A MODEL OF CRISIS MANAGEMENT

By

Fei Li and Jidong Zhou

December 2020

COWLES FOUNDATION DISCUSSION PAPER NO. 2266

COWLES FOUNDATION FOR RESEARCH IN ECONOMICS
YALE UNIVERSITY
Box 208281
New Haven, Connecticut 06520-8281

http://cowles.yale.edu/

A Model of Crisis Management*

Fei Li† Jidong Zhou‡

December, 2020

Abstract

We propose a model of how multiple societies respond to a common crisis. A government faces a "damned-either-way" policy-making dilemma: aggressive intervention contains the crisis, but the resulting good outcome makes people skeptical of the costly response; light intervention worsens the crisis and causes the government to be faulted for not doing enough. This dilemma can be mitigated for the society that encounters the crisis first if another society faces the same crisis afterward. Our model predicts that the later society does not necessarily perform better despite having more information, while the earlier society might benefit from a dynamic counterfactual effect.

Keywords: Crisis Management, Counterfactual Effect, Political Accountability, Public Policy, Pandemic

JEL Classification Codes: D82, D83, F50, H12

*We are grateful to Attila Ambrus, Mark Armstrong, Heski Bar-Isaac, Gary Biglaiser, Florian Ederer, Jonathan Feinstein, Yunzhi Hu, Ilwoo Hwang, Jin Li, Navin Kartik, Leslie Marx, Peter Norman, Benjamin Polak, Larry Samuelson, Yangbo Song, various seminar audiences, and especially Barry Nalebuff for helpful comments.

†University of North Carolina at Chapel Hill. lifei@email.unc.edu

‡Yale University. jidong.zhou@yale.edu

1 Introduction

Many crises can be contained or even prevented if proper measures are taken in a timely manner. However, people are often uncertain about the severity of the threat, and they assess a policymaker's response only after seeing the consequences and updating their opinion accordingly. This can pose a challenge to the policymaker: if an aggressive action is taken and the crisis is prevented, people may then underestimate the severity of the problem and view the costly aggressive action as unnecessary; if a less aggressive action is taken and the crisis gets out of control, people may then blame the policymaker for not having taken the necessary precautions. This leads to a "damned if you do, damned if you don't" *policy-making dilemma*. Such a dilemma is relevant in a wide range of circumstances such as how to respond to infectious diseases or terrorism threats, whether to take precautionary economic measures to prevent a potential economic recession or financial crisis, or how to regulate manufacturing or cyber-security.[1] If policymakers are more concerned about being accused of overreacting, they may then choose a light intervention even if they know the threat is severe, resulting in an inefficient response.[2]

The dilemma would be resolved if the public could observe the counterfactual outcome under alternative policies. This is possible if other jurisdictions or societies face a similar threat but somehow respond differently. For instance, if one society acts aggressively but the other does not, then if the crisis gets out of control in the latter society, the policymaker in the former society can point to the bad outcome in the latter as justification for having taken the aggressive policy.

In this paper, we argue that when multiple societies face a common crisis, they may endogenously adopt distinct policies (even if they are otherwise identical), so that some societies can escape from the policy-making dilemma. To illustrate the idea, we consider the

[1]In his State of the Union Address in 1962, President John F. Kennedy said, "The time to repair the roof is when the sun is shining" to advocate his economic policy for preventing another recession. This was widely supported as the country had just experienced the 1960 recession. Such preventive economic policies, however, are often less popular when the economy is in a good state, as Christine Lagarde commented in her speech "A Time to Repair the Roof" at Harvard University in October 2017.

[2]There is anecdotal evidence that the concern of being blamed for overreacting is well justified. For example, in U.S. history of epidemics, neither Woodrow Wilson's 1918 influenza pandemic failure nor Dwight Eisenhower's misguided response to the 1957 influenza pandemic was faulted. However, in 1976 Gerald Ford was widely criticized and ridiculed thanks to his heavily publicized effort to prevent a new variant of influenza which turned out to be less deadly than expected. See Skidmore (2016) and Stasavage (2020). Also see Healy and Malhotra (2009) for evidence that voters do not reward the incumbent presidential party for disaster preparedness spending.

case when societies handle a crisis sequentially, though, as we point out later, a similar idea applies to the simultaneous-move case as well. When a society that encounters the crisis first takes the precaution and the crisis is contained, the public in the subsequent society, after seeing the outcome in the first society, may then become more optimistic. This boosted optimism makes it harder for the policymaker in the second society to follow suit. Thus the aggressive action in the first society can mislead the public in the second society and prevent its policymaker from adopting the right policy. The resulting adverse outcome, however, can then be used to justify the first policymaker's initially unpopular policy choice. We refer to this as a *dynamic counterfactual effect*.

A consequence of this counterfactual effect is that, all else equal, societies that encounter the crisis later may handle it worse than early-hit societies. Later movers can learn from earlier movers' experiences, but having more information is not necessary a blessing: people may become too optimistic after seeing good outcomes in early-hit societies. Conversely, the early-hit societies, foreseeing the possible counterfactual from subsequent countries, are more willing to adopt costly but more effective policies.

We develop a model of sequential crisis management to capture both the "damned-either-way" policy-making dilemma and the dynamic counterfactual effect. The policy-making dilemma relies on two modeling ingredients: each policy option is more likely to yield an outcome that induces people to believe some alternative option would have worked better;[3] and the policymaker is held accountable for their policy after its consequence is observed. Both are natural in the context of crisis control and prevention. Due to the policy-making dilemma, it is possible that the policymaker panders to public opinion on the severity of the crisis, causing a suboptimal policy choice. The dynamic counterfactual effect further relies on the assumption that people in each society evaluate their policymaker after observing policy consequences in all societies. This is plausible when the policy consequence in each society is realized relatively quickly and is publicly observable.

An obvious application of our model is the ongoing COVID-19 pandemic. Every country's government faces the policy-making dilemma due to the public's initial uncertainty about the severity of the virus. Since some countries were hit earlier than others, it makes the dynamics across countries important. Our model provides a new angle to understand the response disparity among societies, complementing a spectrum of possible explanations from cultural differences to institutional heterogeneity. In particular, our model predicts: (i)

[3]This paper focuses on policies which can influence people's judgment of the severity of the crisis. See Section 5 for a discussion of another type of policies which aim to reduce the damage of a crisis after its severity is already known.

early-hit countries are less hesitant to adopt aggressive and precautionary measures such as massive testing and tracking, mandatory quarantine and even lock-downs from the early stage; (ii) subsequent countries face a stronger political hurdle to take strict measures; (iii) the strict measures adopted by early-hit countries may be initially criticized, but later lauded after the adverse consequences under alternative responses are observed.

There are other relevant examples. For instance, when a terrorism threat spreads internationally, governments need to decide, often in a sequential manner depending on where the threat first emerges, how aggressively to tackle it when the public is initially uncertain about its severity; trading off between free expression and public safety, a social media platform or regulator often falls into a damned-either-way trap, but its hands can be untied by the possible loose regulation and adverse consequence from other platforms or societies.[4]

To deliver the main idea transparently, we adopt two perhaps unconventional assumptions. First, we assume that the public does not believe that the policymaker possesses any superior information on the true state (e.g., due to a mistrust) and so does not attempt to infer information on the state merely from their chosen policy. In other words, the public learns the state only from the policy outcome. This shuts down a potential signaling channel. However, we show in an extension that apart from an extreme case, adding the signaling channel does not affect our main insight. Second, we assume that the public evaluates the policy using their updated belief of the state after seeing the policy outcome. This is different from the usual approach where the public evaluates the policymaker based on some underlying characteristics such as their competence or preferences which can be learned from policy choices. We report such a "reputation" modelling approach in Appendix B and show that it delivers similar results but with somewhat different underlying economics.

Finally, notice that although we choose the sequential-move model to deliver our main message, the counterfactual effect is also present when societies make decisions simultaneously. In that case, asymmetric equilibria with different responses across societies or even mixed-strategy equilibria can arise. For instance, if the policymaker in one society expects the other to take an aggressive action, they then anticipate a more severe domestic policy-making dilemma and so will be more hesitant to adopt the same policy. The policymaker in the other society, anticipating a light action and so a counterfactual from this society, will indeed take the aggressive action.

[4] Despite being criticized for "an act of modern totalitarianism," many tech giants are waging wars to counter the spread of disinformation, hate speech and extremism on social media. See, e.g., https://econ.st/3e2FgmL.

Related literature. With a single society, our model predicts that due to the policy-making dilemma, the policymaker, in spite of knowing that the state is severe, may choose a light intervention, appearing to pander to public opinion when people are initially optimistic. This pander-to-the-prior effect is not new in the literature and can arise in various contexts. For example, in the political economy literature, this can occur when an incumbent politician tries to signal their competency (e.g., Harrington (1993), Canes-Wrone, Herron, and Shotts (2001), and Prat (2005)) or preferences (e.g., Maskin and Tirole (2004)) by choosing a policy which panders to public opinion. This also happens when a firm manager who has a share price concern makes decisions the market wants to see (e.g., Brandenburger and Polak (1996)), or when a media slants its report toward its readers' prior to build a reputation for quality (e.g., Gentzkow and Shapiro (2006)). Also related is the literature on the adverse effect of reputation concern. In particular, Morris (2001) and Ely and Välimäki (2003) show that an agent chooses Pareto-dominated information revelation or action to avoid damaging their reputation.

The policy-making dilemma highlighted in our paper differs from the trade-offs in the aforementioned papers. In those works, either the outcome of an action is unobservable, or when an action matches the state, it is more likely to generate an observable outcome which induces the agent to approve the chosen action. This is opposite to the "damned-either-way" feature of our model: when a strict policy is adopted, a good outcome is realized more likely regardless of the state, which makes people more optimistic about the underlying state and so induces them to disapprove the chosen policy. Moreover, most of these works do not consider a meaningful interaction among multiple decision makers. An exception is Brandenburger and Polak (1996), which we will discuss in more detail in Section 4. But an important difference is that in their model, having multiple sequential decision makers does not restore the earlier movers' incentives to take the efficient action.

The mechanism in our model with multiple societies is the resolution of the belief conflict between the policymaker and the public in a dynamic environment. Hirsch (2016) is a related work in this respect. He studies a two-period model where a principal and an agent initially disagree on the optimal policy to achieve their shared objective. If the principal compromises in the first period and implements the "wrong" policy the agent favors, the agent will then make more effort and eventually learns a more informative signal of the true state. This helps them reach consensus and implement the optimal policy more efficiently in the second period.[5] This idea could also be relevant in our crisis management context:

[5]There are other works which study the interaction between learning and prior disagreement. For instance, Che and Kartik (2009) argue that with conflicting beliefs, individuals will have more incentives to acquire in-

the policymaker can take a light action first and let people learn the true severe state, and then switch to an aggressive action. Such a trial-and-error way to resolve the belief conflict can be rather costly, especially when policy experimentation is time-consuming or an initial incorrect action could have a severe lasting adverse impact. Our model highlights a novel channel to resolve the belief conflict, which is to let the agent learn from the failure of late movers' alternative policy choices.

In our model with multiple societies, strategic players may take some action which appears against their current interest but benefits them in the long run by generating information that influences followers' decisions.[6] Broadly speaking, this insight is related to the strategic experimentation literature. For example, Bolton and Harris (1999) study a strategic bandit problem among multiple long-lived agents. Besides the standard free-rider effect, they also discover an encouragement effect by which each agent has an incentive to experiment more than in the single-agent case, in the hope of generating positive information to incentivize other agents to experiment further in the future. Callander and Hummel (2014) demonstrate that a politician holding on to power temporarily will use preemptive policy experimentation to set the path of their successor's experimentation in their favor.

The rest of the paper is organized as follows: Section 2 introduces the benchmark case with a single society and shows the policy-making dilemma. Section 3 studies the case of multiple societies where the dynamic counterfactual effect arises. Section 4 discusses an important extension, and Section 5 concludes. The omitted proofs and details are relegated to Appendix A. Appendix B contains other extensions and an alternative model.

2 Single Society

Suppose that a society faces a potential crisis, and the crisis can be *severe* or *mild*. There are two players: a policy-maker or government and a representative citizen. We assume that the government has learned the true state, which is severe. The citizen, however, is uncertain about the state, and believes that the crisis is severe with probability $\mu_0 \in (0, 1)$ and is mild with probability $1 - \mu_0$. We assume that the government cannot convince the citizen of the true severity of the crisis and fully resolve their opinion difference, and they agree to

formation to persuade each other. This can, for example, render hiring people with different opinions optimal in organization design.

