

Yale University

EliScholar – A Digital Platform for Scholarly Publishing at Yale

Cowles Foundation Discussion Papers

Cowles Foundation

12-2020

A Model of Sequential Crisis Management

Fei Li

Jidong Zhou

Follow this and additional works at: <https://elischolar.library.yale.edu/cowles-discussion-paper-series>



Part of the Economics Commons


A MODEL OF SEQUENTIAL CRISIS MANAGEMENT

By

Fei Li and Jidong Zhou

April 2022

COWLES FOUNDATION DISCUSSION PAPER NO. 2351



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

<http://cowles.yale.edu/>

A Model of Sequential Crisis Management*

Fei Li[†] Jidong Zhou[‡]

April, 2022

Abstract

We propose a model of how multiple societies respond to a common crisis. A government faces a “damned-either-way” policymaking dilemma: aggressive intervention contains the crisis, but the resulting good outcome makes people skeptical about the costly response; light intervention worsens the crisis and causes the government to be faulted for not doing enough. When multiple societies encounter the crisis sequentially, due to this policymaking dilemma, late societies may underperform despite having more information, while early societies can 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, Robert Axtell, Heski Bar-Isaac, Gary Biglaiser, Florian Ederer, Jonathan Feinstein, Yunzhi Hu, Ilwoo Hwang, Gilat Levy, Jin Li, Navin Kartik, Leslie Marx, Peter Norman, Ben 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

Crises may not end where they begin. The great recession of 2008, emerging in Wall Street, made its way to Asia and Europe. The Covid-19 virus was first reported in Wuhan and then spread across the world. When a crisis spills from one area to another, folk wisdom suggests that late-hit areas can learn from early-hit areas' experiences and respond better. This view, however, overlooks a salient *policymaking dilemma* in crisis management. The public often do not know early in a crisis whether it is a severe threat. If an aggressive action is taken and the crisis is contained, people may underestimate the severity of the problem and become more skeptical about the need for the costly action; if a less aggressive action is taken and the situation gets out of control, people may then blame the policymaker for not having taken the necessary precautions. This leaves the policymaker in a "damned if you do, damned if you don't" dilemma.

If policymakers are deeply concerned about the political ramification, their crisis response can be inadequate and sluggish, especially when the public are initially optimistic. One example is the political fallout following the early response to the subprime crisis. [Bernanke, Geithner, and Paulson \(2019\)](#) document that "the Bear (Sterns) intervention did calm the markets, but it also provided a searing dose of the politics of crisis response. Many politicians and pundit accused us of overreacting to the Darwinian rhythms of capitalism, arguing that the economic impact of an investment bank's failure would be modest — an argument that would be disproved by Lehman's failure six months later." Another example is Gerald Ford's national-wide vaccination campaign to prevent a new variant of influenza in 1976. According to [Skidmore \(2016\)](#) and [Lewis \(2021\)](#), Ford was widely criticized and ridiculed after the influenza turned out to be less deadly than expected.¹

This "damned-either-way" dilemma has a profound policy implication if the crisis spreads across multiple societies sequentially. When a society that encounters the crisis first takes the precaution and the crisis is contained, the public in subsequent societies, after seeing the first society's outcome, may become more optimistic. This boosted optimism can prevent the policymakers there from adopting the right policy if they are more concerned about overreacting. The resulting adverse outcome can in turn provide a policy counterfactual and justify the first policymaker's initially unpopular policy choice, relieving them from the

¹Also, [Healy and Malhotra \(2009\)](#) provide evidence that voters reward the incumbent presidential party for delivering disaster relief spending, but not for investing in disaster preparedness spending. They discuss several possible reasons, including the one that the counterfactual of relief spending is easier to see than the counterfactual of preparedness spending.

policymaking dilemma.

Due to this *dynamic counterfactual effect*, the early-hit society can have a first-mover advantage. That is, all else equal, societies that encounter the crisis later may handle it worse than the early-hit society. Later movers can learn from earlier movers' experiences, but having more information is not necessarily a blessing: people may become too optimistic after seeing good outcomes in early-hit societies. Conversely, the early-hit society, foreseeing the possible counterfactual from subsequent societies, is more willing to adopt a costly but more effective policy.

We develop a parsimonious model to formalize this idea. The policymaking 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; and the policymaker is held accountable for their policy choice after its consequence is observed. Due to the policymaking dilemma, the policymaker may pander 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.

One application of our model is the initial response to the Covid-19 pandemic. Every government faced the policymaking dilemma due to the public's initial uncertainty about the severity of the crisis. The virus spread gradually across regions, making our model applicable. We provide a new angle to understand the witnessed response disparity, complementing a spectrum of possible explanations from cultural differences to institutional heterogeneity. In particular, our model predicts: (i) early-hit countries are less hesitant to adopt aggressive and precautionary measures; (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 by their people, but later lauded after the adverse consequences under alternative responses are observed. These implications are roughly consistent with what happened from January to March in 2020.² Early-hit countries such as China, Singapore, South Korea, and Vietnam adopted strict measures from large-scale lockdown to extensive contact tracing. They quickly flattened the curve, and some even contained the domestic spread. However, many of these policies were widely criticized as clumsy, un-

²From April 2020, the severity of the crisis has been largely recognized worldwide. Once the government and the public reach consensus, the government can take strict policies without facing a political hurdle.

necessary, or even counterproductive,³ and the temporary victory created a false sense of safety that the virus is a regional problem that can be easily tamed.⁴ These misconceptions, despite of being rebutted by reputable epidemiologists, made some rather less costly prevention measures such as wearing masks and social distancing first controversial and then politicized in several countries. It is therefore not surprising that some politicians, despite of being well informed, chose to pander to the opinions of their supporters instead of being more responsive. The inaction soon led to disastrous outcomes, and this was used as a counterfactual to justify the draconian measures adopted by some early-hit countries.

