polynomial goes to ∞). This proves point 2.

The proof of point 3 follows straightforwardly from two observations. First, as $\pi \to \tilde{\pi}$ from the left x^*_n is bounded away from zero (see the proof of point 2. above). Second, the objective function is differentiable in x and π_I when $x > 0$.

Point 4 follows from the observation that objective function is differentiable in x and π_I when $x > 0$, but it is not differentiable at $\tilde{\pi}$ when $x = 0$. The strict inequality then follows from point 1. □

Lemma 7. *Suppose there is no crisis. If $\pi_I < \tilde{\pi}$ then the incumbent always gambles in equilibrium. Otherwise, if $\pi_I > \tilde{\pi}$ then the incumbent never gambles.*

Proof. We break the argument into two cases.

First, assume $\gamma < \frac{1}{2}$, which implies $\pi_C < \tilde{\pi}$.

By Lemma 5, the probability of winning is decreasing in x for $\pi_I > \tilde{\pi}$. Consequently, for any $x > x_I$ the incumbent can choose an $x < x_I$ that yields the same policy utility but a higher probability of winning. Thus, if $\pi_I > \tilde{\pi}$ we must have the incumbent's optimal policy $x^*_n \leq x_I$. Similarly, for $\pi_I < \pi_C$ the incumbent's probability of winning is increasing in x, and the incumbent's optimal policy is $x^*_n \geq x_I$.

Next, consider $\pi_I \in (\pi_C, \tilde{\pi})$. By Lemma 5, the incumbent's probability of winning is maximized at $x = x'$. If $x' < x_I$ then the incumbent's probability of winning is decreasing in x for $x > x_I$. Thus, we must have $x_n^* \leq x_I$. On the other hand, if $x' > x_I$ then the incumbent's probability of winning is increasing in x for $x < x_I$ and so it must be that $x_n^* \geq x_I$. Recall that $x' = \sqrt{\frac{\lambda_N}{2\alpha^2(1-2\overline{\gamma}_N)}}$. We have $\lim_{\pi_I \to \pi_C} x' = \infty$, since $\overline{\gamma}_N = 1/2$ when evaluated at $\pi_I = \pi_C$, and $\lim_{\pi_I \to \tilde{\pi}} x' = 0$, since $\lambda_N = 0$ at $\pi_I = \tilde{\pi}$. Furthermore, $\frac{\partial x'}{\partial \pi_I} < 0$. Thus, there exists a unique π_I that solves $x' = x_I$ and defining $\pi^\dagger$ as this solution yields the result. Second, assume $\gamma > 1/2$, which implies $\pi_C > \tilde{\pi}$. By Lemma 5, the probability of winning is decreasing in x for $\pi_I > \pi_C$. Thus, if $\pi_I > \tilde{\pi}$ then $x_n^* \leq x_I$. Similarly, if $\pi_I < \tilde{\pi}$ then the incumbent's probability of winning is increasing in x and $x_n^* \geq x_I$.

Finally, consider $\pi_I \in (\tilde{\pi}, \pi_C)$. For $x > x'$ the incumbent's probability of winning is increasing in x and bound above by γ. Thus, the incumbent's expected utility for $x \geq x_I$ is bound above by the expected utility to choosing $x = x_I$ and winning with probability γ, i.e., $U(x \geq x_I) < -4x_I^2(1-\gamma) + \gamma\beta$. Since $\pi_I > \tilde{\pi}$, if the incumbent chooses $x = 0$ then the incumbent wins with probability 1. Thus, the incumbent's expected utility for choosing $x < x_I$ is bound below by the expected utility for choosing $x = 0$ and winning, i.e., $U(x < x_I) \geq -x_I^2 + \beta$. Consequently, a sufficient condition to ensure $x_n^* < x_I$ is that

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

which holds by assumption that $\beta > \widehat{\beta}$. Therefore, defining $\pi^\dagger = \tilde{\pi}$ for $\gamma > 1/2$ yields the result. □