[6] In terms of dynamic information spillover across players, our paper is also related to the literature on social learning. In the standard models in that literature (e.g., Banerjee (1992) and Bikhchandani, Hirshleifer, and Welch (1992)), however, early players have no strategic incentives to influence later players' choices.

disagree. This assumption implies that the citizen does not believe that the government has superior information on the state. Then the prior difference $1 - \mu_0$ can be regarded as the citizen's mistrust of the government.

The government has two possible options to handle the crisis: $a \in \{l, h\}$, where l stands for a "low" action or light intervention, and h stands for a "high" action or heavy intervention.[7] An outcome, which can be *good* (denoted by $x = 0$) or *bad* (denoted by $x = 1$), will be realized after the government takes its action. When the state is mild, we assume that the outcome will be good regardless of the government's action. When the state is severe, however, the outcome will depend on the government's action: $x = 0$ with probability q_a for $a \in \{l, h\}$ where

$$0 \le q_l < q_h \le 1 ,$$

i.e., a high action generates a good outcome more likely.[8]

The citizen observes the government action and the outcome, and updates her belief about the state by Bayes' rule. Once a bad outcome occurs, the citizen will be convinced that the state is indeed severe since a mild state always yields a good outcome.[9] A good outcome, however, will make the citizen more optimistic about the state. More precisely, after seeing $a = h$ and $x = 0$, the citizen's posterior belief about the state is

$$T_h(\mu_0) = \frac{\mu_0 q_h}{\mu_0 q_h + 1 - \mu_0} ,$$

where T_h is a Bayesian updating operator. Similarly, after seeing $a = l$ and $x = 0$, the citizen's belief is updated to

$$T_l(\mu_0) = \frac{\mu_0 q_l}{\mu_0 q_l + 1 - \mu_0} .$$

It is clear that

$$T_l(\mu_0) < T_h(\mu_0) \le \mu_0 . \quad (1)$$

[7]For simplicity we assume here that the government takes action only once. In a more realistic setting, the government is perhaps able to make decisions dynamically and the citizen then learns information on the state over time. The high action here is a reduced-form way to capture an in-time response, while the low action corresponds to a sluggish response which squanders the opportunity to keep the crisis under control while allowing the citizen to learn more about the true state.

[8]In some examples (e.g., a pandemic), the citizen's effort also matters for containing the crisis. A more optimistic citizen may make less effort, making the government's action less effective in controlling the crisis. This can strengthen our main point in the two-society model later: making people in the second society more optimistic will not only induce the government there to take the low action but also reduce people's effort there. This will increase the chance of a bad outcome in the second society and so more likely help justify the first government's choice of high action.

[9]As we will discuss in Appendix B, this assumption of bad-news information structure is not crucial for the main sights of this paper.

That is, seeing a good outcome realized when the action is low makes the citizen more optimistic than when the action is high. Also, T_a, $a = h, l$, increases in q_a and approaches the prior μ_0 as q_a goes to 1. That is, a good outcome becomes less informative when an action becomes more effective in containing the crisis. The exact form of belief updating is not important, and what matters for our subsequent analysis is property (1). Note that under our agree-to-disagree assumption, there is no interim Bayesian updating after seeing the government action but before seeing the outcome.

The citizen's utility depends on the government's action a and the outcome x:

$$u(a, x) = -c \times \mathbb{I}_{a=h} - x, \tag{2}$$

where $\mathbb{I}_{a=h}$ is an indicator function and $c > 0$. A heavy intervention imposes a cost c on the citizen, while a light intervention involves a lower cost, which is normalized to 0. When the outcome turns out to be bad, the citizen further suffers a loss, which is normalized to 1. This double normalization makes c the cost difference between two actions relative to the citizen's disutility from the bad outcome. Relative to action l, action h imposes a cost c regardless of the true state but generates a benefit $q_h - q_l$ only when the state is severe. Therefore, the citizen finds action h to be optimal if and only if she is convinced that the state is severe with a sufficiently high probability, i.e., her belief is no less than

$$\hat{\mu} \equiv \frac{c}{q_h - q_l} .$$

To make our problem interesting, we assume henceforth $c < q_h - q_l$ so that $\hat{\mu} \in (0, 1)$. Under this condition, the first-best policy, given the true state is severe, should be heavy intervention. It is easy to see that $\hat{\mu}$ increases in c and q_l, but decreases in q_h. Intuitively, action h will be less favored by the citizen if it is more costly to enforce (higher c) or less effective in containing the crisis (smaller $q_h - q_l$).

The government will be held accountable for its action *after* its consequence is observed. The citizen will evaluate the government's action according to her *posterior* belief, i.e., she prefers to approve the government's action if and only if it maximizes her expected utility based on her updated belief μ. As a consequence, action h will be approved if $\mu \geq \hat{\mu}$, and otherwise action l will be approved. We will discuss more about the citizen's evaluation rule later. The government cares only about the citizen's evaluation (say, for the policy maker's political career such as reelection or personal legacy purpose),[10] and its payoff is 1 if its

[10]This government preference specification applies not only in democracies but also in autocracies where winning public support is critical for the government to legitimate and stabilize its governance.

action is approved by the citizen and is 0 otherwise. The government's objective is therefore to maximize the probability that its action gets approved.[11]

Discussion. Before proceeding, we discuss two main modelling assumptions that enable us to deliver the main insights in a parsimonious way.

(i) Non-common prior and agree-to-disagree. We assume that the government and the citizen hold different priors of the state and agree to disagree.[12] We interpret it as a consequence of the citizen's mistrust of the government. This mistrust prevents the government from convincing the citizen of the true state. With this agree-to-disagree assumption, the citizen does not infer any information on the state directly from the government's action, and so there is no signaling issue in our model. This simplifies the analysis, especially in the dynamic case with multiple societies. For simplicity, we have also assumed that the government holds a degenerated prior belief and know the true state for sure. In Section 4 we will discuss an extension with the signaling channel and a more general prior and show that it delivers similar insights though in a less concise way. We choose to use the model without signaling also because we are inclined to believe that in many cases the public may not be sophisticated enough to make inferences based on their conjectured government policy strategy.

(ii) Policy evaluation and political accountability. The more important assumption is that the citizen uses her posterior of the state (after seeing the policy outcome) to evaluate the government policy. This is the source of the agency problem in our model (i.e., the government may pander to public opinion and choose the inefficient policy).

One interpretation is that the citizen is purely an assessor of the government policy, and she enjoys supporting the policy if it is optimal according to her posterior belief and denouncing it otherwise. That the citizen uses her posterior, instead of her prior, to evaluate the policy is related to the well-known *hindsight bias*: people tend to incorporate the newly available information into their evaluation of a decision, even if they know that the information was not available when the decision was made. This bias is widely documented in the psychology and behavioral economics literature (see, e.g., Fischhoff (1975), and Camerer, Loewenstein, and Weber (1989)), and is plausible in our context when people mistrust the government.

[11]We will discuss a more general government payoff specification in Appendix B. For example, when the government takes action h but a bad outcome is realized, it may suffer from being regarded as having a poor enforcement ability. It is also possible that conditional on being disapproved, the government may have different payoffs, depending on whether it is criticized for overreacting or underreacting. We will show that our main insights are robust to these possible generalizations.

[12]See, for example, Morris (1995) for a comprehensive discussion on the heterogeneous-prior assumption.

Another interpretation is that the policy outcome in our model is just an informative signal of the policy effectiveness realized in the beginning phase of the crisis. If the citizen believes that her opinion of the policy will determine the government's decision of whether to continue the same policy or adopt a new one, it is then rational for her to evaluate the current policy based on her posterior after seeing the signal.

As we will see more clearly later, what really matters for the main results of this paper is that the government's payoff, when it chooses action h (l), is higher if the citizen *ex post* believes the state is more likely severe (mild). We believe this is a sensible feature of the government payoff structure in the context of crisis management.

Notice also that the way we model political accountability differs from the conventional approach. The standard approach assumes the government or policy-maker has some private characteristics such as her preferences or competence. The citizen learns information on her characteristics from her policy choice and if possible also from the consequence, and then decides whether or not to reelect her.[13] On the contrary, our model assumes that the citizen assesses the policy *per se*. Certainly in many circumstances the citizen may care more about some basic characteristics of the policy-maker such as her empathy for the public, and her ability to gather relevant information and to implement policies, etc. In Appendix B, we will report an alternative reputation model in this vein, and show that it delivers similar results but somewhat different underlying economics.

2.1 Analysis

Given the citizen's evaluation rule, the government trades off being blamed for underreacting against for overreacting. Heavy intervention is more able to generate a good outcome, but this also means that it will more likely convince the citizen that the state is mild and the adoption of costly heavy intervention is unnecessary. On the contrary, light intervention will relieve the government from a criticism of overreaction, but it will more likely result in a bad outcome and therefore cause an accusation of underreaction. The more optimistic the citizen is initially, the more optimistic she will be after seeing a good outcome and so the more likely she will favor light intervention. Therefore, intuitively the government has a higher incentive to take the low action when the citizen's prior μ_0 is lower.

[13] This is the so-called forward-looking voting in the retrospective voting literature since the voter uses the information learned from past behavior to select between the incumbent politician and future challengers. See, e.g., the survey by Healy and Malhotra (2013). The other well-known strand in that literature, initiated by Key (1966) and Barro (1973), is about backward-looking voting where the voter sanctions or rewards politicians based on the outcome of their past behavior. Our modelling approach is closer to the latter in spirit.

If the government takes action h, its expected payoff, given the true state is severe, is

$$q_h \mathbb{I}_{T_h(\mu_0) \geq \hat{\mu}} + 1 - q_h \ . \tag{3}$$

When the good outcome is realized, the high action is approved if and only if $T_h(\mu_0) \geq \hat{\mu}$; when the bad outcome is realized, the high action is approved for sure since the true severe state is perfectly revealed. If the government takes action l, its expected payoff is

$$q_l \mathbb{I}_{T_l(\mu_0) < \hat{\mu}} \ . \tag{4}$$

When the good outcome is realized, the low action is approved if and only if $T_l(\mu_0) < \hat{\mu}$; when the bad outcome is realized, the low action is disapproved for sure.

Let $\hat{\mu}_1$ solve

$$T_h(\hat{\mu}_1) = \hat{\mu} \ . \tag{5}$$

It is the prior level from which the citizen's belief will be updated downward to the cut-off level $\hat{\mu}$ after seeing a high action and a good outcome. Similarly, let $\tilde{\mu}_1$ solve $T_l(\tilde{\mu}_1) = \hat{\mu}$. Note that property (1) implies $\hat{\mu} \leq \hat{\mu}_1 < \tilde{\mu}_1$.

Figure 1 plots expressions (3) and (4) and clearly shows the government's trade-off. When the high action is taken and a good outcome is realized (which occurs with probability q_h), the action will be disapproved if $T_h(\mu_0) < \hat{\mu}$, or equivalently if $\mu_0 < \hat{\mu}_1$. In this case, the government will be criticized for overreacting. In contrast, when the low action is taken and a bad outcome is realized (which occurs with probability $1 - q_l$), the government will be faulted for underreacting. Therefore, when

$$q_h > 1 - q_l \ , \tag{6}$$

the risk of being accused of overreacting dominates, and so the government will take action l if $\mu_0 < \hat{\mu}_1$. If $\mu_0 \geq \hat{\mu}_1$, the government will take the first-best action h since the citizen is too pessimistic about the state to disapprove the action even after seeing a good outcome. Figure 1 describes this case under condition (6).[14]

It is also clear from Figure 1 that when the citizen has a sufficiently optimistic prior ($\mu_0 < \hat{\mu}_1$), either action will be disapproved with some probability. This formally captures the aforementioned "damned-either-way" policy-making dilemma. This dilemma arises in our model because each action is more likely to generate an outcome which induces the citizen to believe the alternative action would be better. It does not rely on the randomness of the

[14] In the edge case with $q_h = 1$ and $q_l = 0$, the equality of (6) holds. There are many possible ways to break the tie. For example, we can assume that action h is more costly to enforce for the government than action l.

