Timely Persuasion

Deepal Basak, Zhen Zhou*

August 2019

Abstract

This paper proposes a simple dynamic information disclosure policy that eliminates panic. A panic occurs when some agents take an undesirable action (attack) because they fear that others will do the same, and thus, causing a regime change, even though it was not warranted. We consider a mass of privately informed agents who can attack a regime anytime within a time window. Attack is irreversible, delayed attack is costly, and the delay cost is continuous. We propose a policy called "disaster alert", which at a given date publicly discloses whether the regime is going to change regardless of what the agents do. We show that a timely alert persuades the agents to wait for the alert and not attack if the alert is not triggered, regardless of their private information, and thus, eliminates panic. We demonstrate how this result can be used to design practical policies to reduce, if not eliminate, panic in financial markets.

JEL Classification Numbers: *D02, D82, D83, G28*

Key Words: *Coodination, Information Design, Panic*

*We thank Douglas Gale, Alessandro Pavan, David Pearce, Ennio Stacchetti, Debraj Ray, Laura Veldkamp, Xingye Wu and seminar participants at the Econometric Society meetings, Stony Brook, Peking University, Tsinghua University and NYU Shanghai for their helpful comments and suggestions. We thank Mayank Prakash for excellent research assistance. Basak: Indian School of Business, Email: deepal_basak@isb.edu; Zhou: PBC School of Finance, Tsinghua University, Email: zhouzh@pbcsf.tsinghua.edu.cn

Introduction

In a game of strategic complementarity, often an agent panics and takes an action because he thinks that others will do the same, even when it is not warranted. Imagine some investors who have made direct investments in an emerging market. Suppose an adverse shock hits the emerging economy. The investors have noisy information regarding the severity of the shock. If the shock is too severe, the investors should exit the market. However, even if the actual shock is not severe, the investors could panic and start exiting the market if they think other investors will also exit. Can such panic be avoided? This paper proposes a simple dynamic information disclosure policy that eliminates panic.

The above example can be nicely captured through a canonical regime change game (See Morris and Shin (2003)). A mass of agents decide whether to attack a regime or not. If the regime is strong enough to withstand the aggregate attack, the regime survives, otherwise it fails. The canonical regime change game assumes that the agents move simultaneously. We deviate from this assumption and allow the agents to attack within a time window. It is reasonable to think that after learning about the shock, the investors get a time window to react. We assume that attacking is an irreversible action, delayed attack is costly, and the cost is continuous. For example, the delay cost can be the loss in interest income from not investing elsewhere.

The agents are uncertain about the fundamental strength of the regime and get some noisy private signals about it. The noises may be independent or correlated. We allow for homogeneous or arbitrarily heterogeneous beliefs. Based on these signals, the agents form their beliefs about the fundamental and others' signals. If an agent believes that the regime is not very likely to survive, he attacks. It is possible that many of the other agents do otherwise, and the regime survives. Thus, attacking right away could be a mistake ex-post. We assume that the information structure is such that the agents always believe that attacking right away could be such a mistake with positive probability, regardless of what other do. This is trivially true under standard global game information structure, when the noise distribution has full support. We refer to this assumption as *Doubt*.

There is a principal or an information designer who wants the regime to survive. She commits to a dynamic information disclosure rule: At some date t, she will send a message to the agents based on the exogenous fundamental and the endogenous history of attack until time t. We propose a simple policy, called the "disaster alert." A disaster alert at some date t is triggered by the principal if the regime is no longer strong enough to withstand further attacks, i.e., it is doomed to fail.

Going back to the foreign investment example, this policy is equivalent to saying that the principal will disclose to the investors, based on information up to some date t, whether exiting the market right after that date has become the dominant strategy or not. Thus, when the disaster alert is triggered, the agents will surely attack. On the other hand, if the disaster alert is not triggered, then the agents learn that attacking is not the dominant strategy. However, it may still be wise to attack in case other agents attack after seeing no alert and thus the regime fails because of that. Interestingly though, this strategic uncertainty that others may attack when the alert is not triggered, goes away under endogenous delay. If an agent has waited for the disaster alert, it must be that she will not attack when the alert is not triggered, otherwise there is no positive option value of waiting to justify the cost of a delayed attack.

Thus, once the principal discloses the information, there is no strategic uncertainty left. Agents attack if and only if the alert is triggered. In other words, agents will follow the principal's recommendation. However, this does not mean that agents will always wait for the disaster alert. For example, an agent who receives a very low signal about the fundamental and believes that the disaster alert is very likely to be triggered, may decide to attack based on his private signal rather than wait for the alert.

Since the agents who have waited for the alert will only attack if the alert is triggered, the regime survives for sure when the alert is not triggered. Hence, waiting for the alert and then following the recommendation avoids making a mistake of attacking a regime that survives in the end. On the other hand, attacking right away could be a mistake. In fact, under the *Doubt* assumption, there is a positive probability that attacking right away is a mistake, regardless of what others do. This captures the benefit from waiting for the alert as compared to attacking immediately.

However, delay is costly. The principal can reduce this cost by setting the disaster alert at an earlier date. Since the cost of delay is continuous, for any agent whose information satisfies the *Doubt* assumption, a timely disaster alert policy guarantees that the expected benefit of waiting always outweighs the expected delay cost. Thus, when the principal sets a timely disaster alert, the agents not only follow the principal's recommendation after the alert, but also always wait for the alert regardless of their signals.

This implies that any regime that could have survived if no agent had attacked, will indeed survive in the end. In other words, timely disaster alert eliminates panic.

In our baseline model, we focus on a short time window in which the agents do not receive additional information, and the principal controls the flow of information after the initial time. However, the insight can be extended to the case where agents receive more information from outside sources over time. The extended policy in such a case is to set the timely disaster alert right after the arrival of any new information, which could potentially induce a panic. We show that, as long as the additional learning does not violate the *Doubt* assumption, under the extended policy, no agent would act on their private information. All agents will always wait for the alert and the possibility of panics goes away.

This result shows that panic is a fragile idea. There is a simple way a principal can manipulate the agents and stop them from panicking. The principal does not need to know the private signals each agent receives. She uses a public disclosure policy, and she achieves the first best. More importantly, the policy does not violate the principal's ex-post incentive compatibility. To see this, note that when the alert is triggered, the regime is doomed to fail regardless of what message the principal sends. This means the principal does not need ex-ante commitment to implement such a policy.

We apply this insight to construct a practical policy that could dissuade the investors from leaving an emerging market. The problem is slightly different since the investors earn flow payoff. The disaster alert is not a deliberate policy rather the agents learn about the disaster when they see that the flow payoff has dropped. Thus, there is an endogenous and continuous disaster alert. The policy we propose is simply a tax on capital flight. Suppose the maximum flow payoff an investor can get if he stays is $\overline{r}$, while if he exits he gets a flow payoff of $r < \overline{r}$. Suppose the investor has to pay a tax on the flow payoff he would receive when he leaves. If the tax rate μ is high enough such that $(1 - \mu)\overline{r} < r$, then waiting is costly. We can use our main insight to show that if the fundamental shock does not cause a disaster, investors will not panic and cause a disaster.

In practice, not all the features of a model are likely to hold. We discuss which assumptions are essential for the result and which assumptions can be relaxed. First, we discuss the solution concept – extensive form rationalizability. Under the doubt assumption, if agents are rational and they believe others are rational, there will be no panic. Note that this does not require common knowledge of rationality. Moreover, even if the agents believe that there is a small chance that others are not perfectly rational, and therefore, the regime may not survive even when the disaster alert is not triggered, the result is robust.