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.⁵

To deliver the main idea transparently, we adopt two perhaps unconventional assumptions in the main model. First, we assume that the public and the policymaker agree to disagree on the severity of the crisis (perhaps even after communications). This can be due to the public's mistrust of the government. As a result, the public do not attempt to infer information on the crisis severity from the government's chosen policy. This isolates our mechanism from the familiar signaling channel. Second, we assume that the public evaluate the policy using their updated belief of the state after seeing the policy outcome. This is different from the usual "reputation" approach where the public evaluate the policymaker based on underlying characteristics such as their competence or preferences which can be learned from policy choices. These assumptions greatly simplify our model without losing the key insights. We discuss them in more detail in section 2 after we introduce the model.

Related literature. In the single-society case, our model predicts that the policymaker may pander to the public opinion. This pandering effect is certainly not new. It arises, e.g., when incumbent politicians choose policies to signal their competency (Harrington 1993, Canes-Wrone, Herron, and Shotts 2001, and Prat 2005) or preferences (Maskin and Tirole

³See, e.g., <https://reut.rs/3cQRWxx>, <https://nyti.ms/3CRg0uN>, and chapter 2 of Tooze (2021) for more details about the political ramification after Wuhan lockdown.

⁴See, e.g., <https://nyti.ms/3HZUGHh>.

⁵Despite 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>.

2004), when a firm manager, who has a share price concern, makes decisions the market wants to see (Brandenburger and Polak 1996), or when a media slants its report toward readers' prior in order to build a reputation for quality (Gentzkow and Shapiro 2006). However, the trade-offs in those papers differ from ours. They either assume that the outcome is unobservable, or assume that an action is more likely to be approved when it "matches" the state. 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 of 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), 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.

In our model, the policymaker and the public gradually settle their differences on crisis response. 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 resolve opinion difference and implement the optimal policy more efficiently in the second period.⁶ This idea could also be relevant in our crisis management context: 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.

We assume that people look to what happens in other societies to evaluate policies at home. In this sense, our paper is related to yardstick competition in political economy. In that literature, most papers conclude that competition helps alleviate the agency problem and induces all governments to adopt more benevolent actions (see, e.g., Besley and Case 1995). The economic force in our model is different. We focus on the implications of the dynamic counterfactual effect and predict that comparison across jurisdictions leads to a

⁶There are also other works that study the interaction between learning and prior disagreement. For instance, Che and Kartik (2009) argue that with conflicting beliefs, individuals have stronger incentives to acquire information to persuade each other. This can render hiring people with different opinions optimal in organization design.

policy divergence instead of a convergence.⁷

In terms of dynamic information spillover across players, our paper is also related to the substantial literature on social learning (see, e.g., [Bikhchandani, Hirshleifer, Tamuz, and Welch 2021](#) for a recent survey). Our paper differs from the canonical social-learning models such as [Banerjee \(1992\)](#) and [Bikhchandani, Hirshleifer, and Welch \(1992\)](#), since in our model early players have strategic incentives to influence later players' choices. Moreover, we emphasize the possibility that political accountability and information externality can induce policymakers to sequentially make *different* choices. There are also previous works on social learning with strategic considerations across players (e.g., [Scharfstein and Stein 1990](#), [Ottaviani and Sørensen 2001](#), and [Ali and Kartik 2012](#)). However, these works assume different information structures and emphasize information externality as a potential source of herding.

The strategic dynamic interaction among policymakers in our model is also 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. In a political economy context, [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 the robustness of the main insight when we relax the aforementioned unconventional assumptions. Section 5 concludes. Some omitted proofs and other extensions are relegated to Appendix A.

⁷[Cai and Treisman \(2005\)](#) also discover that competition may cause policy divergences when agents' endowments are heterogeneous. In our model, policymakers are otherwise identical except that they counter the crisis sequentially.

2 Single Society

A society faces a potential crisis which can be *severe* or *mild*. There are two players: a policymaker or government and a (representative) citizen. We assume that the government has learned the true state, which is severe, but it cannot fully convince the citizen of the crisis severity. The citizen believes that the crisis is severe with probability $\mu_0 \in (0, 1)$ and is mild with probability $1 - \mu_0$ (even after a communication between the government and the citizen). This prior difference $1 - \mu_0$ can be regarded as the citizen's mistrust of the government. We assume that the public and the government agree to disagree on the crisis severity, and so the public do not attempt to infer information on the crisis severity from the government's policy choice.