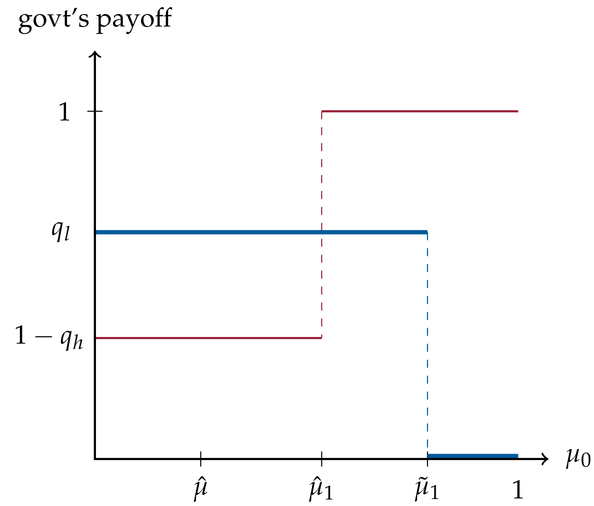

Figure 1: The red thin line corresponds to the government's payoff under action h (expression in (3)); while the blue thick line corresponds to the payoff under action l (expression in (4)), where condition (6) holds and $T_l(\tilde{\mu}_1) = \hat{\mu}$ and $T_h(\hat{\mu}_1) = \hat{\mu}$.

policy outcome. In fact the dilemma is the most prominent when the randomness vanishes, i.e., when both $q_h = 1$ and $q_l = 0$.

When condition (6) fails, the government's concern of being accused of underreacting dominates, and so it will always take the high action which achieves the first-best outcome. In the rest of the paper, we assume (6) and focus on the more interesting case depicted in Figure 1 unless otherwise stated. Then we have the following result:

Proposition 1. *The government takes action h if and only if $\mu_0 \geq \hat{\mu}_1$.*

Therefore, in the single-society case the first-best outcome is achieved if and only if the citizen is initially sufficiently pessimistic; otherwise, the government will pander to public opinion and make a sub-optimal decision.[15] Given $T_h(\cdot)$ increases in q_h and the definition of $\hat{\mu}$ in expression (2), it is easy to see that $\hat{\mu}_1$ decreases in q_h and increases in q_l and c. That is, as expected a higher q_h or a lower q_l or c widens the range of μ_0 in which the government takes the first-best action.

Implication for pandemics. We apply the single-society model to the government's optimal policy in a pandemic crisis. It helps to understand the source of political hurdles of

[15] A similar result holds when the government knows the state is mild. The government will then find it optimal to take the unnecessary high action to comfort the citizen if and only if she is sufficiently paranoid about the threat (i.e., if μ_0 is sufficiently large).

choosing strict measures.

Public mistrust. There has been a great consensus that the public's trust in technocratic expertise and professional elites is crucial in shaping a society's response to a pandemic crisis.[16] A low public trust prevents the government from convincing people of the severity of the crisis, resulting in a large divergence between the government's and people's belief (as captured by $1 - \mu_0$ in our model) and so a suboptimal policy choice. Low public trust may be caused by dysfunctional states and poor leadership, or it is simply a reflection of the polarization of a society.

Cost of strict measures. The government's hesitation to take an aggressive policy also grows as its cost on the public increases. Recall that parameter c corresponds to the cost difference between strict measures and light intervention relative to the citizen's loss from the bad crisis outcome. This relative cost is influenced by many economic and non-economic factors. First, strict measures inevitably cause significant economic damages, threatening the survival of a vast majority of people living paycheck to paycheck in societies with a low saving rate. In such a case the corresponding c should be large. By the same logic, a stimulus payment or tax relief to the public helps to lower the cost of strict measures. Second, in a society with a younger population or with more advanced critical care infrastructure, the damage caused by a pandemic is smaller, making c larger. Third, if a society has a lower tolerance of temporarily restricted civil liberties, it tends to have a larger c.

Doubts about strict measures. The government's political cost of choosing strict measures also depends on people's perceived benefit of doing so ($q_h - q_l$). Characteristics such as geographic isolation and low population density contribute to a large q_l, while controversial views on the effectiveness of strict measures (e.g. wearing face masks) may lead to a small q_h.[17] These will increase the government's incentive to take a light approach.

3 Multiple Societies

Now suppose that two *identical* societies $i = 1, 2$ face the threat of a *common* crisis sequentially. (We will consider the case with more than two societies later.) As in the single-society case, there are two players in each society: a government and a representative citizen. Each government knows the crisis is severe, while the citizen in each society initially believes that

[16] See, for example, Francis Fukuyama, "The thing that determines a country's resistance to the Coronavirus," *The Atlantic*, March 2020.

[17] See, e.g., `https://nyti.ms/2YOMNiS` on "More Americans should probably wear masks for protection" in *The New York Times*.

the crisis is severe with probability $\mu_0 \in (0, 1)$. They agree to disagree as in the single-society case. As before, each government has two possible actions $a_i \in \{l, h\}$ to choose from. The outcome in each society, which is publicly observable, depends only on the state and the government's action in that society.

The timing is as follows: Government 1 moves first and chooses its action a_1. The outcome x_1 in society 1 is then realized. After seeing a_1 and x_1, citizen 2 updates her belief of the state and government 2 chooses its action a_2. Then the outcome x_2 in society 2 is realized. Finally, citizens 1 and 2 evaluate their own government based on the information from *both* societies.

Since a bad outcome in any society perfectly reveals the true state, we need only to specify the updated belief when the outcome is good in both societies. When a citizen sees a high action and a good outcome in both societies, her posterior belief will be $T_h^{[2]}(\mu_0)$, where $T_h^{[2]}$ denotes applying the operator T_h twice. Similarly, her posterior will be $T_l^{[2]}(\mu_0)$ after seeing a low action and a good outcome in both societies, and $T_h \circ T_l(\mu_0)$ after seeing a high action, a low action and two good outcomes. Similar to (1), we have

$$T_l^{[2]} < T_h \circ T_l = T_l \circ T_h < T_h^{[2]}, \tag{7}$$

and $T_a^{[2]} \leq T_a$. As in the single-society model, government i's action will be evaluated according to citizen i's posterior μ. Action h will be approved if $\mu \geq \hat{\mu}$, and otherwise action l will be approved.

It is worth pointing out two implicit assumptions in this two-society model: First, we assume a common prior for the public across the two societies.[18] The case with heterogeneous priors can be analyzed similarly, but does not add particularly new insights. Second, we also assume that the crisis will arise in the second society regardless of the action and the outcome in the first society. This is not crucial as long as a high action or a good outcome in the first society does not completely halt the spread of the crisis.

3.1 Analysis

Let $\hat{\mu}_2$ solve $T_h(\hat{\mu}_2) = \hat{\mu}_1$, or equivalently

$$T_h^{[2]}(\hat{\mu}_2) = \hat{\mu} .$$

This is the prior level from which the citizen's belief will be updated downward to the cut-off level $\hat{\mu}$ after seeing a high action and a good outcome in both societies. Clearly we have

[18] With a common prior, our subsequent analysis remains unchanged if there is only a "common" citizen in both societies, or a common principal facing two agents.

that $\hat{\mu} \leq \hat{\mu}_1 \leq \hat{\mu}_2$. A condition we will often refer to in the subsequent analysis is

$$q_l(q_h + q_l) \leq 1 \,. \tag{8}$$

For any given q_h, this condition holds if q_l is sufficiently small.

The following result reports the equilibrium outcome of the two-society game:

Proposition 2. *When there are two societies,*

(i) if $\mu_0 \geq \hat{\mu}_2$, both governments take action h;

(ii) if $\mu_0 < \hat{\mu}_2$, government 1 takes action h if and only if condition (8) holds, and government 2 takes action h if and only if a bad outcome is realized in the first society.

Proposition 2 identifies the condition under which government 1 is relieved from the policy-making dilemma. When $\mu_0 \geq \hat{\mu}_2$ (which implies $T_h^{[2]}(\mu_0) \geq \hat{\mu}$), citizens are sufficiently pessimistic, and governments can safely choose the first-best action without being blamed for overreacting.

The more interesting case is when $\mu_0 < \hat{\mu}_2$. In this case, the presence of society 2 may help government 1 by providing a counterfactual for citizen 1 to better see the consequence of different policy options. More precisely, let us analyze the problem backward. First, notice that a successful crisis management by government 1 always makes citizen 2 more optimistic. With $\mu_0 < \hat{\mu}_2$, we have $T_{a_1}(\mu_0) < \hat{\mu}_1$ for any action $a_1 \in \{h, l\}$. From the analysis in the single-society case, it is then immediate that government 2 will take action l if $x_1 = 0$ and action h if $x_1 = 1$, regardless of government 1's action. Next, we consider government 1's incentive. If government 1 takes action h, it will be approved if and only if the crisis is out of control in at least one society, which happens with probability $1 - q_h q_l$. Instead, if government 1 takes action l, it will be approved if and only if both societies succeed in containing the crisis, which occurs with probability q_l^2. Therefore, government 1 prefers action h if and only if $1 - q_h q_l \geq q_l^2$, i.e., if condition (8) holds.

By comparing Propositions 1 and 2, it is easy to see that society 1 can either benefit or suffer from the presence of the second society.

Corollary 1. *Having the second society induces government 1 to switch from taking action l to taking action h if $\mu_0 < \hat{\mu}_1$ and (8) holds, and the reverse is true if $\hat{\mu}_1 \leq \mu_0 < \hat{\mu}_2$ and (8) does not hold. In the remaining cases, having the second society has no impact on government 1's policy choice.*

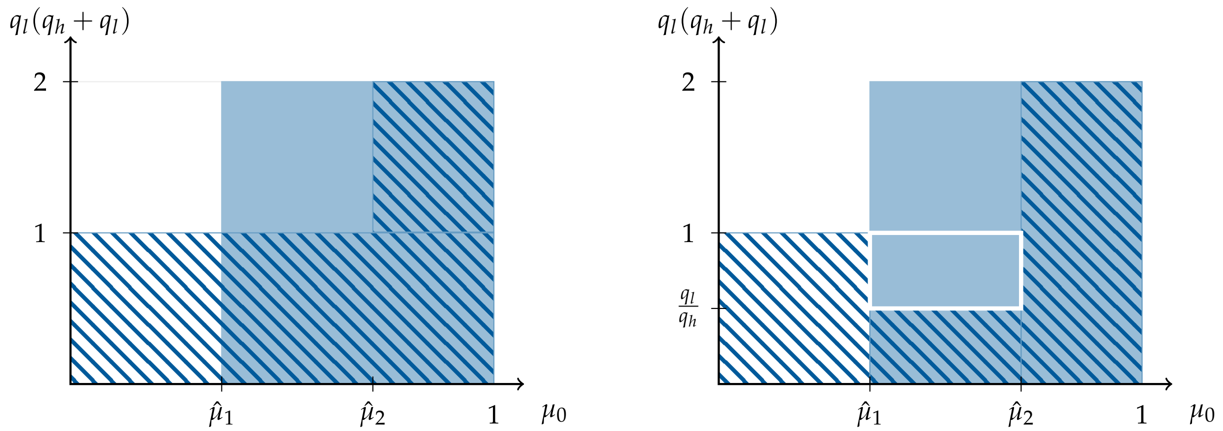

(a) Government 1's choice: equilibrium model.   (b) Government 1's choice: hypothetical situation.

Figure 2: The solid rectangle area describes the set of parameters where government 1 chooses action h in the single-society model. The hatched area in (a) describes the set of parameters where government 1 chooses action h in the two-society model. The region $(0, \hat{\mu}_1) \times [0, 1]$ corresponds to the positive sampling effect, while the region $[\hat{\mu}_1, \hat{\mu}_2) \times [q_l/q_h, 2]$ corresponds to the negative sampling effect. The white-framed region $[\hat{\mu}_1, \hat{\mu}_2) \times [q_l/q_h, 1]$ in (b) corresponds to the strategic effect.

This result is illustrated in Figure 2a. Three forces influence government 1's decision: First, independent of government 1's action, the presence of society 2 increases the chance that the true severe state is revealed. This encourages government 1 to take action h. We call this a *positive sampling effect*. Second, independent of government 1's action, the presence of society 2 also generates the possibility that the good outcome is realized in both societies, in which case citizens will become rather optimistic. This is called a *negative sampling effect*, and it encourages government 1 to take action l instead. Finally, government 1's action can influence citizen 2's interim belief and so government 2's policy. In particular, when it takes action h, it makes it more likely that citizen 2 becomes optimistic so that government 2 takes action l, which increases the chance that the true state is revealed and government 1's policy is justified. We call this third effect a *strategic effect*.

To disentangle the strategic effect from the two sampling effects, we consider the following hypothetical situation: suppose that citizen 2 *cannot* observe what has happened in society 1, but citizen 1 can observe how the crisis unfolds in society 2. In this case, government 1's policy choice is only affected by the sampling effects. Following a similar argument as in the proof of Proposition 2, one can readily show that government 1's action can be different when $\hat{\mu}_1 \leq \mu_0 < \hat{\mu}_2$, in which case it will take action h if and only if

$$q_h(q_h + q_l) \leq 1 , \tag{9}$$

which is equivalent to $q_l(q_h + q_l) \leq q_l/q_h$. This is a more stringent condition than (8). The gap between these two conditions, as illustrated as the white-framed region in Figure 2b, captures the strategic effect. Intuitively, when $\hat{\mu}_1 \leq \mu_0 < \hat{\mu}_2$, if citizen 2 does not observe what has happened in society 1, her pessimistic prior induces government 2 to take action h. While in our model, after seeing $x_1 = 0$ citizen 2 will become sufficiently optimistic ($T_h(\mu_0) < \hat{\mu}_1$) so that government 2 will take action l instead, which increases the chance that the true state is revealed.