Proposition 1*. *If $\pi_I < \pi^\dagger$, then the crisis has a qualitative effect on policymaking: I gambles in normal times, but does not gamble under a crisis, $x_c^* < x_I < x_n^*$. If $\pi_I > \pi^\dagger$, then the crisis has no qualitative effect, $x_c^*, x_n^* < x_I$.*

Proof. Follows from Lemmas 6 and 7. □

Proposition 2. *Suppose $\pi_I > \pi^\dagger$.*

- *When $\gamma > \frac{1}{2}$, the crisis always dampens incentives to control information;*
- *When $\gamma < \frac{1}{2}$, there exists a $\widehat{\pi} \leq \tilde{\pi}$ such that the crisis dampens incentives to control information if $\pi_I > \tilde{\pi}$ and amplifies these incentives if $\pi_I < \widehat{\pi}$.*

Proof. We know from Lemma A5 that $x_n^* < x_c^*$ right above $\tilde{\pi}$. Further, from the proof of Proposition 3, we know that x_c^* is decreasing in π_I and x_n^* is decreasing in π_I (for $\pi_I > \tilde{\pi}$). Thus, it must be the case that $x_n^* < x_c^*$ for any $\pi_I > \tilde{\pi}$. When $\gamma > \frac{1}{2}$, we have that $\pi^\dagger = \tilde{\pi}$ (from Lemma 7). This proves point 1. and the first part of point 2.

Next, suppose $\gamma < \frac{1}{2}$ and $\pi \in (\pi^\dagger, \tilde{\pi}]$. First, notice that Lemmas 6 and 7 imply that $x_n^* > x_I > x_c^*$ for π_I right above $\tilde{\pi}$. Equilibrium policy is continuous in this range by Lemma A5. Thus, there must exist a $\widehat{\pi} \in (\pi^\dagger, \tilde{\pi}]$ s.t. $x_c^* < x_n^*$ for any $\pi < \widehat{\pi}$. Notice that, if x_c^* and x_n^* only cross once (at $\pi_I = \tilde{\pi}$), we have $\widehat{\pi} = \tilde{\pi}$. □

Proposition 3.

1. *The equilibrium crisis policy, x_c^*, is decreasing in π_I.*

2. *There exists a cut-point $\underline{\pi}_I < \tilde{\pi}$, such that: if $\pi_I \in (\underline{\pi}_I, \tilde{\pi})$ then the equilibrium policy under no crisis, x_n^*, is decreasing in π_I. Otherwise, x_n^* is increasing in π_I.*

Proof. Assume there is no crisis in the first period. First, we show that x_n^* must be increasing in π_I for $\pi_I > \tilde{\pi}$. Applying the implicit function theorem we have

$$\frac{\partial x_n^*}{\partial \pi_I} = -(\beta + 4x_I^2)\frac{\frac{\partial^2 P_N}{\partial \pi_I \partial x_1}}{\frac{-2+\partial^2 P_N}{\partial x^2}}.$$

Thus, $\frac{\partial x_n^*}{\partial \pi_I} < 0$ if and only if

$$\frac{\partial^2 P_N}{\partial \pi_I \partial x_1} < 0.$$
$$\Leftrightarrow \frac{1}{2\alpha x^2}\frac{\partial \lambda_N}{\partial \pi_I}\Big(\gamma\phi(\Delta^-) + (1-\gamma)\phi(\Delta^+) + \gamma\phi'(\Delta^-)\Delta^+ + (1-\gamma)\phi'(\Delta^+)\Delta^-\Big) < 0,$$

Since $\frac{\partial \lambda_N}{\partial \pi_I} < 0$, the inequality holds if and only if

$$\gamma\phi(\Delta^-) + (1-\gamma)\phi(\Delta^+) + \gamma\phi'(\Delta^-)\Delta^+ + (1-\gamma)\phi'(\Delta^+)\Delta^- > 0$$
$$\Leftrightarrow \gamma\phi(\Delta^-) + (1-\gamma)\phi(\Delta^+) - \gamma\phi(\Delta^-)\Delta^-\Delta^+ - (1-\gamma)\phi(\Delta^+)\Delta^+\Delta^- > 0$$
$$\Leftrightarrow (1 - \Delta^-\Delta^+)(\gamma\phi(\Delta^-) + (1-\gamma)\phi(\Delta^+)) > 0.$$