The *Doubt* assumption is not only sufficient, but also necessary for this result. If the *Doubt* assumption fails, then panic cannot be eliminated. For example, if the agents commonly believe that the regime will definitely fail if all the agents attack, then a timely disaster alert cannot eliminate the possibility that all the agents attack. However, if the agents have heterogenous beliefs, then even if an agent does not have doubt, he may believe that some other agents doubt, or even if they do not doubt, they believe that some other agents doubt, and so on. In such an environment, the doubt assumption is unnecessary.

The two features of the model that are essential for the result are – (1) continuity of the delay cost and (2) irreversibility of attack. If the delay cost is discontinuous, then the disaster alert could be "too late" in the sense that the agents cannot save their money after learning that the regime is going to change regardless of what they do; and if attack is reversible, then the agents may decide to leave early and come back later if the disaster alert is not triggered, but such actions can trigger the disaster alert, even when it is not warranted. We specialize to an independent noisy signal and uniform prior environment, as is standard in the global game literature, and construct equilibrium in which the agents panic. We show that when the cost of delay is discontinuous, the timely disaster alert policy will reduce panic, but cannot eliminate panic. We characterize the limit to timely persuasion depending on the magnitude of the discontinuity, i.e., how much the the agents can save after they get the early warning that the regime is going to switch.

In reality, we often see policy makers make the point of moving early to assure the market. Consider the stress tests for banks as an information disclosure policy in the financial regulation as in Inostroza and Pavan (2017) and Goldstein and Huang (2016). According to Timothy Geithner, the Secretary of the U.S. Treasury, *"the plan aimed to impose transparency on opaque financial institutions and their opaque assets in order to reduce the uncertainty that was driving the panic"*. The supervisory guidance on Stress Testing published by FED, FDIC and OCC mentioned that *" a banking organization should have the flexibility to conduct new and ad hoc stress tests* ***in a timely manner*** *to address rapidly emerging risks"*. [1] Our model formalizes the argument how timely stress tests help in removing the strategic uncertainty and in eliminating panic.

Related Literature The two most closely related papers are Goldstein and Huang (2016) and Inostroza and Pavan (2017). Similar to this paper, the above mentioned papers also consider a regime change game with privately informed agents. The authors consider an information designer who commits to an information disclosure rule. While Goldstein

[1]See SR Letter 12-7 for Supervisory Guidance on Stress Testing for Banking Organizations with more than 10 Billion in Total Consolidated Assets https://www.federalreserve.gov/supervisionreg/srletters/sr1207.htm.

and Huang (2016) propose a simple stress test policy, Inostroza and Pavan (2017) design the optimal disclosure policy, which under some conditions can be a stress test. In a stress test policy, the designer discloses whether $\theta \geq k$ or not, for some fixed k. Thus, the regimes with $\theta \in (0, k)$ will not survive even when it could have survived if agents were not attacking. Also, under some condition, the principal can do better by supplementing this policy with independent noisy private messages and thus increasing the heterogeneity in the agents' beliefs. Nevertheless, the optimal policy does not eliminate panic. However, if there is a small time window, the principal can eliminate panic. The crucial difference is that the principal can disclose information regarding not only the exogenous fundamental, but also the endogenous history of attack.[2] The disaster alert can be thought of as the weakest stress test ($k = 0$). Under a stress test, there can be multiple equilibria. But when k is sufficiently large, even in the worst equilibrium, i.e., the one in which agents attack most aggressively, they will not attack a regime that passes the stress test. Note that under the stress test $k > 0$, when the regime fails the test, the principal wants to lie to the agents. This violates the ex-post incentive compatibility. This means the principal needs ex-ante commitment power to implement a stress test policy with $k > 0$. This is not the case with disaster alert. If the regime is doomed to fail, the principal cannot do any better by lying.

Basak and Zhou (2019) consider a similar problem in which agents moves sequentially in an exogenous order. The authors show that if the principal runs viability tests (weakest stress tests) sufficiently frequently, then the unique cutoff equilibrium involves no panic. It remains an open question whether this result can be extended to rationalizability. In contrast, in this paper, agents get a short time window and endogenously decide when to attack, and propose a different policy – a one time but timely disaster alert. The result is not limited to cutoff equilibrium. As long as agents are almost rational and they believe others are almost rational, there will be no panic.

This paper contributes to two literature. First, the recent growing literature on dynamic information design. See Kamenica (2018) for a survey of the information design and bayesian persuasion literature. Ely (2017) is the first paper to extend the static bayesian persuasion problem of Kamenica and Gentzkow (2011) to a dynamic setting, but the author only considers history independent disclosure. Makris and Renou (2018) generalizes

[2]If the disaster alert can be set only at time 0 before any agent moves, then it will not eliminate panic. Also, we argued in Section 5.3 that if it already too late to act on the disaster alert (violation of continuity), then a disaster alert policy will have limited success. A simultaneous move regime change game can be thought of as a dynamic regime change game where agents cannot save any money after the disaster alert is triggered.

the correlated equilibrium of Bergemann and Morris (2016) to multistage game. While these papers assume that the players have noisy information about a payoff relevant state, Salcedo (2017) consider a complete informations stage game but introduces uncertainty regarding the game tree that governs the play. Doval and Ely (2019) consider a dynamic information design environment where the designer neither knows the private signals of agents, nor the game tree. Finally, the paper contributes to the dynamic coordination game literature. To model panic, this paper borrows from the global game of regime change literature. See Morris and Shin (2003), Szkup (2017) for recent developments and Angeletos and Lian (2017) for an excellent survey. Similar to Gale (1995), Dasgupta (2007) and Dasgupta, Steiner and Stewart (2012), we extend this canonical regime change game to allow for endogenous delay in attack.

The rest of the paper is organised as follows. Section 1 describes the model. Section 2 demonstrates that a timely disaster alert eliminates panic. Section 3 shows that the result can be extended to arrival of new information over time.

Section V considers a FDI application where the flow payoffs naturally works as a continuous disaster alert. In Section VI, we discuss the relation of this paper to global games literature and other relevant papers. The proofs that are not in the paper can be found in the appendix.

1 Model

Players and Actions The economy is populated by a principal, a continuum of agents, indexed by $i \in [0,1]$, and a regime. A shock hits the regime and it is commonly known. We normalize the date at which the shock hits as 0. Once the shock hits, the agents get a small time window $[0,T]$ to decide whether they want to attack the regime or not. Attacking can be taken as the action of exiting from a market, withdrawal of early investment, attacking a currency regime, making redemption from a mutual fund, etc.

Let us denote the action; attack as 1 and not attack as 0. An agent i chooses $a_i \in [0,T]^{\{0,1\}}$ which describes whether he attacks or not at any date. Attacking is an irreversible action while not attacking is reversible. Hence, if agent i has already attacked by some t, he has no more decision to make. However, if he has not attacked, then he has the option to attack at any time between t and T, or not attack at all. If agent i decides to attack at time $t_0 \in [0,T]$, then $a_{it} = 0$ for any $t \in [0,t_0)$ and $a_{it} = 1$ for any $t \in [t_0,T]$. For any agent i who decides to attack, i.e., $a_{iT} = 1$, let us denote the time of attack as

$t_i \equiv \min\{t \in [0,T] | a_{it} = 1\}$. The time of attack is defined to be $t_i = \infty$ for an agent i who does not attack at all. Hence, a_i can simply be represented by the time of attack $t_i \in [0,T] \cup \infty$. At any t within the time window, the mass of agents who already attack is

$$N_t \equiv \int_{i\in[0,1]} \mathbf{1}\{i | t_i \leq t\} di.$$

By definition, the mass of attacks $N_t \in [0,1]$ is (weakly) increasing in time t.