Next, we compare all parties' (ex ante) welfare between the two societies to see whether there is a first-mover advantage in our model. The citizen's welfare is measured according to the true state.

Corollary 2. *If $\mu_0 \geq \hat{\mu}_2$, both the citizen and the government are equally well across societies. If $\mu_0 < \hat{\mu}_2$ and (8) holds, citizen 1 does better than citizen 2, and government 1 does better than government 2 if and only if $2q_l \leq 1$. If $\mu_0 < \hat{\mu}_2$ and (8) does not hold, both the citizen and the government in society 2 do better than in society 1.*

The first society can influence the second society's belief and action in favor of its own welfare. But the second society has more information when it is its turn to make the decision. This, however, is not always a blessing, given the citizen's welfare is measured according to the true state instead of her own belief. When a good outcome is realized in the first society, it will mislead citizen 2 to be over-optimistic. Of course, when a bad outcome is realized in the first society, it helps the second society. The above result suggests that if the citizen is initially sufficient optimistic ($\mu_0 < \hat{\mu}_2$) and the low action has a sufficiently small chance in containing the crisis (q_l sufficiently small), the first society has the first-mover advantage. (Recall that a smaller q_l amplifies the strategic effect as action l by government 2 will reveal the true state more likely.)

Implication for pandemics. The two-society model also has some useful implications for the pandemic crisis.

Misleading success. A key force in our model is that a successful crisis control in the first society will give people in the second society a false sense of safety. This confidence inflation is greater if, all else equal, the first society is perceived as less developed in health infrastructure (corresponding to smaller q_{a_1}). In this case the political hurdle to strict measures in the second society will be larger.

Information manipulation. We assume free information flow across countries, but it is straightforward to see that governments have incentives to influence the information flow. Government 1, if it has succeeded in containing the crisis, can benefit from downplaying the threat by broadcasting its success to inflate people's doubts in society 2 about the severity of the crisis. Meanwhile, it also has an incentive to broadcast the failure of society 2's light intervention to its own people to justify its strict policy. Government 2, on the other hand, has an incentive to downplay or sow doubt about society 1's success to minimize the confidence inflation among its own people and therefore the political hurdle to strict measures.

Commitment to strict measures. The first mover takes advantage from the aforementioned strategic effect by pushing up the political hurdle faced by government 2. To counter this force, government 2 could make a preemptive commitment to strict measures. As shown in Figure 2b, such a commitment will turn the tables in some circumstances: it can induce government 1 to take the low action and so help justify government 2's aggressive policy.

3.2 More societies

It is not difficult to extend our analysis to the case with n societies. The main insights remain, but the general case also yields some new insights such as the societies in the middle of the sequence may perform the worst. We first report the equilibrium in this general case.

Proposition 3. *Let $\hat{\mu}_n$ solve $T_h^{[n]}(\hat{\mu}_n) = \hat{\mu}$. When there are n societies,*

(i) if $\mu_0 \geq \hat{\mu}_n$, all the governments take action h;

(ii) if $\mu_0 < \hat{\mu}_n$, for any $i = 1, 2, ..., n$, (a) if $x_j = 1$ in at least one predecessor society $j < i$, government i takes action h; (b) if $x_j = 0$ for all $j < i$ or if $i = 1$, government i takes action h if and only if $q_l^{n-i}(q_h + q_l) \leq 1$.

A few simple observations follow. First, since $\hat{\mu}_n$ increases in n, result (i) implies that it becomes harder for all the governments to take action h when there are more societies. Second, when $\mu_0 < \hat{\mu}_n$, having more societies increases the i_{th} government's incentive to take action h more likely if the true state has not been revealed. This is because, when there are more societies, both the positive sampling effect (i.e., the true state is revealed in some subsequent society) and the strategic effect become stronger, while the negative sampling effect (i.e., the good outcome is realized in all the subsequent societies) becomes weaker. Similarly, when $\mu_0 < \hat{\mu}_n$ and the history is good so far, earlier governments are more likely to take action h.

To illustrate the point that the society in the middle of the sequence may perform the worst, consider an example with three societies and suppose $\mu_0 < \hat{\mu}_3$ and $q_l^2(q_h + q_l) \leq 1 < q_l(q_h + q_l)$. From the proposition above, it is easy to see that government 1 will take action

h, which is the best for its citizen. If $x_1 = 1$, then the true state is revealed and so the other two governments will take action h as well, in which case the three societies are equally well. If $x_1 = 0$, government 2 will take action l for sure, which is the worst for its citizen, while government 3 will take action h with some chance (i.e., when $x_2 = 1$), which puts its citizen in the middle of the ranking. Intuitively, the first society enjoys the greatest positive sampling effect and the strategic effect, while the third society has the most information and its government will take the first-best policy if the true state has been revealed.

4 Extension: Asymmetric state information and signaling

This section discusses an extension where the citizen believes that the government possesses superior information concerning the state. As a result, the citizen will attempt to infer the state from the government action as well, based on her rational expectation of the government's policy strategy. This signaling channel was intentionally shut down in our main model. Here we demonstrate that having this signaling channel does not change the main insights except in the polar case when the government perfectly knows the true state.

Let us first consider the single-society case. Let ω denote the state, and it can be bad/severe (B) or good/mild (G). The citizen's prior is $\Pr(\omega = B) = \mu_0$, and the government's prior is $\Pr(\omega = B) = \nu_0$. When $\mu_0 = \nu_0$, we have the common-prior case, but as we will see these two priors play completely separate roles, and so here we present the more general case. Before choosing its policy the government observes a private signal $s \in \{b, g\}$ of the state, and it is commonly known that the signal structure is $\Pr(b|B) = \Pr(g|G) = \delta \in [\frac{1}{2}, 1]$.It is assumed that the government is unable to convey its private information directly to the citizen (e.g., because the state information is too complex to communicate). The other aspects of the model remain unchanged, and we focus on the interior case with $0 < q_l < q_h < 1$. The major difference now is that the citizen can also infer some information on the state from the government's action alone based on her equilibrium belief of the government's strategy. The government strategy is denoted by $\sigma \equiv (\sigma_b, \sigma_g)$, where σ_s is the probability the government takes action h after receiving signal s.

Let

$$\nu_s = \frac{\nu_0 \Pr(s|B)}{\nu_0 \Pr(s|B) + (1 - \nu_0) \Pr(s|G)}$$

be the government's updated belief of the state after receiving signal s. This is what matters for its policy decision. The interesting case is when $\nu_g < \hat{\mu} < \nu_b$, i.e., when the efficient policy, from the government's point of view, is h (l) after seeing a bad (good) signal. This requires δ to be sufficiently high.

For the citizen, what matters for her evaluation of the government policy is her posterior after seeing both the action and the outcome. If the outcome is bad ($x = 1$), it perfectly reveals state B. If the outcome is good ($x = 0$), let $T_a^\sigma(\mu_0)$ denote the citizen's posterior when the government takes action a and it is believed to be using strategy σ. Specifically,

$$T_h^\sigma(\mu_0) = \frac{\mu_0 \sigma_B q_h}{\mu_0 \sigma_B q_h + (1-\mu_0)\sigma_G} = \frac{1}{1 + \frac{1-\mu_0}{\mu_0}\frac{\sigma_G}{\sigma_B}\frac{1}{q_h}},$$

and

$$T_l^\sigma(\mu_0) = \frac{\mu_0(1-\sigma_B) q_l}{\mu_0(1-\sigma_B) q_l + (1-\mu_0)(1-\sigma_G)} = \frac{1}{1 + \frac{1-\mu_0}{\mu_0}\frac{1-\sigma_G}{1-\sigma_B}\frac{1}{q_l}},$$

where $\sigma_B \equiv \delta\sigma_b + (1-\delta)\sigma_g$ and $\sigma_G \equiv (1-\delta)\sigma_b + \delta\sigma_g$ are respectively the expected probability that the government takes action h under strategy (σ_b, σ_g) when the true state is B or G. (We stipulate $1/0 = \infty$.) Both posteriors are increasing in μ_0 as in the baseline model, but now they also depend on the government's policy strategy σ. Let $\hat{\mu}_a^\sigma$ solve $T_a^\sigma(\mu) = \hat{\mu}$ whenever this is well defined, i.e., it is the prior from which the citizen's belief will be updated to the threshold level $\hat{\mu}$ after seeing action a and a good outcome. (Note that $\hat{\mu}_1$ and $\tilde{\mu}_1$ in the baseline model are respectively equal to $\hat{\mu}_h^{1,1}$ and $\hat{\mu}_l^{0,0}$.)

For a given signal s, the government's decision is similar as in the baseline model: its expected payoff is

$$\pi_{s,h} \equiv p_{s,h}\mathbb{I}_{T_h^\sigma(\mu_0)\geq\hat{\mu}} + 1 - p_{s,h}$$

if it takes action h, and is

$$\pi_{s,l} \equiv p_{s,l}\mathbb{I}_{T_l^\sigma(\mu_0)<\hat{\mu}}$$

if it takes action l, where $p_{s,a} \equiv \nu_s q_a + 1 - \nu_s$ is the government's expected probability, after receiving signal s, that a good outcome will be realized under action a. The government's optimal strategy is then determined by comparing $\pi_{s,h}$ and $\pi_{s,l}$, and in equilibrium it should be consistent with σ.

We maintain the assumption $q_h + q_l > 1$ as in the baseline model. Then it is ready to check that we must have $p_{s,h} + p_{s,l} > 1$, and so the payoffs $\pi_{s,a}$, as functions of μ_0, are similar to those in Figure 1 (with q_l replaced by $p_{s,l}$ and $1 - q_h$ replaced by $1 - p_{s,h}$). In particular, for any μ_0 it is impossible that $\pi_{s,h} = \pi_{s,l}$ given our payoff specification. This implies that in our model the government will never play a mixed strategy.

In the polar case with $\delta = 1$ (i.e., when the government's signal perfectly reveals the true state), it is easy to see that there is a separating equilibrium with $\sigma_b = 1$ and $\sigma_g = 0$ (i.e., the government takes action h (l) for sure upon seeing a bad (good) signal). In this equilibrium, the citizen perfectly infers the state from the government's action (and the outcome information becomes redundant), and the government action is always proved. Notice, however, this equilibrium is not "strict" in the sense that given the citizen's belief the government is actually indifferent between the two actions. [19]

Once we go beyond this polar case, however, the separating equilibrium with $\sigma_b = 1$ and $\sigma_g = 0$ can no longer be sustained. The intuition is easy to see when $\delta < 1$ is sufficiently close to 1. Suppose that the citizen holds the belief that the government is playing the separating strategy. If the government takes action h, its expected payoff is 1 regardless of the signal it receives, since the citizen will infer the state is B very likely and so will approve the action even if a good outcome is realized. If the government takes action l, however, its expected payoff must be strictly below 1. This is because no matter what signal it receives, the government is never perfectly sure that the state is G given $\delta < 1$ and so there is always a chance that the state is B and a bad outcome arises, in which case its low action will be disapproved. In other words, action h is always a safer option for the government. This contradicts the separating strategy.

More formally, when $\delta < 1$, in the appendix we prove the following results for any given μ_0: *there are no separating equilibria with* $\sigma_s = 1$ *and* $\sigma_{s'} = 0$*; any equilibrium must be a pure-strategy pooling equilibrium*. In particular, there is a pooling equilibrium with proper off-equilibrium beliefs in which regardless of its private signal, the government always takes action h (i.e., $\sigma_b = \sigma_g = 1$) if $\mu_0 \geq \hat{\mu}_1$ and action l (i.e., $\sigma_b = \sigma_g = 0$) if $\mu_0 < \hat{\mu}_1$, where $\hat{\mu}_1 = \hat{\mu}_h^{1,1}$ takes the same value as (5) in the baseline model. In this equilibrium, the government's policy choice alone does not convey any information on the state, and so the outcome is the same as in our baseline model and features a similar policy-making dilemma.