The government has two options to handle the crisis: $a \in \{l, h\}$, where action l stands for a light intervention, and action h stands for a heavy intervention.⁸ 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, the outcome is good regardless of the government's action. When the state is severe, however, the outcome depends on the government's action: $x = 0$ with probability q_a for $a \in \{l, h\}$, where $0 \leq q_l < q_h \leq 1$, i.e., action h generates a good outcome more likely.⁹ To highlight the consequence of the government's action, we do not model the impact of the citizen's potential action on the outcome.¹⁰

The citizen observes the government action and the outcome, and then updates her belief about the state by Bayes' rule. Once a bad outcome occurs, the citizen will be convinced that the state is severe.¹¹ A good outcome, however, will make the citizen more optimistic about the state. More precisely, after seeing action $a \in \{h, l\}$ and $x = 0$, the citizen's posterior

⁸For 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.

⁹See Acharya, Grillo, and Pei (2021) for a similar information specification in a repeated tax-setting game with random state.

¹⁰In some examples (e.g., a pandemic), the citizen's effort also matters for containing the crisis. A more optimistic citizen tends to 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.

¹¹As we will discuss in Appendix A.2, this assumption of bad-news information structure is not crucial for the main results of this paper.

belief is that the state is severe with probability

$$T_a(\mu_0) = \frac{\mu_0 q_a}{\mu_0 q_a + 1 - \mu_0}. \quad (1)$$

It is clear that

$$T_l(\mu_0) < T_h(\mu_0) \leq \mu_0. \quad (2)$$

That is, seeing a good outcome realized when $a = l$ makes the citizen more optimistic than when $a = h$. Also, T_a increases in q_a and approaches the prior μ_0 as q_a goes to 1. In other words, a good outcome becomes less informative when an action becomes more effective in containing the crisis.

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

$$u(a, x) = -c \times \mathbb{I}_{a=h} - x, \quad (3)$$

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 is bad, the citizen's disutility is 1. 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 optimal if and only if 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: she approves the government's action if and only if it maximizes her expected utility based on her updated belief which is denoted by μ . As a consequence, action h will be approved if $\mu \geq \hat{\mu}$, and otherwise action l will be approved.¹² The government cares only about how the citizen

¹²We assume that the citizen evaluates the government's policy based on its outcome but without accounting for unobserved counterfactuals. This modelling component is related to [Esponda and Pouzo \(2017\)](#) who propose a new voting equilibrium concept which captures a similar feature. In Section 3 with multiple societies,

evaluates its action (say, for the policymaker’s political career such as reelection or personal legacy purpose).¹³ Its payoff is 1 if its action is approved by the citizen and is 0 otherwise. The government’s objective is to maximize the probability that its action gets approved.¹⁴

Before proceeding, we discuss two main assumptions that simplify our model.

(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.¹⁵ 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. 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 report an extension with the signaling channel and a more general prior and show that it delivers similar insights.

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

One interpretation is that the citizen is subject to the well-known *hindsight bias*: people tend to incorporate the newly available information into their evaluation of a decision (possibly made by other people), 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 it prevails in both individual decision problems and principal-agent relationships. For instance, [Camerer, Loewenstein, and Weber \(1989\)](#) (p. 1246) comment that the hindsight bias in a principal-agent relationship is “especially acute in public decision making, in which the principals are a diffuse group of voters and contracts are rarely explicit.” In our context where people mistrust the government, this bias can be even more relevant.

however, citizens take into account the information learned from the counterfactuals *observed* in other societies.

¹³This specification of government preferences applies not only in democracies but also in autocracies where winning public support is critical for the government to legitimate and stabilize its governance.

¹⁴We will discuss a more general government payoff specification in Appendix A.3. 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.

¹⁵See, for example, [Morris \(1995\)](#) for a comprehensive discussion on the heterogeneous-prior assumption.

Another interpretation is that the policy outcome is an informative signal of the policy effectiveness which is realized in the beginning phase of the crisis. If the citizen believes that her opinion of the policy will influence the government’s decision of whether to continue or terminate the current policy, it is then rational for her to support or protest the current policy based on the updated information.

Our modelling of political accountability differs from the conventional approach which assumes that the policymaker has some private information of their characteristics such as their preferences or competence. The citizen learns information on their characteristics from their policy choice and if possible also from the consequence, and then decides whether or not to reelect them.¹⁶ In Section 4, we report an alternative reputation model in this vein, and show that it delivers similar insights. The main advantage of the current approach is its simplicity, and we also believe that the assumed way of how citizens evaluate government policies (mainly relying on hard information instead of strategic reasoning) and its implications are plausible in many political economy settings.

2.1 Analysis

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

$$q_h \mathbb{I}_{T_h(\mu_0) \geq \hat{\mu}} + 1 - q_h . \quad (4)$$

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 revealed. This payoff is depicted as the thin red curve in Figure 1, where $\hat{\mu}_1$ solves

$$T_h(\hat{\mu}_1) = \hat{\mu} , \quad (5)$$

i.e., $\hat{\mu}_1$ 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. When the prior is below $\hat{\mu}_1$, the government’s high action will be criticized for overreacting when a good outcome is

¹⁶This 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. (See also [Fiorina 1981](#).) Our modelling approach is closer to the latter in spirit.




Figure 1: The thin red curve is the government’s expected payoff under action h ; while the thick blue curve is its expected payoff under action l .

realized.