Fundamental States The underlying state of the economy is captured by θ. We refer to it as the fundamental strength of the regime. If the intensity of the shock is more severe, then the strength of fundamental is weaker and thus θ is lower. At time 0, nature draws a state $\theta \in \Theta$, where Θ is a compact subset of $\mathbb{R}$. It is common knowledge that θ is drawn from some distribution Π with smooth density π strictly positive over Θ.

Regime Outcome Let $r \in \{0,1\}$ denotes the fate of the regime. We denote by $r = 0$ the event that the regime survives, and by $r = 1$ the complement event that the regime does not survive. The fate of the regime depends on the fundamental state (θ) and the aggregate attack until the end (N_T). The regime survives, i.e., $r = 0$, if, and only if $R(\theta, N_T) \geq 0$, where $R(.)$ is a continuous function that is increasing in θ and decreasing in N_T.

Payoff The agents are ex-ante identical and expected utility maximzers. If an agent does not attack ($t_i = \infty$), then he gets

$$v(\theta, N_T) = \begin{cases} g(\theta, N_T) & \text{if } r = 0 \\ l(\theta, N_T) & \text{if } r = 1, \end{cases}$$

and if he attacks at time t, then he gets $u(t)$. We normalize $u(0) = 1$. As is standard in the static regime change game, [3]

$$g(\theta, N_T) > 1 > l(\theta, N_T).$$

This captures the fact that if the regime is going to survive ($r = 0$), then not attacking is the desirable action, and if the regime is not going to survive ($r = 1$), then attacking is the

[3] See, for example, Inostroza and Pavan (2017).

desirable action. We allow $g(.)$ and $l(.)$ to be non-monotonic in the arguments, but there exists $\underline{g} > 1$ such that $g(.) \geq \underline{g}$ and there exists $\bar{l} < 1$ such that $l(.) \leq \bar{l}$.

The difference with this static payoff is that the agent has multiple opportunities to attack and delaying attack is costly, i.e., $u(t)$ is decreasing in t. Consider attack as exiting a market. Then, delaying exit means the investor is losing interest he could have earned by investing the money elsewhere. We assume that $u(T) > \bar{l}$. This means that even at the last minute, if the agents learn that the regime will not survive ($r = 1$), attacking is the desirable action. We further assume that $u(t)$ is *Lipschitz continuous* in t. In some applications, this may be too strong an assumption, and we will discuss this in Section 5.

Dominance Region There exists $\underline{\theta}, \overline{\theta} \in \Theta$ such that $R(\underline{\theta}, 0) = R(\overline{\theta}, 1) = 0$. This means that when $\theta \in \Theta^L = \Theta \cap (-\infty, 0)$, the regime cannot cannot survive regardless of whatever strategy the agents take, and when $\theta \in \Theta^U = \Theta \cap [1, +\infty)$, the regime will always survive regardless of whatever strategy the agents take. We refer to Θ^U (or Θ^L) as the upper (or lower) dominance region where not attacking $t_i = \infty$ (or attacking right away $t_i = 0$) is the dominant strategy. We assume that $\Theta^U, \Theta^L \neq \emptyset$.

Exogenous Information In addition to the common prior Π, each agent i receives a signal $s_i \in \mathbb{R}$ about θ before they decide when to attack (if at all). Given any underlying fundamental θ, the signal profile $s(\theta) \in \mathbb{R}^{[0,1]}$ are drawn from a distribution $F(s|\theta)$ with associated density $f(s|\theta)$. Note that this allows for any arbitrarily correlated signals, ranging from independent private signals to public signals. We are interested in a short time window, in which the agents do not receive any more information about the fundamental or observe other agents' actions.

For illustration, we will sometime use independent private noise (as is standard in the global game literature): each agent i receives a noisy signal $s_i = \theta + \sigma\epsilon_i$, where ϵ_i are independent and follows an identical distribution F, and $\sigma > 0$ scales the noise.

Principal The principal's payoff only depends on whether the regime survives or not. She gets 1 if the regime survives and 0 if it does not.[4] The principal does not have access to the agents' noisy private information.

[4] It is easy to generalize to the case where the principal also wants to minimize aggregate attack condition the regime's survival.

Disclosure Policy For any $\tau \in [0, T]$, the principal can disclose some information to the agents based on the exogenous fundamental θ and the endogenous attack so far N_τ. We consider a continuous time model which means that an agent who has not attacked by time τ, gets the opportunity to attack again at $(\tau + dt)$, where $dt \to 0$. At any date τ we allow for a sequence of events to occur. Accordingly, we define τ^- and τ^+. At τ^-, an agent can attack, while at τ^+, the principal can disclose information. There is no time discounting between τ^- and τ^+. Let $\mathcal{S}$ be a compact metric space defining the set of possible disclosures to the agents, and $m_i(\tau, \theta, N_\tau) \in \mathcal{S}$ be the message to agent i. A general disclosure policy is $\mathit{\Gamma} = (\pi, \mathcal{S})$ consists of the set of disclosed messages $\mathcal{S}$ and the disclosure rule $\pi : [0, T] \times \mathit{\Theta} \times [0, 1] \to \mathit{\Delta}(\mathcal{S}^{[0,1]})$. The feature of endogenous move enables the principal to select the time of disclosure, and to make this information disclosure policy history dependent.

Robust Design We use rationalizability in extensive form game à la Pearce (1984) as our solution concept. Given a disclosure policy $\mathit{\Gamma}$, let $\mathcal{R}(\mathit{\Gamma})$ be the set of all possible rationalizable strategy profiles $a \equiv (a_i(s_i))$. Define

$$\mathit{\Theta}^F(\mathit{\Gamma}) := \{\theta \in \mathit{\Theta} | \theta < N_T(a) \text{ for some } \; a \in \mathcal{R}(\mathit{\Gamma})\}.$$

Thus, if $\theta \notin \mathit{\Theta}^F(\mathit{\Gamma})$, then the regime will survive regardless of whatever rationalizable strategies the agents play, and if $\theta \in \mathit{\Theta}^F(\mathit{\Gamma})$, then the regime may not survive. The principal's objective is

$$\min_{\mathit{\Gamma}} \mathit{\Pi}(\mathit{\Theta}^F(\mathit{\Gamma})).$$

That is, the principal anticipates, state by state, the "worst possible" outcome that is consistent with the agents playing some rationalizable strategy, and chooses the policy $\mathit{\Gamma}$ to minimize the ex-ante chance that the regime may not survive.