Now consider the two-society case where the state is common but each government receives an independent signal of the state with precision $\delta < 1$. We show that having the second society can enable the first government to take the efficient separating strategy $\sigma_b = 1$ and $\sigma_g = 0$. The single-society argument implies that there is only a pooling equilibrium in the second society, and suppose it takes the form specified above. Then as we show in the appendix, the desired outcome arises when $T_h^{1,1} \circ T_h^{1,0}(\mu_0) < \hat{\mu}$ and $p_{b,l}(p_{b,h} + p_{b,l}) < 1 < p_{g,l}(p_{g,h} + p_{g,l})$. These conditions are qualitatively similar to those in Proposition 2 in the baseline model. The first condition, which holds if μ_0 is sufficiently low, ensures that after seeing a high action and a good outcome in society 1, people in society

[19] Another subtlety is that on the off-equilibrium path when the government deviates to action l at state B, if a bad outcome is realized, the action and the outcome generate completely opposite inferences so that the citizen's Bayesian updating is not well defined.

2 will be optimistic enough so that their government will take action l. The second condition holds when q_l is sufficiently low, and ν_b (ν_g) is sufficiently high (low), which is the case if δ is sufficiently high. It ensures that government 1 will take the efficient strategy, anticipating a good outcome will induce a low action in society 2. Intuitively, when government 1 is sufficiently confident that the state is bad (i.e., ν_b is high), it believes that an induced low action in society 2 will tend to generate a bad outcome, which will help justify its choice of high action; in contrast, when government 1 is sufficiently confident that the state is good (i.e., ν_g is low), it believes that even a low action will tend to generate a good outcome, and together with the same likely outcome in society 2 this will justify its choice of low action.

Discussion. This extended model with signaling is related to Brandenburger and Polak (1996) (BP thereafter). They study how a firm may make decisions the market wants to see. In their model, the firm receives a private signal of the state and then takes an action (e.g., an investment decision) which generates a profit if it matches the state. After seeing the firm's action (but not the resulting outcome), the market updates its belief on the state and then assesses the firm. The market assessment determines the firm's share price which the firm aims to maximize. BP show that if the prior is skewed to one state, there is no equilibrium in which the firm plays a separating strategy and maximizes its own expected profit. Instead in any equilibrium the firm's decision panders to the market's prior to some extent. This remains true even in the case with multiple firms which observe independent signals and make sequential decisions and are all evaluated by the market in the end.

Our model shares some features with BP: the market (the public) assesses a firm (a government) based on its posterior of the state instead of some underlying fundamental (e.g., the manager/government's ability), and the actions taken by early firms (governments) influence later firms' (governments') information and decisions. Nevertheless, the two papers differ in several important aspects. First, in BP both states and actions are symmetric (i.e., they can be relabelled), while in our model they are asymmetric. This asymmetry is natural in our context of crisis prevention (e.g., regardless of the state the high action prevents a crisis more likely than the low action). Second, in BP the market does not observe the outcome of the action and it only infers the state from the firm's action, while the observable outcome plays an important role in our model. Third, the above two differences imply that the policy-making dilemma in our single-society model does not occur in BP, and the result that having another society can restore the first government's incentive to take the efficient action does not arise in BP either.

5 Conclusion

This paper provides a framework for studying crisis management with multiple jurisdictions or societies. We first highlight a "damned-either-way" policy-making dilemma: sufficient precautions can contain a crisis, but people may then become skeptical of the severity of the problem and question the costly response; light intervention is less costly but often fails to control the crisis, and people will then accuse the policymaker of underreacting. Such a dilemma can raise the political cost for the government to take an efficient policy. We then argue that the dilemma will be mitigated if people can see the counterfactual policy outcome. One possibility is when another society faces the same crisis afterward. The success under an aggressive policy in the first society boosts the optimism of people in the second society, increasing the chance for the second society to adopt a light approach and experience an outbreak, which in turn justifies the first society's policy choice. This helps explain, for example, why similar societies might respond to a common crisis differently, and why societies that handle the crisis later may perform worse despite having more information.

This paper has focused on "preventive" policies that if succeed cause the public to question the severity of a potential crisis. In the pandemic example, they are policies such as wide testing and tracking, mandatory quarantine, travel bans and even strict lock-downs from the early stage. Another type of policies which we do not study in this paper are "mitigating" policies that aim to reduce the damage of a crisis when it already outbreaks and its severity is already known. In the pandemic example, they are policies such as stockpiling ventilators, subsidizing medicine and vaccine research, and stimulus payment. When the government is constrained by the policy-making dilemma from implementing preventive policies, it may then reply more on mitigating measures.[20] The counterfactual effect predicts that early-hit countries tend to focus more on preventive measures while later countries on mitigating measures.

Another interesting perspective is to consider countries with different cultures, institutions, or public infrastructures, etc. Depending on which countries are hit first by a crisis, the dynamics of crisis management may vary significantly, resulting in rather different welfare outcomes.

[20] Fox and Van Weelden (2015) study a model of crisis prevention when a policy maker can allocate effort across multiple tasks, but the essential economic force there is different from ours.

A Appendix: omitted proofs and details

Proof of Proposition 2. (i) Suppose $\mu_0 \geq \hat{\mu}_2$. If government 1 takes action h and the outcome is $x_1 = 1$, both citizens learn the true state, in which case government 1's action will be approved and government 2 will take action h as well. If government 1 takes action h and the outcome is $x_1 = 0$, then citizen 2's interim belief will be $T_h(\mu_0) \geq \hat{\mu}_1$, in which case government 2 will take action h as well. The two citizens' posterior will then be at least $T_h^{[2]}(\mu_0) \geq \hat{\mu}$, and so both governments' actions will be approved. Therefore, if government 1 takes action h, it will always get approved and have an expected payoff 1.

Conversely, if government 1 takes action l, with probability $1 - q_l$ the true severe state will be revealed, in which case its action will be disapproved, and so its payoff is at most $q_l < 1$. Therefore, government 1's optimal choice is action h. The above argument then implies that government 2 will take action h as well.

(ii) Suppose now $\mu_0 < \hat{\mu}_2$. If government 1 takes action l, it will be approved if and only if $x_1 = x_2 = 0$. When $x_1 = 0$, citizen 2's interim belief will be updated to $T_l(\mu_0) < \hat{\mu}_1$, and so government 2 will take action l. Therefore, $x_1 = x_2 = 0$ occurs with probability q_l^2, and this is government 1's expected payoff.

If government 1 takes action h, with probability $1 - q_h$, $x_1 = 1$, in which case government 1's payoff is 1. With probability q_h, $x_1 = 0$, in which case citizen 2's interim belief will be $T_h(\mu_0) < \hat{\mu}_1$, and so government 2 will take action l. For government 1' action h to be approved, we need $x_2 = 1$, which happens with probability $1 - q_l$. (Otherwise, the citizen's posterior would be $T_h \circ T_l(\mu_0) < \hat{\mu}$ and she would not approve action h.) Hence, government 1's expected payoff is $(1 - q_h) + q_h(1 - q_l) = 1 - q_h q_l$.

Therefore, government 1's optimal choice is action h if and only if $q_l^2 \leq 1 - q_h q_l$, which is equivalent to (8). □

Proof of Corollary 2. When $\mu_0 \geq \hat{\mu}_2$, both governments take the same action h, and so all parties' expected payoff must be the same across societies.

When $\mu_0 < \hat{\mu}_2$ and (8) holds, government 1 takes action h while government 2 takes action h if and only if $x_1 = 1$. So citizen 1 must do better given the high action is the first-best action. As we have shown in the proof of Proposition 2, in this case government 1's payoff is $1 - q_h q_l$. Government 2's payoff is $1 - q_h + q_h q_l$. (When $x_1 = 1$, government 2 will take action h, in which its payoff is 1. When $x_1 = 0$, government 2 will take action l, in which case its action will be approved if and only if $x_2 = 0$.) Comparing these two payoffs yields the condition stated in the result.

When $\mu_0 < \hat{\mu}_2$ and (8) does not hold, government 1 takes action l while government 2 takes h if $x_1 = 1$. So the citizen in society 2 must do better. As we have shown in the proof of Proposition 2, in this case government 1's payoff is q_l^2. Government 2's payoff is at least q_l^2 because when $x_1 = 0$ government 2 will take action l, and this will be approved by the citizen if $x_2 = 0$. □

Proof of Proposition 3. (i) Let us use induction and suppose the claim is true when there are $n-1$ societies. If $a_1 = h$, then all the citizens in the subsequent societies will update their interim beliefs to $\mu = 1$ (if $x_1 = 1$) or $\mu = T_h^{[1]}(\mu_0) \geq \hat{\mu}_{n-1}$ (if $x_1 = 0$). In either case, according to the induction assumption, all the subsequent governments will take action h. When $x_1 = 1$, government 1 gets 1; when $x_1 = 0$, it gets 1 as well because even if $x_j = 0$ for all $j > 1$ the posterior will be $T_h^{[n]}(\mu_0) \geq \hat{\mu}$. Hence, government 1's expected payoff, when it takes the high action, is 1. If $a_1 = l$, government 1 gets zero if $x_1 = 1$, and so its payoff is at most q_l. Therefore, government 1 should take action h.

(ii) Part (a) is obvious given a bad outcome in any society reveals the true severe state. Again we use induction and suppose (b) is true when there are $n-1$ societies. Consider government 1's decision when there are n societies. There are $n-1$ cases. We call the case of $q_l(q_h + q_l) \leq 1$ "case 1," the case of $q_l^k(q_h + q_l) \leq 1 < q_l^{k-1}(q_h + q_l)$ "case k" if $2 \leq k \leq n-2$, and the case of $q_l^{n-2}(q_h + q_l) > 1$ "case $n-1$."

If government 1 takes action h, it will be approved if and only if the true severe state is revealed at some point. With probability $1 - q_h$, $x_1 = 1$, in which case government 1 gets 1. With probability q_h, $x_1 = 0$, in which case the citizens in the subsequent societies have an interim belief $T_h^{[1]}(\mu_0) < \hat{\mu}_{n-1}$ and so the induction assumption can be applied. Then government 1's payoff depends on how many subsequent governments will take action h and how many will take action l. In case 1, all the subsequent governments but the last one will take action h if the history is good so far. So among them the chance that the true state is revealed is $1 - q_h^{n-2} q_l$. Then government 1's payoff is $1 - q_h + q_h(1 - q_h^{n-2} q_l) = 1 - q_h^{n-1} q_l$. In case k, all the subsequent governments but the last k will take action h if the history is good so far. Then government 1's payoff is $1 - q_h^{n-k} q_l^k$. In case $n-1$, all the subsequent government will take action l if the history is good so far. Then government 1's payoff is $1 - q_h q_l^{n-1}$.

If government 1 takes action l instead, it will be approved if and only if the true state is never revealed. With probability $1 - q_l$, $x_1 = 1$, in which case its payoff is zero. With probability q_l, $x_1 = 0$, in which case the citizens in the subsequent societies have an interim belief $T_l^{[1]}(\mu_0) < \hat{\mu}_{n-1}$ and so the induction assumption can be applied. The analysis is then similar as above. In case 1, all the subsequent governments but the last one will take action h if the history is good so far, and so the chance that $x_i = 0$ among all the subsequent societies is $q_h^{n-2}q_l$. Thus, government 1's payoff is $q_h^{n-2}q_l^2$. In case k, all the subsequent governments but the last k will take action h if the history is good so far, and so government 1's payoff is $q_h^{n-k-1}q_l^{k+1}$. In case $n-1$, all the subsequent government will take action l if the history is good so far, and so government 1's payoff is q_l^n.

It is then straightforward to verify: in case 1, government 1 prefers action h if and only if $q_h^{n-2}q_l(q_h+q_l) \leq 1$, which is implied by the condition of case 1; in case k, government 1 prefers action h if and only if $q_h^{n-k-1}q_l^k(q_h+q_l) \leq \tau$, which is also implied by the condition of case k; in case $n-1$, government 1 prefers action h if and only if $q_l^{n-1}(q_h+q_l) \leq 1$, which implies the condition of case $n-1$. Therefore, we can conclude that government 1 will take action h if and only if $q_l^{n-1}(q_h+q_l) \leq 1$. This completes the proof. □

Omitted details in Section 4. In the single-society case, we show that if $\delta < 1$ there is no equilibrium with $\sigma_b = 1$ and $\sigma_g = 0$. (The other counter-intuitive case with $\sigma_b = 0$ and $\sigma_g = 1$ is easy to rule out.) If the citizen believes that the government is playing the above separating strategy, her posterior after seeing a good outcome will be

$$T_h^{1,0}(\mu_0) = \frac{1}{1+\frac{1-\mu_0}{\mu_0}\frac{1-\delta}{\delta}\frac{1}{q_h}}; \;\; T_l^{1,0}(\mu_0) = \frac{1}{1+\frac{1-\mu_0}{\mu_0}\frac{\delta}{1-\delta}\frac{1}{q_l}}.$$

When $\delta < 1$, both are well-behaved strictly increasing functions. It is also clear that given $\delta \geq \frac{1}{2}$ and $q_h > q_l$, we must have $\frac{1-\delta}{\delta}\frac{1}{q_h} < \frac{\delta}{1-\delta}\frac{1}{q_l}$ and so $T_h^{1,0}(\mu_0) > T_l^{1,0}(\mu_0)$. This implies $\hat{\mu}_h^{1,0} < \hat{\mu}_l^{1,0}$. From a graph similar to Figure 1, it is ready to see that regardless of signal s, we have $\pi_{s,h} > \pi_{s,l}$ if $\mu_0 \geq \hat{\mu}_h^{1,0}$ and $\pi_{s,h} < \pi_{s,l}$ otherwise. Therefore, for a given μ_0, the separating strategy cannot be sustained in equilibrium. Essentially, this is because given $p_{s,h} + p_{s,l} > 1$, the ranking of $\pi_{s,h}$ and $\pi_{s,l}$ is independent of the signal s.