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

$$4\alpha^2 x^2(1 + \alpha^2 x^2) > \lambda_N^2. \tag{17}$$

Lemma A5 point 1. implies that to the right of $\tilde{\pi}$, x_n^* must be increasing in π_I. Thus, (17) cannot hold. Consider now increasing π_I away from $\tilde{\pi}$. This increases both sides of the inequality.

If the LHS remains smaller than the RHS for all $\pi_I \in (\tilde{\pi}, 1]$, then this implies that x_n^* is increasing in π_I over this range. Suppose instead that there exists a point $\pi' > \tilde{\pi}$ s.t. (17) fails for $\pi < \pi'$, and holds with equality at $\pi = \pi'$. Then, increasing π_I away from π' increases the RHS decreases the LHS. But then, this implies that (17) continues to fail for any $\pi \geq \pi'$. Thus, it must be the case that (17) fails for all $\pi_I \in (\tilde{\pi}, 1]$, and x_n^* is increasing in π_I over this range.

Next, consider $\pi_I \leq \tilde{\pi}$. By construction, at $\pi_I = \tilde{\pi}_I$ we have $\lambda_N = 0$. As $x_n^* > 0$, inequality (17) must hold. Since x_n^* is continuous in π_I in this range, it must also hold for all π sufficiently close to $\tilde{\pi}_I$, i.e., for $\pi_I \in (\underline{\pi}_I, \tilde{\pi}]$. We now show that (17) can only hold over this interval. Note that λ_N^2 is increasing as π_I moves away from $\tilde{\pi}_I$. First, consider decreasing π_I down from $\tilde{\pi}$. Doing so increases λ_N^2 and increases x_n^*, which increases $4\alpha^2 x^{*2}_n(1 + \alpha^2 x^{*2}_n)$. If $4\alpha^2 x^{*2}_n(1 + \alpha^2 x^{*2}_n)$ remains greater than λ_N^2 for all $\pi_I < \tilde{\pi}_I$ then $\frac{\partial x_n^*}{\partial \pi_I} < 0$ for all $\pi_I \in (0, \tilde{\pi}_I)$. Next, assume that at some point λ_N^2 crosses $4\alpha^2 x_n^{*2}(1 + \alpha^2 x_n^{*2})$, call it $\underline{\pi}$. Then $\frac{\partial x_n^*}{\partial \pi_I} \geq 0$ and decreasing π_I further from $\underline{\pi}_I$ decreases x_n^* which decrease $4\alpha^2 x_n^{*2}(1 + \alpha^2 x_n^{*2})$ and increases λ_N^2. Thus, $4\alpha^2 x_n^{*2}(1 + \alpha^2 x_n^{*2}) > \lambda_N^2$ for all $\pi_I < \underline{\pi}_I$. Moreover, $\underline{\pi}_I$ is unique by uniqueness of the equilibrium policy choice for $\pi_I < \pi^\dagger$.

Assume there is a crisis. Applying the implicit function theorem yields

$$\frac{\partial x_n^*}{\partial \pi_I} = -(\beta + 4x_I^2)\frac{\frac{\partial P_S}{\partial x_1} - \frac{\partial P_F}{\partial x_1}}{-2 + \pi_I \frac{\partial^2 P_S}{\partial x^2} + (1 - \pi_I)\frac{\partial^2 P_F}{\partial x^2}}.$$

Thus, $\frac{\partial x_n^*}{\partial \pi_I} \leq 0$ if and only if $\frac{\partial P_S}{\partial x_1} - \frac{\partial P_F}{\partial x_1} \leq 0$. By Lemma 5 we have $\frac{\partial P_S}{\partial x_1} \leq 0$ and $\frac{\partial P_F}{\partial x_1} = 0$. Therefore, $\frac{\partial x^*}{\partial \pi_I} \leq 0$, as required. □