If the government takes action l , its expected payoff is

$$q_l \mathbb{I}_{T_l(\mu_0) < \hat{\mu}}. \quad (6)$$

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. This payoff is depicted as the thick blue curve in Figure 1, where $\tilde{\mu}_1$ solves $T_l(\tilde{\mu}_1) = \hat{\mu}$. (Note that property (2) implies $\hat{\mu} \leq \hat{\mu}_1 < \tilde{\mu}_1$.) When the prior is below $\tilde{\mu}_1$, the government’s low action will be criticized for underreacting when a bad outcome is realized.

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

In the rest of the paper, we maintain the following assumption and focus on the case depicted in Figure 1:

Assumption 1. *In the single-society case, the risk that the government is accused of overreacting*

under action h is greater than the risk of underreacting under action l , i.e.,

$$q_h > 1 - q_l. \quad (7)$$

When Assumption 1 fails to hold, we have the uninteresting case where 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.

Then as indicated in Figure 1, the government will take the inefficient 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.¹⁷ That is, we have the following result:

Proposition 1. *The government takes action h if and only if $\mu_0 \geq \hat{\mu}_1$, where $\hat{\mu}_1$ is defined in (5).*

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.¹⁸ Given $T_h(\cdot)$ increases in q_h , it is easy to see from the definition of $\hat{\mu}_1$ that, similar as $\hat{\mu}$, the cutoff $\hat{\mu}_1$ also decreases in q_h and increases in q_l and c . Therefore, when the high action is more effective in containing the crisis (i.e., when $q_h - q_l$ is greater) or when the high action imposes less cost on the citizen (i.e., when c is smaller), the range of μ_0 widens 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 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.¹⁹ 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

¹⁷In the edge case with $q_h = 1$ and $q_l = 0$, the equality of (7) 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 .

¹⁸A 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).

¹⁹See, for example, Francis Fukuyama, "The thing that determines a country's resistance to the Coronavirus," *The Atlantic*, March 2020.

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 masks) may lead to a small q_h .²⁰ 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.) Each society has two players: a government and a citizen. Each government knows the crisis is severe, while the citizens initially believe that the crisis is severe with probability μ_0 . They agree to disagree as in the single-society case. As before, each government has two options, $a_i \in \{l, h\}$. The outcome in each society depends only on the state and the government's action in that society. We have assumed that citizens hold a common prior. The case with heterogeneous priors can be analyzed similarly, but does not add useful new insights.

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.

²⁰See, e.g., <https://nyti.ms/2YOMNiS> on "More Americans should probably wear masks for protection" in *The New York Times*.

Finally, citizens 1 and 2, after observing actions and outcomes in *both* societies, evaluate their own government. An implicit assumption we have made is 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 intervention or a good outcome in the first society does not completely halt the spread of the crisis.

Since a bad outcome in any society perfectly reveals the true state, we only need 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 defined in (1) 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 (2), we have

$$T_l^{[2]} < T_h \circ T_l = T_l \circ T_h < T_h^{[2]}, \quad (8)$$

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.

3.1 Analysis

Let $\hat{\mu}_2$ solve

$$T_h(\hat{\mu}_2) = \hat{\mu}_1, \quad (9)$$

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 an action h and a good outcome in both societies. Clearly we have that $\hat{\mu} \leq \hat{\mu}_1 \leq \hat{\mu}_2$. A condition we will often refer to in the subsequent analysis is

$$q_l(q_h + q_l) \leq 1. \quad (10)$$

For any given q_h , this condition holds if q_l is sufficiently small (i.e., if a low action is very unlikely to contain the crisis). This condition is essential for the counterfactual effect to work: when q_l is smaller, if government 2 takes action l due to the boosted optimism among its citizens after seeing a good outcome in the first society, the outcome is more likely to be bad, in which case government 1's high action will be justified.

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

Proposition 2. *When there are two societies,*

- (i) *if $\mu_0 \geq \hat{\mu}_2$, where $\hat{\mu}_2$ is defined in (9), both governments take action h ;*
- (ii) *if $\mu_0 < \hat{\mu}_2$, government 1 takes action h if and only if condition (10) holds, and government 2 takes action h if and only if a bad outcome is realized in the first society.*

Proof. (i) Suppose that $\mu_0 \geq \hat{\mu}_2$. If government 1 takes action l , it will be disapproved at least when a bad outcome is realized, and so its payoff is at most $q_l < 1$ regardless of what will happen in society 2. However, if government 1 takes action h , it will always get approved and have an expected payoff 1. This is because even if a good outcome $x_1 = 0$ is realized, citizen 2's interim belief will be still pessimistic enough ($T_h(\mu_0) \geq \hat{\mu}_1$) so that according to Proposition 1 government 2 will take action h as well. Therefore, the two citizens' posterior will be at least $T_h^{[2]}(\mu_0) \geq \hat{\mu}$, and so both governments' actions will be approved.

(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 according to Proposition 1. 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's 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 (10). □

Proposition 2 identifies the condition under which government 1 is relieved from the policymaking dilemma. When $\mu_0 \geq \hat{\mu}_2$, 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.

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.