Note that when $\theta \in \mathit{\Theta}^L$, the regime fails irrespective of the size of the attack. Hence, any disclosure policy $\mathit{\Gamma}$ cannot endure such a regime, i.e., $\mathit{\Theta}^L \subseteq \mathit{\Theta}^F(\mathit{\Gamma})$. A regime could also fails even when it is not warranted ($\theta \notin \mathit{\Theta}^L$) because the agents attack thinking that others will attack. We refer to this as *panic-based attacks*. Let us define $\mathit{\Theta}^P(\mathit{\Gamma}) := \mathit{\Theta}^F(\mathit{\Gamma}) \setminus \mathit{\Theta}^L$ for any policy $\mathit{\Gamma}$ as the set of fundamental in which the regime can fail because of panic-based attacks. If $\mathit{\Theta}^P(\mathit{\Gamma}) = \emptyset$, then we say that the policy $\mathit{\Gamma}$ eliminates panic.

2 Main Result

We restrict our attention to a simple information disclosure policy. We show that, under some reasonable assumption on agent's exogenous information structure, this simple disclosure policy eliminates panic. In other words, this policy induces the agents to perfectly coordinate their actions and never attack a regime when it is not warranted ($\theta \notin \Theta^L$).

Disaster Alert

We refer to the following disclosure policy as *disaster alert*. The principal only discloses information once at some $\tau \in [0, T]$. The public signal d^τ is generated based on the underlying fundamental θ and the history of attacks N_τ as follows

$$d^\tau(\theta, N_\tau) = \begin{cases} 1 \text{ if } R(\theta, N_\tau) < 0 \\ 0 \text{ otherwise.} \end{cases}$$

We denote this binary public disclosure policy as Γ^τ. Upon receiving the signal $d^\tau = 1$, agents understand the regime cannot survive in the end, i.e., $R(\theta, N_T) < 0$ (Since N_t is weakly increasing in t). Hence, for agents who have not attacked, it is the dominant strategy to attack at time $\tau + dt$. In this sense $d^\tau = 1$ acts as an alert for disaster. On the other hand, if the alert is not triggered, or $d^\tau = 0$, agents understand that $R(\theta, N_\tau) \geq 0$, and thus, the regime will survive if no agent attacks the regime after time τ. In the spirit of Bayesian Persuasion, this can be thought of as the principal sending a recommendation at time τ to the agents to attack when the disaster alert is triggered, and not attack otherwise.

Option Value of Waiting

Under the policy Γ^τ, agents will only have one chance to get new information at τ^+. Hence, attacking at any time $t_i \in (0, \tau^-]$ is dominated by attacking at time $t_i = 0$ since delayed attack is costly. Similarly, after receiving the new information d^τ, attacking at anytime $t_i \in (\tau + dt, T]$ is dominated by attacking immediately after the disclosure at $\tau + dt$.

Lemma 1 *(Option Value) Under the disclosure policy Γ^τ, for any noisy signal s_i, the only rationalizable strategies are*

- $\mathcal{A}$: *attack at time* 0*, i.e.,* $a_{it}(s_i) = 1$ *for all* $t \in [0, T]$*, and*

$\mathcal{W}$: *wait until time τ for the disaster alert, i.e., $a_{it}(s_i) = 0$ for all $t \in [0, \tau + dt)$, and then follow the principal's recommendation $a_{it}(s_i, d^\tau) = d^\tau$ for $t \in [\tau + dt, T]$.*

Proof. First, it is not rational for agents to attack at any time other than 0 and $\tau + dt$. Secondly, when $d^\tau = 1$, the dominant action is attack. Hence, the only possible strategies are: $\mathcal{A}$, $\mathcal{W}$ and attacking at time $\tau + dt$ independent of d^τ. Let us call this third strategy $\mathcal{W}'$. Note that the strategy of $\mathcal{W}'$ generates a payoff of $u(\tau + dt)$, which is strictly less than $u(0)$. Thus, it is strictly dominated by $\mathcal{A}$. ■

For any agent who decides to wait for the disclosure (instead of attacking immediately), the information that will be disclosed right after time τ must be valuable to him. That means he will never take the same action regardless of the future disclosed information. Otherwise, there is no option value associated with the information arriving in the future and hence he will not wait. Chamley and Gale (1994) and Gul and Lundholm (1995) made a similar argument in the context of social learning in which an agent can learn from others' actions, but such actions do not affect his payoff. The intuition is simple - consider two agents deciding whether to attack at date 1 or date 2. If an agent waits to see whether the other agent attacks or not, it must be that he will take different actions conditional on whether the other agent attacks at date 1 or not. Otherwise, there is no positive option value of waiting.

In this paper, the principal controls the information flow after date 0. If the disaster alert is triggered ($d^\tau = 1$), then attacking is the dominant strategy for an agent. Therefore, for positive option value of waiting, it must be that the agent will not attack when the alert is not triggered.

This implies that when the disaster alert is not triggered ($d^\tau = 0$), the regime will survive in the end. This is because there is no further attack after time τ when $d^\tau = 0$, i.e., $N_T = N_\tau$. For that reason, no alert ($R(\theta, N_\tau) \geq 0$) implies the survival of the regime, i.e., $R(\theta, N_T) \geq 0$. Consequently, under the policy Γ^τ, the agents who decide to wait for the information disclosure, perfectly coordinate their actions. This completely removes the strategic uncertainty after time τ since the agents understand perfectly what other agents would do after getting the new information.

This is in sharp contrast with the static regime change game, in which the agents move simultaneously. In a static regime change game, the strategic uncertainty cannot be removed by publicly disclosing that attacking is not the dominant strategy. To see this, consider a disaster alert before the agents make any decision. This alert is triggered if $\theta < 0$. If this alert is not triggered, then the agents know that $\theta \geq 0$, or the regime will survive if no agent attacks. When this is publicly known, one possible equilibrium outcome is that no agent attacks and the regime survives. However, this is not the unique rationalizable strategy. If an agent receives a low signal and believes others will attack, then he will attack as well. In fact, Angeletos, Hellwig and Pavan (2007) show that there are many other possible equilibria in which the regime could fail because of panic-based attacks. However, under endogenous timing, if an agent with low signal decides to attack even after the disaster alert is not triggered ($d^{\tau} = 0$), then he is better off not waiting for the alert at all. So, it follows from Pearce (1984)'s extensive form rationalizability that either an agent attacks right away, or waits for the alert, and if he waits for the alert, then he follows the principal's recommendation afterwards. However, it is possible that the agent do not wait for the alert and attack right away, and by doing so, they trigger the alarm.

Reasonable Doubt

The following assumption restricts the information generating process F. It says that regardless of whatever noisy signal an agent receives, he always assigns some positive chance that $\theta \in \Theta^U$. In other words, he always has some doubt that attacking is a mistake regardless of what other agents do.

Assumption 1 *(Doubt) There exists $\varepsilon > 0$ such that, any agent i with noisy signal s_i believes that*

$$\mathbb{P}(\theta \in \Theta^U | s_i) = \frac{\int_{\theta \in \Theta^U} f_i(s_i|\theta)\pi(\theta)d\theta}{\int_{\theta \in \Theta} f_i(s_i|\theta)\pi(\theta)d\theta} > \varepsilon,$$

where $f_i(s_i|\theta) = marg_{s_{-i}} f(s|\theta)$.

In particular, if f_i has full support and it is bounded away from 0, then the above assumption holds true. [5]

Timely Disaster Alert

Lemma 2 *(Timely Alert) There exists $\hat{\tau} > 0$, such that under the disclosure policy Γ^{τ}, where $\tau < \hat{\tau}$, for any signal structure satisfying Assumption 1, the only rationalizable strategy for an agent with signal s_i is $\mathcal{W}$, i.e., wait for the disaster alert and then follow the principal's recommendation.*

[5] A sufficient condition that validates this doubt assumption is that $f_i(s_i|\theta) > \frac{\varepsilon}{\Pi(\theta^U)}$ for all $i \in [0, 1]$, $s_i \in \mathbb{S}$ and $\theta \in \Theta^U$.