Proposition 4* *There exists a $\pi^d \leq \tilde{\pi}$ s.t.*

- *if $\pi_I > \tilde{\pi}$, then $x_c^* > x_n^*$;*
- *if $\pi_I < \pi^d$, then $x_c^* < x_n^*$.*

Proof. The first point follows from the proof of Proposition 2. Consider instead $\pi_I < \tilde{\pi}$. When $\gamma > \frac{1}{2}$, we have that $\tilde{\pi} = \pi^\dagger$. This implies $x_c^* < x_I < x_n^*$ for any $\pi_I < \tilde{\pi}$. Finally, suppose $\gamma < \frac{1}{2}$. From Lemmas 7 and Proposition 1, $x_c^* < x_I < x_n^*$ for all $\pi_I < \pi^\dagger < \tilde{\pi}$. Finally, from the proof of Proposition 2, we have that $x_c^* < x_n^*$ for all $\pi_I < \widehat{\pi} \leq \tilde{\pi}$ (where $\widehat{\pi} = \tilde{\pi}$ if x_c^* and x_n^* cross only once, and$\widehat{\pi} < \tilde{\pi}$ otherwise). This concludes the proof of the second point. We note that if x_c^* and x_n^* can cross at most once, we have $\pi^d = \tilde{\pi}$, and this is always the case when $\gamma > 1/2$. □

Corollary 3. $x_n^{max} - x_n^{min} > x_c^{max} - x_c^{min}$.

Proof. Follows from Proposition 3. □

Proposition 5*. *Suppose office benefit is sufficiently large. If $\pi_I \in [\gamma, \tilde{\pi}]$, then there exist parameter values such that endogenous crises emerge in equilibrium. Otherwise, if $\pi_I \notin [\gamma, \tilde{\pi}]$, then endogenous crises never emerge.*

Proof. To start, recall that in normal times the incumbent's equilibrium policy solves

$$\frac{-2(x - x_I)}{\beta + 4x_I^2} + \frac{\partial P_N}{\partial x} = 0.$$

As $\beta \to \infty$ the LHS of the FOC goes to $\frac{\partial P_N}{\partial x}$. Since the incumbent's problem is continuous in β we must have that x_n^* approaches the policy that simply maximizes I's winning probability. Denote this policy as $x_p^* = \text{argmax}\, P_N(x)$. Moreover, we must have that $\lim_{\beta\to\infty} U_n(x_n^*) = (\beta + 2x_I)P_n(x_p^*)$. To see this, consider the ratio

$$\begin{aligned}\frac{U_N(x_n^*)}{(\beta + 2x_I)P_N(x_p^*)} &= \frac{-|x_n^* - x_I| + (\beta + 2x_I)P_N(x^*n)}{(\beta + 2x_I)P_n(x_p^*)} \\ &= \frac{P_N(x_n^*)}{P_N(x_p^*)} - \frac{|x_n^* - x_I|}{(\beta + 2x_I)P_N(x_P^*)}.\end{aligned}$$

Thus, $\lim_{\beta\to\infty} \frac{U_n(x_n^*)}{(\beta+2x_I)P_N(x_p^*)} = 1$. A similar argument yields that $\lim_{\beta\to\infty} \frac{U_c(x_n^*)}{(\beta+2x_I)\pi_I P_S(0)} = 1$ (note that $x = 0$ maximizes $(\beta + 4x_I^2)\pi_I P_S(0)$). Consequently, for β sufficiently high, $U_n(x_n^*) < U_c(x_c^*)$ if and only if $(\beta + 4x_I^2)P_N(x_p^*) < (\beta + 4x_I^2)\pi_I P_S(0)$, which reduces to $P_N(x_p^*) < \pi_I$.

We now show that if $\pi_I > \tilde{\pi}$ then the incumbent never wants to generate a crisis for β sufficiently high. In this case, $P_N(x_p^*) = P_N(0) = 1 > \pi_I$, and so the incumbent's equilibrium payoff is higher with no crisis.