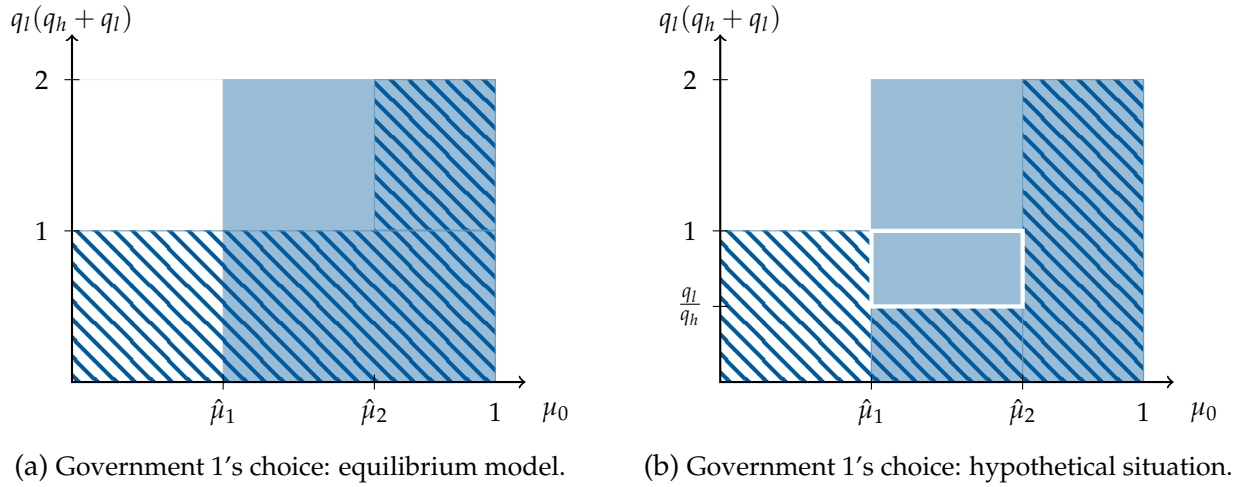


Figure 2: The solid rectangle area describes the set of parameters where government 1 chooses action h in the single-society model. The hatched area in (a) describes the set of parameters where government 1 chooses action h in the two-society model. The region $(0, \hat{\mu}_1) \times [0, 1]$ corresponds 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.

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 (10) holds, and the reverse is true if $\hat{\mu}_1 \leq \mu_0 < \hat{\mu}_2$ and (10) does not hold. In the remaining cases, having the second society has no impact on government 1's policy choice.*

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 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

$$q_h(q_h + q_l) \leq 1, \quad (11)$$

which is equivalent to $q_l(q_h + q_l) \leq q_l/q_h$. This is a more stringent condition than (10). 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 (10) 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 (10) does not hold, both the citizen and the government in society 2 do better than in society 1.*

Proof. 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 (10) 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 case 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 $x_2 = 0$.) The former is larger if and only if $2q_l \leq 1$.

When $\mu_0 < \hat{\mu}_2$ and (10) 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 $x_2 = 0$. \square

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 sufficiently 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 gives 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.

U-shape public opinion. Another implication of our model is that the first society's public approval of strict measures exhibits a U-shape pattern over time. After the crisis being contained by draconian measures, the upsurge of the accusation for overreacting follows. The blame then fades after the counterfactual is observed in the second society.

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.

Timing of re-election. The adverse information externality to the second society relies on the assumption that the second policymaker cannot find a possible counterfactual to justify her efficient policy choice before being evaluated. This assumption gives us another testable implication. When an administration or a policymaker is soon up for re-election, they are more vulnerable to the negative externality and make suboptimal decisions.

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 $i = 1, 2, \dots, n$, (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 $q_l^{n-i}(q_h + q_l) \leq 1$.*

Proof. (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 $1 - q_hq_l^{n-1}$.

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 $q_h^{n-k-1}q_l^k(q_h + q_l) \leq 1$, 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. \square

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 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 Discussions

In this section, we discuss the robustness of the main insight when the two main unconventional assumptions of the model are relaxed. The omitted proofs in this section can be found in the appendix.

4.1 Asymmetric state information and signaling

Suppose now that 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's 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 adding this signaling channel does not change the main insight 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 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

$$v_s = \frac{v_0 \Pr(s|B)}{v_0 \Pr(s|B) + (1 - v_0) \Pr(s|G)}$$

be the government's updated belief that the state is bad after receiving signal s . This is what matters for its policy decision. The interesting case is when $v_g < \hat{\mu} < v_b$, 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 $\hat{\mu}_l^{0,0}$.)

Let

$$p_{s,a} \equiv v_s q_a + 1 - v_s$$

denote the government's expected probability, after receiving signal s , that a good outcome will be realized if it takes action a . For a given signal s , the government's 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 . Similar to the baseline model, 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 Assumption 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 $\pi_{s,a}$ defined above, 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.

We mainly focus on the generic case of $\delta < 1$, i.e., when the government's signal is not perfect.²¹ As shown in Proposition 4 below, there is no separating equilibrium where the government's policy choice perfectly reveals the signal it receives, and the only equilibrium is pooling. (The intuition is discussed below after Proposition 4.) Therefore, in the single-society case, 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 $\delta < 1$. In this case, we show in Proposition 4 below that having the second society can restore the first government's incentive to take the efficient separating strategy, which is again similar to the result in the baseline model.