Proof. Consider an agent who has decided to wait for the disclosure (plays $\mathcal{W}$). If $R(\theta, N_\tau) < 0$, the alert triggers ($d^\tau = 1$). Then, he attacks at time $\tau + dt$ and gets $u(\tau + dt)$. Otherwise, $R(\theta, N_\tau) \geq 0$ and the alert does not trigger ($d^\tau = 0$). We know from Lemma 1 that the agent will not attack. Thus, the expected payoff from playing $\mathcal{W}$ is

$$\mathbb{P}(d^\tau = 1|s_i)u(\tau + dt) + \mathbb{P}(d^\tau = 0|s_i)E\left(v(\theta, N_T)|s_i, d^\tau = 0\right).$$

It follows from Lemma 1 that no agent who has waited will attack, i.e., $N_T = N_\tau$. This implies that when the alert is not triggered ($d^\tau = 0$), $R(\theta, N_T = N_\tau) \geq 0$, i.e., the regime survives ($r = 0$). Recall that if $r = 0$, then by not attacking the agent gets $g(\theta, N_T = N_\tau)$. Finally, since $dt \to 0$, the expected payoff from playing $\mathcal{W}$ simplifies to

$$\mathbb{P}(d^\tau = 1|s_i)u(\tau) + \mathbb{P}(d^\tau = 0|s_i)E\left(g(\theta, N_T = N_\tau)|s_i\right).$$

While, the expected payoff from attacking immediately ($\mathcal{A}$) is $u(0) = 1$. Hence, the expected payoff difference from strategy $\mathcal{W}$ as compared to $\mathcal{A}$ is

$$D(\mathit{\Gamma}^\tau, s_i) = \mathbb{P}(d^\tau = 1|s_i)(u(\tau) - u(0)) + \mathbb{P}(d^\tau = 0|s_i)(E(g(\theta, N_\tau)|s_i) - 1). \tag{1}$$

Since $u(t)$ is Lipschitz continuous, $u(0) - u(\tau + dt) \leq K\tau$ for some positive finite K. It follows from Assumption 1 that $\mathbb{P}(d^\tau = 0|s_i) > \varepsilon$. Also, recall that $g(\theta, N) \geq \underline{g} > 1$. Therefore,

$$D(\mathit{\Gamma}^\tau, s_i) > -(1 - \varepsilon)K\tau + \varepsilon(\underline{g} - 1).$$

Define

$$\hat{\tau} := \frac{\varepsilon}{1 - \varepsilon}\left(\frac{\underline{g} - 1}{K}\right).$$

If $\tau < \hat{\tau}$, then $D(\mathit{\Gamma}^\tau, s_i) > 0$ for any s_i. This implies an agent prefer to wait for the alert and then follow the principal's recommendation ($\mathcal{W}$) rather than attacking right away ($\mathcal{A}$), regardless of his private signal. ■

The tradeoff agents face when choosing between attacking immediately ($\mathcal{A}$) and the strategy of wait and see ($\mathcal{W}$) is as follows. A cost $(u(0) - u(\tau))$ (lost interest income) is associated with a delayed attack when the alert ($d^\tau = 1$) is triggered. While, taking the strategy of wait and see can prevent agents from making a mistake by moving early, i.e., attacking a regime that survives. The benefit from this more informed choice is at least $(\underline{g} - 1)$. Assumption 1 guarantees that regardless of the signal s_i, an agent assigns positive probability that the alert will not be triggered regardless of what other agents do, i.e., attacking is definitely a mistake. Because of that, the benefit from waiting is strictly positive regardless of what other agents would do. Lemma 2 shows that, if the disaster alert can be set in a timely manner, the cost of delay will be limited (strictly lower than the expected benefit), which makes the strategy of wait and see ($\mathcal{W}$) a strict dominant strategy.

No Panic

Theorem 1 *(No Panic) Under the disclosure policy $\varGamma^{\tau}$ with $\tau < \hat{\tau}$, if the information structure satisfies Assumption 1, then there is no panic, i.e.,*

$$\varTheta^{P}(\varGamma^{\tau}) = \emptyset.$$

Proof. The only rationalizable strategy is $\mathcal{W}$ when $\tau < \hat{\tau}$ (Lemma 2). Thus, $N_{\tau} = 0$. For any regime with $\theta \notin \varTheta^{L}$, $R(\theta, N_{\tau} = 0) \geq 0$. Hence, no alert will be triggered and no further attack happens, or $N_{T} = N_{\tau} = 0$ (Lemma 1). Therefore, any regime with $\theta \notin \varTheta^{L}$ survives (since $R(\theta, N_{T}) \geq 0$). Hence, $\varTheta^{F}(\varGamma^{\tau}) = \varTheta^{L}$ and $\varTheta^{P}(\varGamma^{\tau}) = \emptyset$. ■

Theorem 1 follows immediately from Lemma 1 and Lemma 2. Under a timely disaster alert, all agents, regardless of their noisy signal, would wait for the disclosure and follow the recommendation of the principal afterwards. Since all the agents are waiting, any regime that can survive without attack ($\theta \notin \varTheta^{L}$) will not trigger the alert. Since the alert is not triggered, agents will follow the principal's recommendation and not attack. Thus, any regime that can survive without any attack, will survive for sure. In other words, timely disaster alert eliminates panic.

Under strategic complementarity, "runs" are common. Runs could be based on fundamental ($\theta \in \varTheta^{L}$), but it could also be because of panic ($\theta \in \varTheta^{P}$). The principal cannot save a regime when run is based on fundamental, but the above theorem shows that panics can be eliminated. More importantly, it can be eliminated using a very "simple" disclosure policy. Recall that the principal does not have access to the agents' private signals. The disaster alert policy does not require disclosure conditional on such private signals. Moreover, the policy is a binary and public disclosure – whether the alert is triggered or not, and it does not send private messages to the agents. Finally, the alert is triggered only when $R(\theta, N_{\tau}) < 0$, i.e., the regime cannot survive regardless of what the agents do. This means the principal cannot get a higher payoff by misreporting. Thus, the principal's ex-post incentive compatibility holds.

In static regime change game, the optimal disclosure policy cannot eliminate panic. Moreover, such polices requires ex-ante commitment (as is standard in the bayesian persuasion or information design literature). For example, Inostroza and Pavan (2017) show that in some situations, the optimal disclosure policy is a tough stress test, i.e., disclose whether $\theta \geq k$ or nor, for some $k > 0$. It is clear that, under this policy, if the regime fails the stress test $\theta < k$, the principal would want to misreport. In sharp contrast, if there is a small time window in which the agents can attack (irreversibly), while delayed attack is costly, then panic can be eliminated by a very simple public disclosure policy.

3 New Information and Repeated Disaster Alert

So far, we focus on a short time window $[0, T]$ in which the agents could react to a bad news. In sharp contrast to the simultaneous move game, we show that the principal can exploit this endogenous timing and stop agents from panicking. The insight has nothing to do with the length of the time window. But if the time window is not small, we may expect that new information may arrive over time. For example, if the time window is a month, the agents may receive weekly updates regarding the fundamental. In this case, the agents cannot be certain that even if the timely disaster alert is not triggered, agents will not panic later and attack when they receive new information. Here, we consider the optimal disclosure policy with exogenous arrival of new information.