We have explained in the main text that there is no mixed-strategy equilibrium. Hence, only pure-strategy pooling equilibria remain possible. Suppose first the citizen believes that the government's policy strategy is $\sigma_b = \sigma_g = 1$ (i.e., it always takes action h regardless of its signal). Then

$$T_h^{1,1}(\mu_0) = \frac{1}{1+\frac{1-\mu_0}{\mu_0}\frac{1}{q_h}}\,. \tag{10}$$

Let us specify the off-equilibrium belief so that $T_l^{1,1}(\mu_0) < T_h^{1,1}(\mu_0)$. (This is reasonable since without the signaling channel a low action with a good outcome is more convincing evidence that the state is good.) Then we have $\hat{\mu}_h^{1,1} < \hat{\mu}_l^{1,1}$, and so the government will indeed always take action h if $\mu_0 \geq \hat{\mu}_h^{1,1}$. Now consider the case when the citizen believes that the government's policy strategy is $\sigma_b = \sigma_g = 0$ (i.e., it always takes action l regardless of its signal). Then

$$T_l^{0,0}(\mu_0) = \frac{1}{1 + \frac{1-\mu_0}{\mu_0}\frac{1}{q_l}} \ . \tag{11}$$

A reasonable off-equilibrium belief is $T_h^{0,0}(\mu_0) > T_l^{0,0}(\mu_0)$, in which case we have $\hat{\mu}_h^{0,0} < \hat{\mu}_l^{0,0}$ and so the government will indeed always take action l if $\mu_0 < \hat{\mu}_h^{0,0}$. If we assume $T_h^{0,0}(\mu_0)$ takes the same form as (10) (which can be justified if both σ_b and σ_g converge to 0 at the same speed), then $\hat{\mu}_h^{0,0} = \hat{\mu}_h^{1,1}$. This is the pooling equilibrium described in the main text.

In the two-society case, let us consider the possibility of the equilibrium where government 1 adopts the efficient separating strategy $\sigma_b = 1$ and $\sigma_g = 0$. When the citizen in either society believes that government 1 is taking this strategy, her belief of the state, after seeing action a and a good outcome, is updated to $T_a^{1,0}(\mu_0)$. From the definition of $\hat{\mu}_h^{1,1}$, we can see that $T_a^{1,0}(\mu_0) < \hat{\mu}_h^{1,1}$ if and only if $T_h^{1,1} \circ T_a^{1,0}(\mu_0) < \hat{\mu}$. This is true for both $a = h$ and $a = l$ if

$$T_h^{1,1} \circ T_h^{1,0}(\mu_0) < \hat{\mu} \ . \tag{12}$$

Under this condition, following a similar argument as in the baseline model, we can see that when government 1 takes action h after seeing signal s, its expected payoff is

$$\pi_{s,h} = 1 - p_{s,h}p_{s,l}\left(1 - \mathbb{I}_{T_l^{0,0} \circ T_h^{1,0}(\mu_0) \geq \hat{\mu}}\right) \ ;$$

when it takes action l after seeing signal s, its expected payoff is

$$\pi_{s,l} = p_{s,l}^2 \mathbb{I}_{T_l^{0,0} \circ T_l^{1,0}(\mu_0) < \hat{\mu}} \ .$$

Notice that $T_h^{1,1} > T_l^{0,0}$ and $T_h^{1,0} > T_l^{1,0}$, and so (12) implies $T_l^{0,0} \circ T_l^{1,0}(\mu_0) < T_l^{0,0} \circ T_h^{1,0}(\mu_0) < \hat{\mu}$. Therefore, $\pi_{b,h} > \pi_{b,l}$ if $p_{b,l}(p_{b,h} + p_{b,l}) < 1$, and $\pi_{g,h} < \pi_{g,l}$ if $p_{g,l}(p_{g,h} + p_{g,l}) > 1$.

B Appendix: other extensions and discussions

In this appendix, we report two other extensions: one with a more general government payoff structure, and the other with a more general information structure; and we also explore an alternative "reputation" model.

B.1 More general government payoff

In the baseline model, we assume a simple payoff structure for the government: it gets 1 if its action is approved and 0 otherwise. We now consider a more general payoff structure as in the table below:

	$x = 0$	$x = 1$
a approved	1	$\beta \in [0, 1]$
a disapproved	α_o if $a = h$ and $T_h(\mu_0) < \hat{\mu}$ α_u^+ if $a = l$ and $T_l(\mu_0) \geq \hat{\mu}$	$\alpha_u^- \equiv 0$

In the first cell, the action is approved and the outcome is good, in which case the government gets the highest possible payoff 1. In the second cell, the action is approved but the outcome is bad, which can happen only if $a = h$. In this case the citizen may doubt the government's enforcement ability, and we assume the government's payoff is $\beta \in [0, 1]$. In the third cell, the action leads to a good outcome but it is disapproved. If the action is h, the government must be criticized for overreacting, in which case its payoff is $\alpha_o < 1$; if the action is l, the government must be criticized (perhaps mildly) for underreacting, in which case its payoff is $\alpha_u^+ < 1$. In the last cell, the action is disapproved and the outcome is bad, which can happen only if $a = l$. In this case the government should suffer from a more severe criticism for underreacting, and let its payoff be $\alpha_u^- \leq \alpha_u^+$ and we normalize it to 0. In sum, we assume the parameters satisfy $0 = \alpha_u^- \leq \alpha_u^+, \beta \leq 1$ and $\alpha_o < 1$. In particular, $\alpha_o < 0$ is allowed to reflect the possibility that the citizen strongly dislikes overreaction. Note that our baseline model corresponds to the case with $\beta = 1$ and the three α parameters being 0.

The single-society case can be analyzed similarly as before. The government's expected payoff, if it takes action h, is $q_h(\mathbb{I}_{T_h(\mu_0)\geq\hat{\mu}} + \alpha_o \mathbb{I}_{T_h(\mu_0)<\hat{\mu}}) + (1 - q_h)\beta$, and otherwise it is $q_l(\mathbb{I}_{T_l(\mu_0)<\hat{\mu}} + \alpha_u^+ \mathbb{I}_{T_h(\mu_0)\geq\hat{\mu}})$. By a similar argument as in the baseline case, one can check that Proposition 1 (i.e., the government takes action h if and only if $\mu_0 \geq \hat{\mu}_1$) still holds if we replace condition (6) by

$$q_l > q_h \alpha_o + (1 - q_h)\beta \,.$$

This is easier to be satisfied when α_o, the payoff associated with overreaction, is smaller. (If $\alpha_o < 0$, this condition is even satisfied in the polar case with $q_l = 0$ and $q_h = 1$.) Given the cut-off rule in the single-society case, the main economic force in the two-society model remains unchanged as well. For example, when the three α parameters are zero, Proposition 2 still holds if we replace condition (8) by $q_l(q_h + q_l) \leq \beta + (1-\beta)q_h$.

The government's payoff can be generalized in other aspects as well. For instance, the citizen's prior may also directly affect her evaluation of the government's policy, and the so-called "outcome bias" (i.e., a good outcome will be praised while a bad outcome will be criticized regardless of the action) may also play some role. Also, the government may directly care about the citizen's welfare to some extent. However, provided that the evaluation component based on the citizen's posterior is sufficiently important, our main insights should carry over.[21]

B.2 Beyond bad-news information structure

In the baseline model we assume that when the state is mild, the outcome is always good regardless of the government action. Now we relax this assumption and let the outcome under the mild state be stochastic as well. More specifically, suppose $x = 0$ with probability q'_a for $a \in \{l, h\}$ under the mild state. It is natural to assume $q'_a > q_a$ and

$$q_h - q_l > q'_h - q'_l \,. \tag{13}$$

The latter implies the "marginal" effect of taking the high action in containing the crisis is higher when the state is severe.

As in the baseline model, let $\hat{\mu}$ be the threshold in the citizen's evaluation rule. It now solves

$$c = \mu(q_h - q_l) + (1-\mu)(q'_h - q'_l) \tag{14}$$

since the high action can also lower the chance of a bad outcome under the mild state. Under condition (13) the citizen will approve action h if and only if her posterior is greater than $\hat{\mu}$.

[21] Generally, we can define $\tilde{v}_a(x, \mu, \mu_0)$ as the government's payoff when it takes action a, the realized outcome is x, the citizen's posterior is μ, and the citizen's prior is μ_0. Since μ is function of (a, x, μ_0), we can rewrite the payoff function as $v_a(x, \mu_0)$. Let $\bar{v}_a(\mu_0) \equiv q_h v_a(0, \mu_0) + (1-q_h)v_a(1, \mu_0)$ be the expected payoff function associated with action a. Then we have the cut-off result if $\bar{v}_h(\mu_0)$ increases in μ_0, $\bar{v}_l(\mu_0)$ decreases in μ_0, $\bar{v}_h(0) < \bar{v}_l(0)$, and $\bar{v}_h(1) > \bar{v}_l(1)$. At this level of generality, of course it can be complex to specify the primitive conditions for all these conditions to be satisfied.

Let $T_{a,x}(\mu_0)$ be the citizen's posterior of the state after seeing action a and outcome x. When a good outcome is realized, we have

$$T_{h,0}(\mu_0) = \frac{\mu_0 q_h}{\mu_0 q_h + (1-\mu_0)q'_h};\; T_{l,0}(\mu_0) = \frac{\mu_0 q_l}{\mu_0 q_l + (1-\mu_0)q'_l}\,.$$

Both are less than μ_0 since observing a good outcome makes the citizen more optimistic. The opposite is true when a bad outcome is realized, in which case we have

$$T_{h,1}(\mu_0) = \frac{\mu_0(1-q_h)}{\mu_0(1-q_h) + (1-\mu_0)(1-q'_h)};\; T_{l,1}(\mu_0) = \frac{\mu_0(1-q_l)}{\mu_0(1-q_l) + (1-\mu_0)(1-q'_l)}\,.$$

Notice that condition (13) implies $q'_h/q'_l < q_h/q_l$ and so $T_{l,0}(\mu_0) < T_{h,0}(\mu_0)$.[22] Let $\hat{\mu}_{a,x}$ solve $T_{a,x}(\hat{\mu}_{a,x}) = \hat{\mu}$. Then we have

$$\hat{\mu}_{h,1}, \hat{\mu}_{l,1} < \hat{\mu} < \hat{\mu}_{h,0} < \hat{\mu}_{l,0}\,.$$

Notice that $\hat{\mu}_{h,0}$ and $\hat{\mu}_{l,0}$ are the counterparts of $\hat{\mu}_1$ and $\tilde{\mu}_1$ in the baseline model.

Given the government knows the true state is severe, its expected payoff if it takes action h is

$$q_h \mathbb{I}_{T_{h,0}(\mu_0)\geq\hat{\mu}} + (1-q_h)\mathbb{I}_{T_{h,1}(\mu_0)\geq\hat{\mu}}\,,$$

and its expected payoff if it takes action l is

$$q_l \mathbb{I}_{T_{l,0}(\mu_0)<\hat{\mu}} + (1-q_l)\mathbb{I}_{T_{l,1}(\mu_0)<\hat{\mu}}\,.$$

The main difference, compared to the baseline case, is that now action l can also be approved when the outcome is bad, which occurs when the citizen was initially very optimistic. As a result, no action will dominate the other over all possible priors μ_0 as illustrated in Figure 3. (In the baseline model with $q'_h = q'_l = 1$, both $\hat{\mu}_{h,1}$ and $\hat{\mu}_{l,1}$ degenerate at 0. In that case action h dominates in Figure 3b.)