Proposition 4. *Suppose the government's signal is not perfect (i.e., $\delta < 1$).*

- (i) *In the single-society case, 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 in which regardless of its private signal, the government 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 $\hat{\mu}_1$ takes the same value as in the baseline model.*
- (ii) *In the two-society case, there is an equilibrium in which the first government takes the separating strategy $\sigma_b = 1$ and $\sigma_g = 0$ if $T_h^{1,1} \circ T_h^{1,0}(\mu_0) < \hat{\mu}$ and $p_{b,l}(p_{b,h} + p_{b,l}) < 1 < p_{g,l}(p_{g,h} + p_{g,l})$.*

²¹In the baseline model without the signaling channel, for convenience we have focused on the case of $\delta = 1$. This simplification does not affect the main insight there. In this extension, however, as we discuss below, there can exist a qualitatively different equilibrium in the edge case of $\delta = 1$.

The intuition of no separating equilibrium with $\sigma_b = 1$ and $\sigma_g = 0$ in the single-society case is easy to see in the case 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.

The conditions in result (ii) are qualitatively similar to those in Proposition 2 in the baseline model. The first condition $T_h^{1,1} \circ T_h^{1,0}(\mu_0) < \hat{\mu}$, 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 2 will be optimistic enough so that their government will take action l . The second condition $p_{b,l}(p_{b,h} + p_{b,l}) < 1 < p_{g,l}(p_{g,h} + p_{g,l})$ holds when q_l is sufficiently low, and v_b (v_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., v_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., v_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. In the edge case with $\delta = 1$ (i.e., when the government's signal perfectly reveals the true state), if the citizen ignores the outcome information once she can perfectly infer the state from the government's policy choice, 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, since the citizen perfectly infers the state from the government's action, the government action is always approved and so it has no incentive to deviate. 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.

4.2 A reputation model

As we discussed before, 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.

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 .²² 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,²³ 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.²⁴

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 (3) when her gov-

²²This behavioral-type approach is standard in the reputation literature. See, e.g., [Kreps and Wilson \(1982\)](#) for classic examples and [Mailath and Samuelson \(2015\)](#) for a comprehensive survey.

²³In this extension, unlike in the previous one, considering imperfect signals does not change the main results qualitatively.

²⁴The 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 \(2004\)](#).

ernment's action is a and the outcome is x .²⁵ 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 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 $\lambda_{a,x}^\sigma$ 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 $\lambda_{l,x}^\sigma$ is always 1. When $a = h$ and $x = 1$, the citizen learns that the state must be severe, and so

$$\lambda_{h,1}^\sigma = \frac{\lambda_0 \sigma (1 - q_h)}{\lambda_0 \sigma (1 - q_h) + (1 - \lambda_0)(1 - q_h)} = \frac{1}{1 + \frac{1 - \lambda_0}{\lambda_0} \frac{1}{\sigma}}; \quad (12)$$

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

$$\lambda_{h,0}^\sigma = \frac{\lambda_0 \mu_0 \sigma q_h}{\lambda_0 \mu_0 \sigma q_h + (1 - \lambda_0)(\mu_0 q_h + 1 - \mu_0)} = \frac{1}{1 + \frac{1 - \lambda_0}{\lambda_0} \frac{1}{\sigma} \left(1 + \frac{1 - \mu_0}{\mu_0} \frac{1}{q_h}\right)}. \quad (13)$$

²⁵Notice that if $\gamma = 0$ and each government only cares about its reputation, then action l becomes their dominant strategy as it perfectly signals competence.

(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 $\lambda_{h,0}^\sigma \leq \lambda_{h,1}^\sigma \leq \lambda_0$. 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 $\lambda_{h,1}^\sigma = \lambda_0$, i.e., the high action causes no reputation damage.

If the government takes action h , its expected payoff is $q_h \lambda_{h,0}^\sigma + (1 - q_h) \lambda_{h,1}^\sigma - \gamma(1 - q_h + c)$. 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


$$q_h \lambda_{h,0}^\sigma + (1 - q_h) \lambda_{h,1}^\sigma \geq 1 - \gamma(q_h - q_l - c). \quad (14)$$

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 5 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 5. *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 (14) at $\sigma = 1$; in the two-society case, a competent government 1 takes the high action in the severe state for a larger set of (λ_0, μ_0) .*

Figure 3a illustrates a competent government’s optimal policy in the severe state in the



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

Figure 3: 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.

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 (14) at $\sigma = 1$. (The boundary is decreasing because both $\lambda_{h,0}^\sigma$ and $\lambda_{h,1}^\sigma$ 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 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 3b 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.²⁶

5 Conclusion

This paper provides a framework for studying crisis management with multiple jurisdictions or societies. We first highlight a “damned-either-way” policymaking dilemma: sufficient precautions can contain a crisis, but people may then become skeptical of the severity of the problem and question the costly response; a 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 can be mitigated if 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.

Notice that although we choose the sequential-move model to deliver our main message, the counterfactual effect in our paper 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 policymaking 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.

Our paper has focused on “preventive” policies that if succeed cause the public to question the severity of a potential crisis. 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. When the government is constrained by the policymaking dilemma from implementing preventive policies, it may then reply more

²⁶A 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.

on mitigating measures.²⁷ The counterfactual effect highlighted in this paper 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.