Suppose that agent i receives a noisy private signal s_i^0 at date 0^- (as before) and s_i^1 at date t_1^- for some $t_1 \in (0, T)$ about the fundamental θ. The agents can act based on the private signal they receive as early as in the same period, i.e., at date 0^- and t_1^-. We maintain the same assumptions regarding the fundamental state and the noisy information. We generalize the doubt assumption to

$$\exists \epsilon > 0, P(\theta \in \Theta^U | s_i^0, s_i^1) > \epsilon \text{ for any } s_i^0 \text{ and } s_i^1.$$

Restricting to two signals is without the loss of generality. The general case for more than two signals can be derived analogously.

Extended Disaster Alert Policy With new arrival of exogenous information, a natural extension of the one-shot disclosure policy is to set the disaster alert right after any new information arrives. In this new information environment, the principal sets a disaster alert at τ as well as at $t_1+\tau$. To reduce burden of notation, we use the same notation Γ^τ to capture this modified policy. $d^\tau=1$ means the first disaster alter is triggered at time τ^+ and $d^{t_1+\tau}=1$ means the second disaster alert is triggered at time $(t_1+\tau)^+$.

Theorem 2 *Under the extended disclosure policy Γ^τ with $\tau<\hat{\tau}$, there is no panic even when new noisy information arrives over time.*

Proof. See Appendix. ■

The basic argument is an extension of our main result. First of all, if the first alert has been triggered at time τ, i.e., $d^\tau=1$, then all agents have attacked. Thus, there is no need to think about any decision making after the new information arrives. Let us assume otherwise and start our analysis from time t_1. Consider any agent who has not attacked before time t_1 and receives the new information. For and possible signal s_i^0 and s_i^1, he either attacks at time t_1, or waits for the second alert and then attacks iff $d^{t_1+\tau}=1$. Thus, as in Lemma 1, all agents who have waited for the second alert, will not attack when $d^{t_1+\tau}=0$, and thus the regime survives when $d^{t_1+\tau}=0$. The agent gets e^{-rt_1} from attacking at t_1, while waiting gives

$$\mathbb{P}(d^{t_1+\tau}=1|s_i^0,s_i^1,d^\tau=0)e^{-r(t_1+\tau)}+\mathbb{P}(d^{t_1+\tau}=0|s_i^0,s_i^1,d^\tau=0)g$$

Since

$$\mathbb{P}(d^{t_1+\tau}=0|s_i^0,s_i^1,d^\tau=0)=\frac{\mathbb{P}(d^{t_1+\tau}=0|s_i^0,s_i^1)}{\mathbb{P}(d^\tau=0|s_i^0,s_i^1)}\geq\mathbb{P}(\theta\in\Theta^U|s_i^0,s_i^1)>\epsilon$$

for all possible s_i^0 and s_i^1, it is dominant strategy for the agents to wait for second disaster alert and not attack at t_1 (as in Lemma 2). Thus, any agent who does not attack after the first alert is not triggered ($d^\tau=0$), will wait for the next disclosure ($d^{t_1+\tau}$) and only attack when the second alert is triggered ($d^{t_1+\tau}=1$).

Now let us move to time 0. It follows from Lemma 1 and the above argument that an agent will either play : ($\mathcal{A}$) attack at time 0, or ($\mathcal{W}^2$) wait for the first disclosure (d^τ), then attack immediately if the alert is triggered ($d^\tau=1$), otherwise wait for the second alert ($d^{t_1+\tau}$), and then attack immediately if the second alert is triggered ($d^{t_1+\tau}=1$), otherwise do not attack at all. This means $d^{t_1+\tau}=0$ whenever $d^\tau=0$. The strategy $\mathcal{A}$ generates a payoff of 1. While, the strategy $\mathcal{W}^2$ generates

$$\mathbb{P}(d^\tau = 1|s_i^0)e^{-rdt} + \mathbb{P}(d^\tau = 0|s_i^0)g$$

It follows from the same argument as in Lemma 2 that an agent will always play $\mathcal{W}^2$ rather than $\mathcal{A}$. When $\theta \geq 0$, regardless of their signal s_i^0, the agents do not attack immediately. Hence, the first alert is never triggered, or $d^\tau = 0$. Therefore, the agents do not attack after the first alert. At time t_1, regardless of the new information s_i^1, the agents do not attack and again wait for the next disaster alert. This means the second disaster alert will not be triggered either, or $d^{t_1+\tau} = d^\tau = 0$. Hence, the agents do not attack at all. This shows that when there are disaster alerts in place right after every date when there is a risk of panic, then the agents never panic.

In the following Section, we consider an application to capital outflow. The problem is slightly different from our benchmark setting. Nevertheless, we show that a practical policy can be designed using the main insight.

4 Application: Taxing capital flight

Government in the emerging economies often imposes tax on capital flight. Clearly a tax would discourage the investors from exiting the market. However, this could discourage the investors to invest in the first place (See Mathevet and Steiner (2013)). We make a more subtle argument. A tax on capital flight makes waiting costly, i.e., if an investor exits then he should exit earlier than later. Then, it follows from the main insight from this paper that when the investors learn that it is not a disaster, they are assured that it will not becomes one, and consequently they do not panic.[6]

Consider a continuum of foreign investors who have invested in an emerging economy. As in our benchmark set up, suppose that a shock hits the economy, and the investors have noisy signals about the fundamental. They decide whether to exit (attack) or stay (not attack). They can exit at any time $t \in [0, T]$. Exit is irreversible. However, unlike in our benchmark setting, the agents receive flow payoff from staying. If an investor has not exited by time $t \in [0, T]$, he receives a flow payoff at time t depending on the underlying

[6]Although we do not model ex-ante investment, since panic is eliminated, such a tax will not discourage ex-ante investment.

fundamental θ and aggregate exit so far N_t as follows.

$$\tilde{r}(\theta, N_t) = \begin{cases} \overline{r} & \text{if } R(\theta, N_t) \geq 0 \\ \underline{r} & \text{if } R(\theta, N_t) < 0. \end{cases}$$

If $R(\theta, N_T) \geq 0$ (the regime survives), then the investors who does not exit will earn a flow payoff $\overline{r}$ forever. But if $R(\theta, N_t)$ becomes less tha 0 at some t, then the investor who does not exit will start getting a flow payoff $\underline{r}$ from time t onwards. Thus, if the investor stays, then his payoff is

$$v(\theta, (N_t)) = \int_{t=0}^{T} e^{-\beta t} \tilde{r}(\theta, N_t) dt + \int_{T}^{\infty} e^{-\beta t} \tilde{r}(\theta, N_T) dt,$$

where $\beta > 0$ is the discount rate. On the other hand, if an investor exits at some date t, then he switches to a safe investment project, which yields a fixed flow return of $r > 0$, where

$$\overline{r} > r > \underline{r}.$$

However, the investor has to pay a tax at a rate $\mu \in (0, 1)$ on the flow payoff the investor has earned until time t. Thus, his payoff from withdrawing at time $t_i \in [0, T]$ is

$$u(\theta, t_i, (N_t)) \equiv \int_{t=0}^{t_i} e^{-\beta t} (1 - \mu) \tilde{r}(\theta, N_t) dt + \int_{t_i}^{\infty} e^{-\beta t} r dt.$$

Note that unlike in the benchmark set up, the regime can change at any time t, rather than only at T. The flow payoff acts as endogenous disaster alert – whenever the flow payoff from keeping the investment at the emerging market becomes $\underline{r}$, the agents learn that the regime has changed. Also under no tax ($\mu = 0$), unlike in our benchmark set up, delayed attack may not be costly since the investor could earn $\overline{r}$ rather than r by exiting later.