Next, suppose $\pi_I < \gamma$. In this case, $\max P_N(x) \geq \gamma > \pi_I$, and so the incumbent's equilibrium policy is higher without a crisis. Moreover, this implies that if $\gamma > \tilde{\pi}_I$ then the incumbent is never better under a crisis.

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

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

Lemma 8. *Suppose there is a crisis. In any perfect Bayesian equilibrium, we always have that the incumbent chooses different policies following each signal, $x^a_c(0) \neq x^a_c(1)$. Furthermore, $x^a_c(0) = x^b_c(\pi_I = \psi_0)$ and $x^a_c(1) = x^b_c(\pi_I = \psi_1)$.*

Proof. Suppose a crisis emerges in the first period. Then, the incumbent's type will always be revealed via the crisis outcome, thereby making the interim posterior $\mu_\theta(x_1)$ electorally irrelevant.[10] Therefore, the incumbent's policy choice influences his reelection chances only via experimentation and the voter learning on the policy dimension. Thus, in equilibrium the incumbent will always act as if there was no asymmetry of information between him and the voter, and implement the dynamically optimal policy given the interim posterior ψ_ζ. □

Corollary 4. *The expected policy in times of crisis is decreasing in π_I.*

Proof. The ex-ante expected policy is given by

$$\begin{aligned} E[x^*_c] = &[\pi_I p(\zeta = 1|\theta = 1) + (1 - \pi_I)p(\zeta = 1|\theta = 0)]x^*_c(\psi_1) \\ &+[\pi_I p(\zeta = 0|\theta = 1) + (1 - \pi_I)p(\zeta = 0|\theta = 0)]x^*_c(\psi_0) \end{aligned}$$

This yields

$$\begin{aligned} \frac{\partial E[x^*_c]}{\partial \pi_I} = &[p(\zeta = 1|\theta = 1) - p(\zeta = 1|\theta = 0)]x^*_c(\psi_1) + [\pi_I p(\zeta = 1|\theta = 1) + (1 - \pi_I)p(\zeta = 1|\theta = 0)]\frac{\partial x^*_c(\psi_1)}{\partial \pi_I} \\ &+[p(\zeta = 0|\theta = 1) - p(\zeta = 0|\theta = 0)]x^*_c(\psi_0) + [\pi_I p(\zeta = 0|\theta = 1) + (1 - \pi_I)p(\zeta = 0|\theta = 0)]\frac{\partial x^*_c(\psi_0)}{\partial \pi_I}. \end{aligned}$$

We know from the analysis in the main body that $x^*_c(\psi_\theta)$ is decreasing in ψ_θ, and therefore decreasing in π_I. Notice that this also implies that $x^*_c(0) > x^*_c(1)$. Further, recall that $p(\zeta = 1|\theta = 1) = p(\zeta = 0|\theta = 0) > p(\zeta = 0|\theta = 1) = p(\zeta = 1|\theta = 0)$. Therefore, we have that

$$[p(\zeta = 1|\theta = 1) - p(\zeta = 1|\theta = 0)]x^*_c(\psi_1) + [p(\zeta = 0|\theta = 1) - p(\zeta = 0|\theta = 0)]x^*_c(\psi_0) < 0,$$

and

[10] Here we are assuming that, if the voter's interim posterior and the posterior conditional on crisis outcome are incompatible, the voter only updates conditioning on o^c_1. Specifically, if the voter forms interim posterior that the incumbent is competent but then observes a failed crisis, we assume that she reaches final beliefs that $\theta_I = 0$.

$$\frac{\partial E[x_c^*]}{\partial \pi_I} < 0.$$

□

Lemma 9. *There always exists an equilibrium where, regardless of his private signal, the incumbent adopts $x_n^b(\pi_I)$.*