A Appendix: Omitted proofs and other extensions

In this appendix, we report the omitted proofs in Section 4, and two additional extensions: one with a more general information structure, and the other with a more general government payoff structure.

A.1 Omitted proofs in Section 4

Proof of Proposition 4. (i) In the single-society case, we prove that if $\delta < 1$ there is no equilibrium with $\sigma_b = 1$ and $\sigma_g = 0$. (The proof for the other case of $\sigma_b = 0$ and $\sigma_g = 1$ is similar and so omitted.)

For the sake of contradiction, suppose that such an equilibrium exists. The citizen believes that the government is playing the above separating strategy, her posterior after seeing a good outcome will be

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

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 .

²⁷Fox 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.

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

$$T_h^{1,1}(\mu_0) = \frac{1}{1 + \frac{1-\mu_0}{\mu_0} \frac{1}{q_h}}. \quad (15)$$

The off-equilibrium belief is: $T_l^{1,1}(\mu_0) < T_h^{1,1}(\mu_0)$, i.e., a low action with a good outcome is more convincing evidence that the state is mild than a high action with a good outcome. Then we have $\hat{\mu}_h^{1,1} < \hat{\mu}_l^{1,1}$, and so the government will indeed always take action h if $\mu_0 \geq \hat{\mu}_h^{1,1}$.

Consider now 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

$$T_l^{0,0}(\mu_0) = \frac{1}{1 + \frac{1-\mu_0}{\mu_0} \frac{1}{q_l}}. \quad (16)$$

The off-equilibrium belief here is similar as in the previous case: $T_l^{0,0}(\mu_0) < T_h^{0,0}(\mu_0)$. Then 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 (15) (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 proposition.

(ii) 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

$$T_h^{1,1} \circ T_h^{1,0}(\mu_0) < \hat{\mu}. \quad (17)$$

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 (17) 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 $p_{g,l}(p_{g,h} + p_{g,l}) > 1$. \square

Proof of Proposition 5. Single society. Note that the left-hand side of (14) strictly increases in σ . If the opposite of (14) 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 (14) 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 (14) 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 (14). (Such an interior solution of σ always exists in this case since the left-hand side of (14) 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 $\lambda_{h,1}^\sigma$ as defined in (12). 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 $\lambda_{h,1}^\sigma$; if $x_2 = 0$, let $\lambda_{h,0;a_2,0}^\sigma$ be government 1's reputation, which will be specified later. Therefore, government 1's expected payoff is $q_h[\lambda_0 \Lambda_{a_2}^\sigma + (1 - \lambda_0) \Lambda_h^\sigma] + (1 - q_h) \lambda_{h,1}^\sigma - \gamma(1 - q_h + c)$, where

$$\Lambda_a^\sigma \equiv q_a \lambda_{h,0;a,0}^\sigma + (1 - q_a) \lambda_{h,1}^\sigma$$

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

$$q_h [\lambda_0 \Lambda_{a_2}^\sigma + (1 - \lambda_0) \Lambda_h^\sigma] + (1 - q_h) \lambda_{h,1}^\sigma > 1 - \gamma(q_h - q_l - c). \quad (18)$$

Compared to condition (14) in the single-society case, the difference is the square-bracket term (which was simply $\lambda_{h,0}^\sigma$ 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 .²⁸ If $(\lambda_0, T_h(\mu_0))$ is in the light blue area in Figure 3a, 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

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

If $(\lambda_0, T_h(\mu_0))$ is in the dark-blue area in Figure 3a, 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

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

(It is ready to see that $\lambda_{h,0;l,0}^\sigma < \lambda_{h,0;h,0}^\sigma$ 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 $\lambda_{h,0;h,0}^\sigma$.

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 $\lambda_0 \Lambda_{a_2}^\sigma + (1 - \lambda_0) \Lambda_h^\sigma > \lambda_{h,0}^\sigma$, so that (18) holds at $\sigma = 1$ for a larger set of (λ_0, μ_0) than that for (14) at $\sigma = 1$. This is true if

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

²⁸More precisely,

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

for both $a_2 = l$ and h . This is immediate once one notices that $q_l/\lambda_{h,0;l,0}^\sigma + (1 - q_l)/\lambda_{h,1}^\sigma = 1/\lambda_{h,0}^\sigma$ and $q_h/\lambda_{h,0;h,0}^\sigma + (1 - q_h)/\lambda_{h,1}^\sigma < 1/\lambda_{h,0}^\sigma$ and then apply the Jessen's inequality.²⁹ \square

A.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's 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

$$q_h - q_l > q'_h - q'_l. \quad (19)$$

The latter implies that 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

$$c = \hat{\mu}(q_h - q_l) + (1 - \hat{\mu})(q'_h - q'_l) \quad (20)$$

since the high action can also lower the chance of a bad outcome under the mild state. Under condition (19) the citizen will approve action h if and only if her posterior is greater than $\hat{\mu}$.

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

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

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

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

Notice that condition (19) implies $q'_h/q'_l < q_h/q_l$ and so $T_{l,0}(\mu_0) < T_{h,0}(\mu_0)$.³⁰ Let $\hat{\mu}_{a,x}$ solve

²⁹This is actually a consequence of a more general martingale property in our setup: given $\lambda_{h,1}^\sigma = \lambda_{h,0;a_2,1}^\sigma$, we have $\mathbb{E}[\lambda_{h,0;a_2,x_2}^\sigma | \text{severe state}] > \mathbb{E}[\lambda_{h,0;a_2,x_2}^\sigma] = \lambda_{h,0}^\sigma$.