Assumption 2 *(1)* $\mu > 1 - \frac{r}{\overline{r}}$ *and (2)* $T < \frac{1}{\beta} \ln(1 + \frac{r - \underline{r}}{\overline{r}})$.

The first restriction in Assumption 2 ensures that the tax rate is high enough such that $(1 - \mu)\overline{r} < r$. This implies that u is decreasing in t, i.e., if an investor exits, he should exit as early as possible. The second restriction in Assumption 2 ensures that the time window is sufficiently small. Otherwise, an investor may accumulate significant flow payoff over time and may not want to exit because of the significant exit tax.

Corollary 1 *In the capital outflow game, under Assumption 2, the investors do not panic. That is, $\Theta^P = \emptyset$.*

Proof. See Appendix. ■

After seeing $\tilde{r}(\theta, N_t) = \underline{r}$, an investor will exit right the next instance. But importantly, since a delayed exit is costly (Assumption 2), any agent who did not exit early would only exit later when $\tilde{r}(\theta, N_t) = \underline{r}$. In other words, the option value argument holds here. Since, the disaster alert is continuously in pace, it follows from Theorem 2 that all agents who believe that the flow payoff in future can be $\overline{r}$ with positive probability regardless of what others do (Assumption 1) would never exit unless $\tilde{r} = \underline{r}$ is realized. This completely eliminates the panic.

This result may sound surprising. However, once we understand the role of disaster alert policy in an endogenous move coordination game, it is not hard to see why investors do not panic. The shock that hits the emerging market may not be severe and exit is not warranted. However, if the investors believe that other investors will exit, they will exit as well. This generates panic. However, in an endogenous move game, when the realization of the flow payoff at each date within the short time window tells them whether it is dominant to exit or not, they would prefer to wait and then exit if the economic environment is proved to be really bad. This eliminates panic.

5 Discussion

This paper considers a canonical regime change game where agents' private signals could be arbitrarily correlated, and the principal does not have access to these private signals. She adopts a simple policy – a timely disaster alert, and she does not need ex-ante commitment to enforce such a policy. Yet, surprisingly, the policy completely eliminates panic. Even when new information arrives over time, a timely alert each time a new information arrives stops the agents from panicking. The readers may wonder that perhaps the result depends "too much" on the assumption of rationality, and if a rational agent fears that some agents will make mistakes, he may still panic.

5.1 Rationalizability

We use extensive form raionalizability to argue that an agent who has waited for the disaster alert will not attack after the alert is not triggered (Lemma 1).[7] This removes the strategic uncertainty after the alert and we leverage this to show that the agents will wait for the alert. In fact, common knowledge of rationality is not essential. As long as an agents is rational Lemma 1 holds true, and as long as an agent believes that others are rational Lemma 2 holds true. However, if there is a small chance η that even after the disaster alert is not triggered, the agents who have waited will behave irrationally, and attack, then the agents may not want to wait for such alert. Below we argue that for small η, a timely disaster alert will eliminate panic. For simplicity, let us assume $g(\theta, N_T) = g$ and $l(\theta, N_T) = l$.

Proposition 1 *Suppose that there is η fraction of agents who are irrational in the sense that they may attack after the alert is not triggered even though they have waited for the alert (play $\mathcal{W}'$). If $\eta < \frac{g-1}{g-l}$, then there exists a $\hat{\tau}$ such that Γ^τ with $\tau < \hat{\tau}$ eliminates panic.*

Proof. Consider an agent with signal s_i. The net payoff from playing $\mathcal{W}$ as compared to $\mathcal{A}$ is

$$D(\Gamma^\tau, s_i) = \mathbb{P}(d^\tau = 1|s_i)(u(\tau) - u(0)) + \mathbb{P}(d^\tau = 0|s_i)((1-\eta)(g-1) + \eta(l-1)).$$

Using Lipschitz continuity of $u(t)$ and Assumption 1, we have

$$D(\Gamma^\tau, s_i) \geq -(1-\varepsilon)K\tau + \varepsilon((1-\eta)(g-1) + \eta(l-1)).$$

If $\eta < \frac{g-1}{g-l}$, then a timely disaster alert policy Γ^τ where $\tau < \frac{\varepsilon}{1-\varepsilon}\frac{(1-\eta)(g-1)-\eta(1-l)}{K}$ eliminates panic. ■

This shows that if it is common knowledge that the agents are mostly rational, a timely disaster alert will eliminate panic. Next, we revisit the Doubt assumption.

[7] Ben-Porath and Dekel (1992) (also see Kohlberg and Mertens (1986) and Van Damme (1989)) make a similar forward induction argument that if an agent has a better outside option, then by sacrificing this option, he can send a signal to his opponent about his future action. However, if both agents have such options, then Ben-Porath and Dekel (1992) also argue that "simultaneous singling need not select the mutually preferred outcome."

5.2 (Un)necessary Doubt

We show that if agents have "doubt" ($\mathbb{P}(\theta \in \Theta^U | s_i) > 0$), then regardless of the information structure a timely disaster alert can eliminate panic. This assumption is stronger than saying that there is a upper dominance region. It says that regardless of the private signal an agent believes that θ could be in the upper dominance region. This assumption can be weakened depending on the heterogeneity of agent's beliefs. Consider the following simple examples.

Example 1 *Nature draws $\theta \in \Theta = [-1, 2]$ from uniform distribution, and the agents receive independent private signals $s_i \in \{l, m, h\}$ according to the following conditional distribution*

$f_i(s_i\|\theta)$	l	m	h
$\theta \in [-1, 0)$	p	$\frac{1}{2}(1-p)$	$\frac{1}{2}(1-p)$
$\theta \in [0, 1)$	$\frac{1}{2}(1-p)$	p	$\frac{1}{2}(1-p)$
$\theta \in [1, 2]$	$\frac{1}{2}(1-p)$	$\frac{1}{2}(1-p)$	p

Example 2 *Nature draws $\theta \in \Theta = [-1, 2]$ from uniform distribution, and the agents receive independent private signals $s_i \in \{l, m, h\}$ according to the following conditional distribution*

$f_i(s_i\|\theta)$	l	m	h
$\theta \in [-1, 0)$	p	$(1-p)$	0
$\theta \in [0, 1)$	$\frac{1}{2}(1-p)$	p	$\frac{1}{2}(1-p)$
$\theta \in [1, 2]$	0	$(1-p)$	p

Suppose that $p \in (\frac{1}{2}, 1 - 2\varepsilon)$. Under example 1, an agent always assigns probability at least ε to $\theta \in \Theta^U = [1, 2]$ regardless of her private information. Thus, assumption 1 holds true. Nevertheless, it is a restrictive assumption and may not hold true in general. In example 2, $f(s_i|\theta)$ does not have full support. This assumption is violated because the agent who receives $s_i = l$ does not believe that θ can be greater than 1.