Proof. First, notice that fixing the voter's interim posterior $\mu_\theta(x_1)$, both types have the same dynamically optimal policy. Thus, fixing the voter's belief at the prior, the incumbent's optimal policy is the same as in the baseline, regardless of the private signal. Suppose that, following a deviation from the conjectured equilibrium the voter forms interim beliefs $\mu_\theta = \psi_0$. Denote $x_1^d(\psi_0)$ the policy maximizing the incumbent's utility, conditional on the voter forming interim beliefs ψ_0. By definition of $x_n^*(\pi_I)$ (and the envelope theorem), it has to be the case that the incumbent's utility is higher at $x_1^*(\pi_I)$ than at $x_1^d(\psi_0)$. Thus, neither type has a profitable deviation from the conjectured □

Corollary (A2). *Suppose no crisis emerges in the first period. Among all possible pooling equilibria, the one where the incumbent adopts $x_n^*(\pi_I)$ is the one that maximizes the incumbent's utility.*

Proof. Follows from the definition of $x_n^*(\pi_I)$ and the observation that the private signal ζ does not influence the incumbent's expected utility when no crisis arises. □

Lemma 10. *In any separating equilibrium, both types are indifferent between policies on the equilibrium path.*

Proof. Recall that, for any policy x_1, fixing the voter's interim beliefs $\mu_\theta(x_1)$, both types have the same expected utility. Thus, if one type strictly prefers the conjectured equilibrium policy to imitating the other type, separation cannot be sustained. In any separating equilibrium, therefore, both types have to be indifferent between policies on the equilibrium path. □

Proposition 6. *Suppose no crisis emerges in the first period. Among all perfect Bayesian equilibria, the one where the incumbent adopts $x_n^a(0) = x_n^a(1) = x_n^b(\pi_I)$ maximizes the expected utility of both types of the incumbent.*

Proof. First, suppose that the equilibrium is separating. Then, the $\zeta = 0$ type must be at his dynamic optimum $x_n^*(\psi_0)$. Here, the envelope theorem implies that his expected utility must be lower than in the pooling equilibrium where the voter's interim is at the prior $\pi_I > \psi_0$ and $x_n^*(\pi_I)$. Further, we know from the previous Lemma that in any separating equilibria both types are indifferent between policies on the equilibrium path. Recall that, fixing the voter's interim, the two types have the same expected utility. This implies that the $\zeta = 1$ type is also better off in the equilibrium identified in Lemma 9.

Combining the above with Corollary A2, we have the stated result. □

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.

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

Ashworth, Scott, Ethan Bueno de Mesquita and Amanda Friedenberg. 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." *The Review of Economic Studies* 78(2):487–522.

Bils, Peter. 2020. "Overreacting and Posturing: How Accountability and Ideology Shape Executive Policy Responses.".

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

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

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.

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.

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. 1996. "The political economy of delayed reform." *The Journal of Policy Reform* 1(1):25–46.

Drazen, Allan. 2000. *Political Economy in Macroeconomics.* Princeton University Press.

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." *The American economic review* pp. 1146–1155.

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

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

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

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

Healy, Andrew and Neil Malhotra. 2013. "Retrospective voting reconsidered." *Annual Review of Political Science* 16:285–306.

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

Izzo, Federica. 2020. "Who runs when?".

Judd, Gleason. 2017. "Showing off: promise and peril in unilateral policymaking." *Quarterly Journal of Political Science* 12(2):241–68.

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* .

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.

Mahmalat, Mounir and Declan Curran. 2018. "Do Crises Induce Reform? A Critical Review of Conception, Methodology and Empirical Evidence of the 'Crisis Hypothesis'." *Journal of Economic Surveys* 32(3):613–648.

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.

Smith, Alastair. 1996. "Diversionary foreign policy in democratic systems." *International Studies Quarterly* 40(1):133–153.

Stewart, Alexander J, Nolan McCarty and Joanna J Bryson. 2020. "Polarization under rising inequality and economic decline." *Science advances* 6(50):eabd4201.

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

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

Tarar, Ahmer. 2006. "Diversionary incentives and the bargaining approach to war." *International Studies Quarterly* 50(1):169–188.

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.

Williamson, John. 1994. "In search of a manual for technopols." *The political economy of policy reform* pp. 11–28.