If $q_l > 1 - q_h$ as in the baseline case, the government takes action h if and only if $\mu_0 \geq \hat{\mu}_{h,0}$ as illustrated in Figure 3a; in contrast, if $q_l < 1 - q_h$, the government takes action h if and only if $\mu_0 \geq \max\{\hat{\mu}_{h,1}, \hat{\mu}_{l,1}\}$ as illustrated on Figure 3b. In the former case, overreaction arises under action h more likely than underreaction under action l, so the government takes h less likely than the citizen herself would do according to her prior; in the latter case, however, the opposite is true, so that the government takes h more likely than the citizen herself would do. It is also clear that when the citizen's prior is rather extreme, the government will take the

[22] But condition (13) does not necessarily imply $(1-q'_h)/(1-q'_l) < (1-q_h)/(1-q_l)$, and so the ranking between $T_{h,1}(\mu_0)$ and $T_{l,1}(\mu_0)$ is unclear.

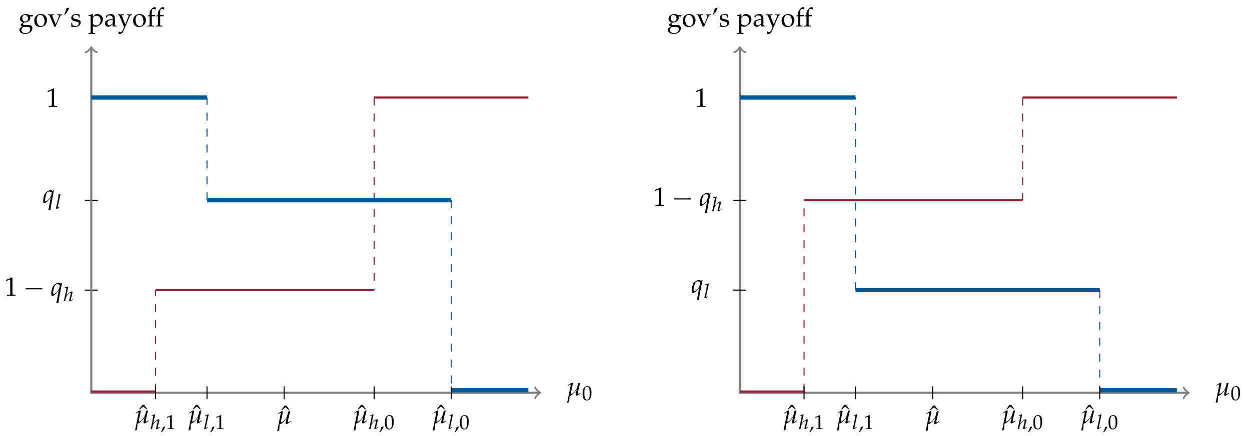

(a) $1-q_h < q_l$

(b) $1-q_h > q_l$

Figure 3: Government's payoff under general information structure: the red thin line corresponds to the government's payoff under action h, while the blue thick line corresponds to the payoff under action l.

action consistent with her prior, in which case it always gets approved. The policy-making dilemma now arises when μ_0 is in the middle range.

Since the single-society case still features a cut-off rule as in the baseline model, the main economic force in the two-society case remains unchanged as well. However, with more societies it is possible for a belief-and-action cycle to arise, which differs from the baseline case. For example, consider the case when a low action in an early society leads to a bad outcome. Since the bad outcome is no longer conclusive about the state, if it leads to a high action and a good outcome in the next society, societies afterward can become optimistic enough to adopt a low action again.

B.3 A reputation model

As we discussed in the main text, a more conventional approach to model political accountability is to introduce a government's private type that is payoff relevant to citizens. In this section, we explore a modelling approach in this vein which can generate similar main results but with somewhat different economics and empirical implications.

There are two societies, where each government can be either *competent* or *incompetent*. The competence type is independent across the two governments. A competent government is a strategic player who chooses an action $a \in \{l, h\}$ to maximize its payoff as specified below given its information on the state, while an incompetent government is a "behavioral" player who mechanically commits to action h. This behavioral-type approach is standard in the reputation literature. See Kreps and Wilson (1982) for classic examples and Mailath and Samuelson (2015) for a comprehensive survey. In our crisis management context, the assumption for the behavioral type can be justified if an incompetent government is unable to efficiently acquire the state information and its enforcement ability is extremely poor. If it takes the low action, a third catastrophic outcome will take place when the state is severe, causing massive damage to both the society and itself. Consequently, provided it believes there is a chance that the true state is severe, an incompetent government always takes the high action.

Each government privately observes its competence type, and if it is competent it also observes a private signal of the true state. For simplicity we assume the signal perfectly reveals the true state, but no government can creditably reveal its information to its citizen. Each citizen's prior is that a government is competent with probability λ_0 and the state is severe with probability μ_0, and they know that their government, if competent, observes a perfect signal.[23]

The other aspects of the model remain the same as before, except for each government's payoff structure. Let λ denote a citizen's posterior belief that her government is competent, i.e., the government's *reputation*. Her government's payoff is then

$$\lambda + \gamma u(a, x)$$

for some constant $\gamma > 0$, where $u(a, x)$ is the citizen's payoff defined in (2) when her government's action is a and the outcome is x.[24] As standard in the political economy literature, the reputation concern can be justified by introducing a post-crisis reelection in each society: the citizen prefers a competent government and chooses between the incumbent government and a challenger whose reputation is uniformly distributed on $[0, 1]$. This specification implies that each government is motivated by both its citizen's welfare and the perks of office.

The strategy of a government specifies a competent government's policy choice in each state, conditional on the action and the outcome in the previous society (if any). Citizens observe actions and outcomes in both societies and form their beliefs about the state and the types of governments. A Perfect Bayesian Equilibrium consists of governments' strategies

[23]The feature that the citizen is uncertain about both an underlying state and the policy maker's type is similar to, for example, Coate and Morris (1995) and Maskin and Tirole (2004).

[24]Notice that if $\gamma = 0$ and each government only cares about its reputation, then action l becomes their dominant strategy as it perfectly signals competence.

and citizens' beliefs that satisfy the following properties. First, citizens' beliefs are consistent with governments' strategies in the sense that they are generated by Bayesian updating wherever possible. Second, each government's strategy is optimal given citizens' beliefs.

Let us first consider the single-society case. When the state is mild, a competent government will choose action l. This is because the low action is a perfect signal of competence, and it is also the best policy for the citizen given the outcome is always good under a mild state. What needs to be pinned down is a competent government's strategy when the state is severe. Let $\sigma \in [0,1]$ be the probability that it chooses action h. The government's trade-off is between its desire to separate itself from an incompetent type (which favors action l) and the citizen's welfare (which favors action h).

Let $\lambda^{\sigma}_{a,x}$ denote the citizen's posterior of the government's type after seeing action a and outcome x given the competent government's policy strategy σ when the state is severe. When $a = l$, the posterior $\lambda^{\sigma}_{l,x}$ is always 1. When $a = h$ and $x = 1$, the citizen learns that the state must be severe, and so

$$\lambda^{\sigma}_{h,1} = \frac{\lambda_0 \sigma (1-q_h)}{\lambda_0 \sigma (1-q_h) + (1-\lambda_0)(1-q_h)} = \frac{1}{1 + \frac{1-\lambda_0}{\lambda_0}\frac{1}{\sigma}}; \tag{15}$$

when $a = h$ and $x = 0$, we have

$$\lambda^{\sigma}_{h,0} = \frac{\lambda_0 \mu_0 \sigma q_h}{\lambda_0 \mu_0 \sigma q_h + (1-\lambda_0)(\mu_0 q_h + 1 - \mu_0)} = \frac{1}{1 + \frac{1-\lambda_0}{\lambda_0}\frac{1}{\sigma}\left(1 + \frac{1-\mu_0}{\mu_0}\frac{1}{q_h}\right)}. \tag{16}$$

(We stipulate $1/0 = \infty$ so that $\sigma = 0$ is permitted.) Note that the government will take action h only if it is competent and the state is severe or if it is incompetent.

The following two observations are important for both our subsequent analysis and the key insights in this reputation model: First, we have $\lambda^{\sigma}_{h,0} \le \lambda^{\sigma}_{h,1} \le \lambda_0$. Given an incompetent government always takes the high action, h is a signal of incompetence, and that is why both posteriors become smaller than λ_0. Meanwhile, when $x = 1$, the citizen learns the state is severe, in which case h is less a signal of incompetence given the competent government is more likely to take h in the severe state than in the mild state. Second, both posteriors are increasing in λ_0, μ_0 and σ. In particular, when the citizen believes the state is more likely to be severe or when she believes the competent government takes action h more often in the severe state, she regards h less as a signal of incompetence. When $x = 1$ and $\sigma = 1$, we have $\lambda^{\sigma}_{h,1} = \lambda_0$, i.e., the high action causes no reputation damage.

If the government takes action h, its expected payoff is $q_h \lambda^{\sigma}_{h,0} + (1-q_h)\lambda^{\sigma}_{h,1} - \gamma(1 - q_h + c)$. In this case it bears the reputation cost and also imposes a cost c on the citizen, but the citizen is less likely to suffer from a bad outcome. If the government chooses action l, its expected payoff is $1 - \gamma(1 - q_l)$. In this case it bears no reputation cost, but the citizen is more likely to suffer from a bad outcome. (From the reputation perspective, there is no the feature of "damned if you do, damned if you don't" in this model, but it remains from the perspective of the government's payoff.) The first payoff is higher if and only if

$$q_h \lambda_{h,0}^{\sigma} + (1 - q_h)\lambda_{h,1}^{\sigma} \geq 1 - \gamma(q_h - q_l - c). \tag{17}$$

Notice that the left-hand side strictly increases in σ. With this observation we can characterize a unique (stable) equilibrium in the single-society case as reported in Proposition 4 below.

In the two-society case, a competent government 2 will act similarly as in the single-society case except that its citizen has an updated interim belief on the state after seeing what has happened in society 1. Then by backward induction we can similarly analyze a competent government 1's decision.

Proposition 4. *In the single-society case, there is a unique (stable) equilibrium in which a competent government takes the high action in the severe state if and only if (λ_0, μ_0) satisfies (17) at $\sigma = 1$; in the two-society case, a similar result holds for a competent government 1 but for a larger set of (λ_0, μ_0).*

Proof. Single society. Note that the left-hand side of (17) strictly increases in σ. If the opposite of (17) holds at $\sigma = 1$, the competent government always takes the low action in the severe state. Then we must have $\sigma = 0$ in equilibrium. In contrast, if (17) holds at $\sigma = 1$, it is an equilibrium that the competent government always takes the high action in the severe state, i.e., $\sigma = 1$. If the right-hand side of (17) is positive, there is also another equilibrium where the competent government plays a mixed strategy with $\sigma \in (0, 1)$ which solves the equality of (17). (Such an interior solution of σ always exists in this case since the left-hand side of (17) equals zero at $\sigma = 0$.) However, this equilibrium is unstable in the sense that if the citizen expects a slightly different σ, the competent government will take either the high or the low action for sure.

Two societies. Suppose both citizens expect a competent government 1 to take action h with probability $\sigma \in [0, 1]$ in the severe state. If the competent government 1 takes l in the severe state, its type is revealed perfectly. Then its expected payoff is independent of society 2 and is exactly the same as in the single-society case, i.e., $1 - \gamma(1 - q_l)$.

If the competent government 1 takes h in the severe state and if $x_1 = 1$, then its payoff is also independent of society 2 since the bad outcome already reveals the true severe state. In this case, its reputation is $\lambda^{\sigma}_{h,1}$ as defined in (15). If $x_1 = 0$, however, government 1's expected payoff will depend on government 2's policy and its outcome. An incompetent government 2 will always take action h; a competent government 2 will take a deterministic action a_2 as already shown in the single-society case. If $x_2 = 1$, the severe state is revealed, and then government 1's reputation is $\lambda^{\sigma}_{h,1}$; if $x_2 = 0$, let $\lambda^{\sigma}_{h,0;a_2,0}$ be government 1's reputation, which will be specified later. Therefore, government 1's expected payoff is $q_h[\lambda_0 \Lambda^{\sigma}_{a_2} + (1 - \lambda_0)\Lambda^{\sigma}_{h}] + (1 - q_h)\lambda^{\sigma}_{h,1} - \gamma(1 - q_h + c)$, where

$$\Lambda^{\sigma}_{a} \equiv q_a \lambda^{\sigma}_{h,0;a,0} + (1 - q_a)\lambda^{\sigma}_{h,1}$$

is government 1's expected reputation when government 2 takes action a conditional on $a_1 = h$ and $x = 0$. Therefore, a competent government 1 prefers h if and only if

$$q_h \left[\lambda_0 \Lambda^{\sigma}_{a_2} + (1 - \lambda_0)\Lambda^{\sigma}_{h}\right] + (1 - q_h)\lambda^{\sigma}_{h,1} > 1 - \gamma(q_h - q_l - c) \,. \tag{18}$$

Compared to condition (17) in the single-society case, the difference is the square-bracket term (which was simply $\lambda^{\sigma}_{h,0}$ in the single-society case), and it reflects how the presence of society 2 affects government 1's payoff.