³⁰But condition (19) 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.

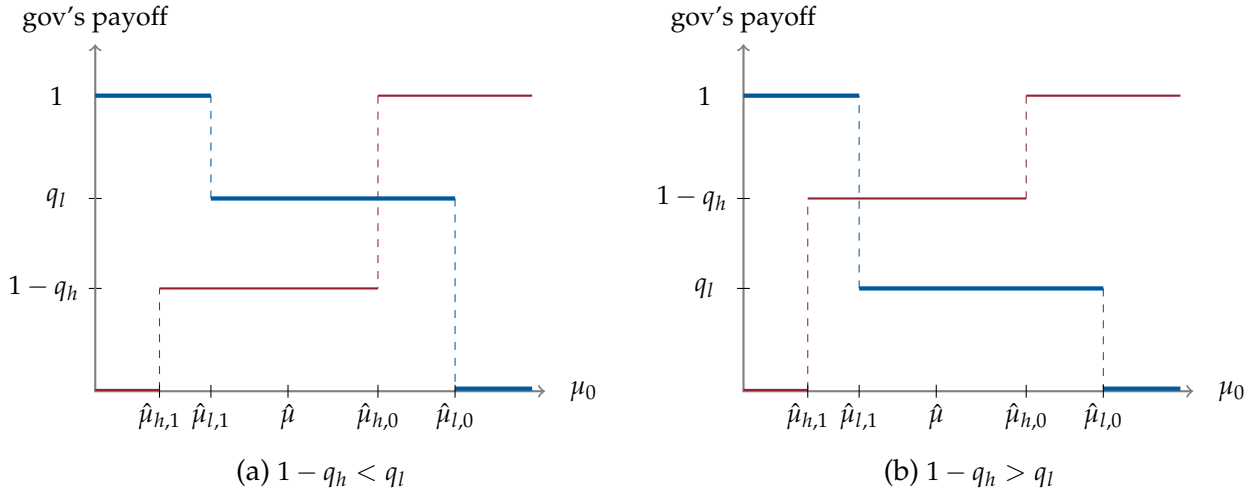


Figure 4: Government's payoff under general information structure: the red thin line corresponds to the government's payoff under action h , while the blue thick line corresponds to the payoff under action l .

$T_{a,x}(\hat{\mu}_{a,x}) = \hat{\mu}$. 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 4. (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 4b.)

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 4a; 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 4b. 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 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.

A.3 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 $q_l(\mathbb{I}_{T_l(\mu_0) < \hat{\mu}} + \alpha_u^+ \mathbb{I}_{T_l(\mu_0) \geq \hat{\mu}})$. 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 (7) 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 (10) 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.³¹

References

- ACHARYA, A., E. GRILLO, AND D. PEI (2021): "Building Trust in Government: The Opportunities and Challenges of Adverse Shocks," Stanford University, Working Paper.
- ALI, S. N., AND N. KARTIK (2012): "Herding with collective preferences," *Economic Theory*, 51(3), 601–626.

³¹Generally, we can define $\bar{v}_a(x, \mu, \mu_0)$ 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, it can be complex to specify the primitive conditions for all these conditions to be satisfied.

- 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.
- BERNANKE, B. S., T. F. GEITHNER, AND H. M. PAULSON (2019): *Firefighting: the financial crisis and its lessons*. Penguin Books.
- BESLEY, T., AND A. CASE (1995): "Incumbent Behavior: Vote-Seeking, Tax-Setting, and Yardstick Competition," *The American Economic Review*, pp. 25–45.
- BIKHCHANDANI, S., D. HIRSHLEIFER, O. TAMUZ, AND I. WELCH (2021): "Information Cascades and Social Learning," Discussion paper, National Bureau of Economic Research.
- 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.
- CAI, H., AND D. TREISMAN (2005): "Does competition for capital discipline governments? Decentralization, globalization, and public policy," *American Economic Review*, 95(3), 817–830.
- 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.
- ESPONDA, I., AND D. POUZO (2017): "Conditional Retrospective Voting in Large Elections," *American Economic Journal: Microeconomics*, 9(2), 54–75.
- FIORINA, M. (1981): *Retrospective Voting in American National Elections*. Yale University Press.
- 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.
- LEWIS, M. (2021): *The Premonition: A Pandemic Story*. Norton.

- 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.
- OTTAVIANI, M., AND P. SØRENSEN (2001): "Information aggregation in debate: who should speak first?," *Journal of Public Economics*, 81(3), 393–421.
- PRAT, A. (2005): "The Wrong Kind of Transparency," *American economic review*, 95(3), 862–877.
- SCHARFSTEIN, D. S., AND J. C. STEIN (1990): "Herd Behavior and Investment," *American Economic Review*, pp. 465–479.
- SKIDMORE, M. J. (2016): *Presidents, Pandemics, and Politics*. Springer.
- TOOZE, A. (2021): *Shutdown: How Covid Shook the World's Economy*. Viking.