Notice that given $a_1 = h$ and $x_1 = 0$, a competent government 2 will act as in the single-society case with primitives $(\lambda_0, T_h(\mu_0))$, where $T_h(\mu_0)$ is citizen 2's posterior of the state given she believes that government 1's strategy is σ.[25] If $(\lambda_0, T_h(\mu_0))$ is in the light blue area in Figure 4a, a competent government 2 will choose $a_2 = l$ for sure. In this case, if $x_2 = 0$, citizen 1's posterior of government 1's type is

$$\lambda^{\sigma}_{h,0;l,0} = \frac{\lambda_0 \mu_0 \sigma q_h \lambda_0 q_l}{\lambda_0 \mu_0 \sigma q_h \lambda_0 q_l + (1 - \lambda_0)[\mu_0 q_h \lambda_0 q_l + (1 - \mu_0)\lambda_0]} = \frac{1}{1 + \frac{1-\lambda_0}{\lambda_0}\frac{1}{\sigma}\left(1 + \frac{1-\mu_0}{\mu_0}\frac{1}{q_h q_l}\right)} \,.$$

If $(\lambda_0, T_h(\mu_0))$ is in the dark-blue area in Figure 4a, a competent government 2 will choose $a_2 = h$ for sure. In this case, if $x_2 = 0$, citizen 1's posterior of government 1's type is

$$\lambda^{\sigma}_{h,0;h,0} = \frac{\lambda_0 \mu_0 \sigma q_h q_h}{\lambda_0 \mu_0 \sigma q_h q_h + (1 - \lambda_0)[\mu_0 q_h q_h + (1 - \mu_0)(1 - \lambda_0)]} = \frac{1}{1 + \frac{1-\lambda_0}{\lambda_0}\frac{1}{\sigma}\left(1 + \frac{1-\mu_0}{\mu_0}\frac{1-\lambda_0}{q_h^2}\right)} \,.$$

(It is ready to see that $\lambda^{\sigma}_{h,0;l,0} < \lambda^{\sigma}_{h,0;h,0}$ given $\frac{1}{q_l} > \frac{1-\lambda_0}{q_h}$ and both are less than λ_0 as expected.) When government 2 is incompetent and its high action leads to $x_2 = 0$, citizen 1's posterior of its government's type is also $\lambda^{\sigma}_{h,0;h,0}$.

[25] More precisely,

$$T_h(\mu_0) = \frac{\mu_0 \sigma q_h}{\mu_0 \sigma q_h + (1 - \mu_0)(1 - \lambda_0)} = \frac{1}{1 + \frac{1-\mu_0}{\mu_0}\frac{1-\lambda_0}{\sigma q_h}} \,.$$

Given the left-hand side of (18) strictly increases in σ, the same argument as in the single-society case implies that there is a unique (stable) equilibrium where a competent government 1 takes h in the severe state if and only if (18) holds at $\sigma = 1$.

We now show that $\lambda_0 \Lambda^\sigma_{a_2} + (1-\lambda_0)\Lambda^\sigma_h > \lambda^\sigma_{h,0}$, so that (18) holds at $\sigma = 1$ for a larger set of (λ_0, μ_0) than that for (17) at $\sigma = 1$. This is true if

$$q_{a_2}\lambda^\sigma_{h,0;a_2,0} + (1-q_{a_2})\lambda^\sigma_{h,1} > \lambda^\sigma_{h,0}$$

for both $a_2 = l$ and h. This is immediate once one notices that $q_l/\lambda^\sigma_{h,0;l,0} + (1-q_l)/\lambda^\sigma_{h,1} = 1/\lambda^\sigma_{h,0}$ and $q_h/\lambda^\sigma_{h,0;h,0} + (1-q_h)/\lambda^\sigma_{h,1} < 1/\lambda^\sigma_{h,0}$ and then apply the Jessen's inequality.[26] □

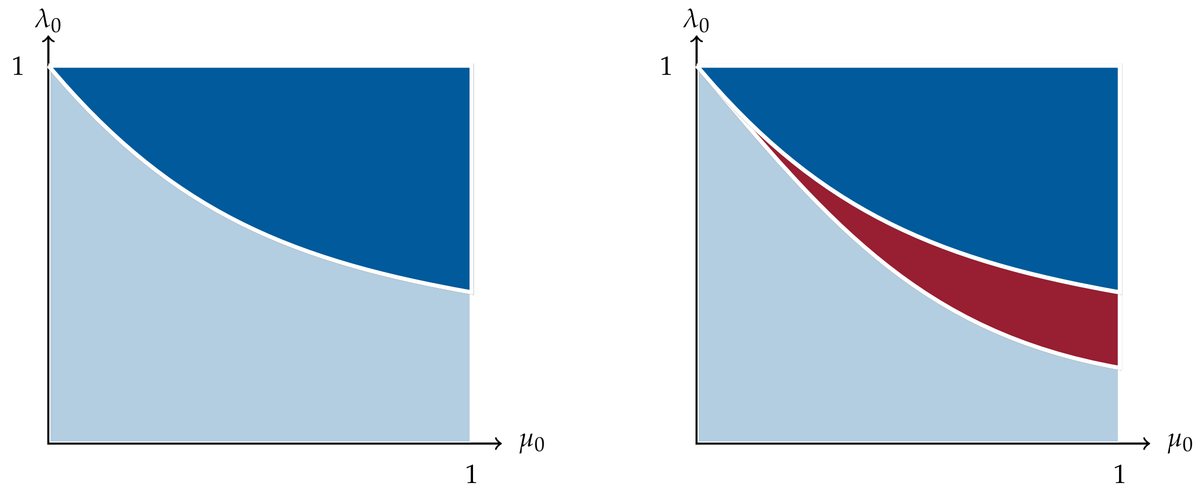

(a) Government's decision in the single-society case (b) Government 1's decision in the two-society case

Figure 4: Illustration of a competent government's decision in the reputation model: The dark-blue area indicates the prior pairs (λ_0, μ_0) under which h is the government's optimal action, while the light-blue area corresponds to the prior pairs under which l is the optimal action. The red area corresponds to the extra prior pairs under which h becomes the optimal action for government 1 in the two-society case.

Figure 4a illustrates a competent government's optimal policy in the severe state in the single-society case: it takes the first-best high action if and only if both λ_0 and μ_0 are sufficiently high (so that the reputation damage caused by the high action is small enough), where the boundary is determined by the equality of (17) at $\sigma = 1$. (The boundary is decreasing because both $\lambda^\sigma_{h,0}$ and $\lambda^\sigma_{h,1}$ are increasing in λ_0 and μ_0.) In particular, for a given λ_0 which is not too small, the competent government takes the high action if and only if μ_0 is

[26] This is actually a consequence of a more general martingale property in our setup: given $\lambda^\sigma_{h,1} = \lambda^\sigma_{h,0;a_2,1}$, we have $\mathbb{E}[\lambda^\sigma_{h,0;a_2,x_2}|\text{severe state}] > \mathbb{E}[\lambda^\sigma_{h,0;a_2,x_2}] = \lambda^\sigma_{h,0}$.

greater than some threshold as in our baseline model. The underlying economics, though, is different: in our baseline model, what discourages the government from taking the high action in the severe state is the prospect of being criticized for overreacting when a good outcome is realized; here in this reputation model, it is because taking a high action is regarded as a signal of incompetence.

Figure 4b illustrates a competent government 1's optimal policy in the severe state in the two-society case: the presence of the second society expands the range of λ_0 and μ_0 in which it takes the first-best high action. The intuition is as follows: Each government's reputation is now influenced by the action and outcome in both societies. A good outcome in society 1 leads to a lower updated μ_0 in society 2, which tends to induce a low action and so likely a bad outcome there. A bad outcome in society 2 helps reveal the severe state. Once citizen 1 is eventually convinced that the state is severe, she will regard its government's high action less as a signal of incompetence. This mitigates the reputation concern and encourages a competent government 1 to adopt the high action in the severe state. This result is qualitatively similar to what we saw in our baseline model.[27]

[27] A subtle difference here is that the negative sampling effect (which arises when a good outcome is also realized in society 2) is always dominated jointly by the positive sampling and the strategic effect. This is due to the payoff-structure difference for the government.

References

BANERJEE, A. V. (1992): "A Simple Model of Herd Behavior," *Quarterly Journal of Economics*, 107(3), 797–817.

BARRO, R. (1973): "The Control of Politicians: An Economic Model," *Public Choice*, 14(1), 19–42.

BIKHCHANDANI, S., D. HIRSHLEIFER, AND I. WELCH (1992): "A Theory of Fads, Fashion, Custom, and Cultural Change as Informational Cascades," *Journal of Political Economy*, 100(5), 992–1026.

BOLTON, P., AND C. HARRIS (1999): "Strategic Experimentation," *Econometrica*, 67(2), 349–374.

BRANDENBURGER, A., AND B. POLAK (1996): "When Managers Cover Their Posteriors: Making the Decisions the Market Wants to See," *RAND Journal of Economics*, 27(3), 523–541.

CALLANDER, S., AND P. HUMMEL (2014): "Preemptive Policy Experimentation," *Econometrica*, 82(4), 1509–1528.

CAMERER, C., G. LOEWENSTEIN, AND M. WEBER (1989): "The Curse of Knowledge in Economic Settings: An Experimental Analysis," *Journal of Political Economy*, 97(5), 1232–1254.

CANES-WRONE, B., M. C. HERRON, AND K. W. SHOTTS (2001): "Leadership and Pandering: A Theory of Executive Policymaking," *American Journal of Political Science*, 45(3), 532–550.

CHE, Y.-K., AND N. KARTIK (2009): "Opinions as Incentives," *Journal of Political Economy*, 117(5), 815–860.

COATE, S., AND S. MORRIS (1995): "On the Form of Transfers to Special Interests," *Journal of Political Economy*, 103(6), 1210–1235.

ELY, J. C., AND J. VÄLIMÄKI (2003): "Bad Reputation," *Quarterly Journal of Economics*, 118(3), 785–814.

FISCHHOFF, B. (1975): "Hindsight Is Not Equal to Foresight: The Effect of Outcome Knowledge on Judgment under Uncertainty," *Journal of Experimental Psychology: Human Perception and Performance*, 1(3), 288–299.

FOX, J., AND R. VAN WEELDEN (2015): "Hoping for the Best, Unprepared for the Worst," *Journal of Public Economics*, 130, 59–65.

GENTZKOW, M., AND J. M. SHAPIRO (2006): "Media Bias and Reputation," *Journal of Political Economy*, 114(2), 280–316.

HARRINGTON, J. E. (1993): "Economic Policy, Economic Performance, and Elections," *American Economic Review*, 83(1), 27–42.

HEALY, A., AND N. MALHOTRA (2009): "Myopic Voters and Natural Disaster Policy," *American Political Science Review*, pp. 387–406.

——— (2013): "Retrospective Voting Reconsidered," *Annual Review of Political Science*, 16, 285–306.

HIRSCH, A. V. (2016): "Experimentation and Persuasion in Political Organizations," *American Political Science Review*, 110(01), 68–84.

KEY, V. (1966): *The Responsible Electorate: Rationality in Presidential Voting, 1936-1960*. Harvard University Press.

KREPS, D. M., AND R. WILSON (1982): "Reputation and Imperfect Information," *Journal of Economic Theory*, 27(2), 253–279.

MAILATH, G. J., AND L. SAMUELSON (2015): "Reputations in Repeated Games," in *Handbook of Game Theory with Economic Applications*, vol. 4, pp. 165–238. Elsevier.

MASKIN, E., AND J. TIROLE (2004): "The Politician and the Judge: Accountability in Government," *American Economic Review*, 94(4), 1034–1054.

MORRIS, S. (1995): "The Common Prior Assumption in Economic Theory," *Economics and philosophy*, 11, 227–253.

——— (2001): "Political Correctness," *Journal of Political Economy*, 109(2), 231–265.

PRAT, A. (2005): "The Wrong Kind of Transparency," *American economic review*, 95(3), 862–877.

SKIDMORE, M. J. (2016): *Presidents, Pandemics, and Politics*. Springer.

STASAVAGE, D. (2020): "Democracy, Autocracy, and Emergency Threats: Lessons for COVID-19 From the Last Thousand Years," *International Organization*, pp. 1–17.