Consider the information structure in example 2. If an agent receives the signal l, then the agent believes that $\theta < 1$. Then, this agent knows that attacking immediately is not a mistake if others also attack immediately. Suppose that all agents receive the public signal $s = l$. Then, even under any timely disclosure policy Γ^τ, a possible equilibrium outcome is that all the agents attack immediately. This shows that under the general information structure, assumption 1 is indeed a necessary and sufficient condition for our main result.

However, this does not mean that a given information structure must satisfy Assumption 1 for the result to be true. Let us go back to the simple example 2, but now suppose that the signals are not public. In particular, let us suppose that the signals are conditionally independent.

First, note that the agents who receive $s_i = m$ or h believes that $P(\theta \geq 1|s_i) > 0$ and hence will not attack immediately (Lemma 2). Now consider the agent who receives signal $s_i = l$. She believes $P(\theta \geq 1|l) = 0$, but since the signals are not public, she believes that others may have received signal $s_{-i} = m$ or h, and not attacking. Hence, the maximum attack at time 0 is from the agents who have received $s_{-i} = l$ and the disaster alert will never be triggered if θ is greater than the fraction of those agents. Thus,

$$\begin{aligned}\mathbb{P}(d^\tau = 0|l) &\geq \mathbb{P}(\theta \geq \mathbb{P}(s_{-i} = l)|l) = \mathbb{P}(\theta \geq \mathbb{P}(s_{-i} = l)|\theta \in [0,1))\mathbb{P}(\theta \in [0,1)|l) \\ &= P(\theta \geq \frac{1}{2}(1-p)|\theta \in [0,1)) \times \frac{1-p}{1+p} = \frac{1-p}{2}.\end{aligned}$$

This shows that even the agent who receives $s_i = l$ believes that there is a strictly positive probability that the disaster alert will not be triggered. That means attacking immediately (strategy $\mathcal{A}$) might be a mistake and thus, when the alert is set in a timely manner, the strategy of wait and see ($\mathcal{W}$) is the dominant one.

Now suppose that the information structure is as follows : with probability α, $s_j = s_i$ and with probability $(1-\alpha)$ the s_j is a conditionally independent signal. If $\alpha = 1$, then this captures the public information case, and if $\alpha = 0$, this captures the conditionally independent signal case. $\alpha \in (0,1)$ captures the heterogeneity in agents' beliefs. Then, for any given $\alpha < 1$,

$$\mathbb{P}(d^\tau = 0|l) \geq (1-\alpha)\frac{1-p}{2}.$$

This shows that if α is away from 1, i.e., there is enough heterogeneity in agents' beliefs, the principal can eliminate panic. In this sense, the doubt assumption is unnecessary.

Next, we discuss two essential features of our model – (1) continuity of $u(.)$ and (2) irreversibility of attack. We argue that in the absence of these two assumptions, a timely disaster alert may not eliminate panic. To show this formally, we construct equilibrium with panic. To explicitly chacterize this limitation, we specilize to the following standard global game environment. Agents have common prior about the underlying fundamental θ is distributed according to $U[\underline{\theta}, \bar{\theta}]$. In addition, agent $i \in [0,1]$ receives noisy private signal, denoted by $s_i = \theta + \sigma\epsilon_i$, where the error terms ϵ_i are conditionally independent and identically distributed with zero mean. Let $F : [-1/2, 1/2] \to [0, 1]$ be the distribution and f be the density of the error. $\sigma > 0$ scales the random noise ϵ_i. We assume that $\underline{\theta} \leq -\sigma$ and $\bar{\theta} > 1 + \sigma$.[8] Note that this information environment violates Assumption 1. For example, an agent who receives a signal $s_i < 1 - \frac{1}{2}\sigma$, believes that $\mathbb{P}(\theta \geq 1|s_i) = 0$.

5.3 Limits to timely persuasion

Lipschitz continuity of the delay cost ensures that the cost of waiting for the disaster alert can be made arbitrarily small. This is important for eliminating panic. However, in some applications, this assumption may not hold and this limits the extend to which panic can be reduced through timely disaster alert. For simplicity, we assume $g(\theta, N_T) = g$ and $l(\theta, N_T) = l$. However, unlike in our benchmark setting, we allow for discontinuity in the payoff when the agent attacks. In particular,

$$u(t, \theta, N_t) = \begin{cases} u_0(t) \text{ if } R(\theta, N_t) \geq 0 \\ u_1(t) \text{ if } R(\theta, N_t) < 0, \end{cases}$$

where both $u_0(t)$ and $u_1(t)$ are continuous and decreasing in t, and $u_0(t) \geq u_1(t)$. Suppose that at time t, an agent waits for a small time τ. Suppose that the other agents attack between $[t + t + \tau]$, resulting in $R(\theta, N_{t+\tau}) < 0$. Then, his payoff from attacking falls from $u_0(t)$ to $u_1(t + \tau)$. If $u_0(t) > u_1(t)$, then the cost of waiting for a small time may not be arbitrarily small. We maintain the assumption that even at the last minute, if the agent learns that $R(\theta, N_{T-dt}) < 0$, attacking is the dominant acton, i.e., $u_0(T) \geq u_1(T) \geq l$.

Proposition 2 *Under the disclosure policy Γ^τ, the regime does not survive because of panic when $\theta \in \Theta^P(\Gamma^\tau) = [0, \hat{\theta}^\tau]$, where*

$$\hat{\theta}^\tau = \frac{1}{1 + \frac{g-1}{1-u_1(\tau)}}.$$

Consider the two extreme cases. If there is no discontinuity (as in our benchmark model), i.e., $u_1(\tau) \to u_0(0) = 1$ as $\tau \to 0$, then $\hat{\theta}^\tau \to 0$. Thus, a timely disaster alert can eliminate panic. Now suppose $u_1(\tau) \to l$ as $\tau \to 0$, then $\hat{\theta}^\tau \to \frac{1-l}{g-l}$. This limit is the same

[8] This implies that for any $\theta \in [\underline{\theta}, \bar{\theta}]$, the probability of receiving private signal s is distributed via $F\left((s - \theta)/\sigma\right)$. Given the uniform prior, for any $s \in [\underline{\theta} - \sigma/2, \bar{\theta} + \sigma/2]$, θ is distributed via $1 - F\left((s - \theta)/\sigma\right)$.

threshold cutoff that arises even without any disaster alert (see Morris and Shin (2003)). Thus, a disaster alert is completely ineffective.

In general, the timely disaster alert is more effective when $\lim_{\tau\to 0} u_1(\tau)$ is higher. To see the intuition, consider the case when there is no disaster alert. Then, an agent either attack immediately, or does not attack at all. It follows from standard global game argument that there is a unique equilibrium where the agents follow a cutoff strategy – attack immediately if $s_i < \hat{s}$, otherwise do not attack. Accordingly, the regime survives if and only if $\theta \geq \hat{\theta}$, where

$$\hat{\theta} = \frac{1}{1 + \frac{g-1}{1-l}}.$$

Now suppose there is a disaster alert. If the agents follow the same strategy: attack if and only if $s_i < \hat{s}$, then the alert will be triggered if and only if $\theta < \hat{\theta}$. It follows from Lemma 1 that the agent who has not attacked, will attack if and only if the alert is triggered. If $u_1(\tau) = l$, then after the alert is triggered, it is too late to work on it, i.e., the agent gets the same if he attacks or not. Thus, from an ex-ante perspective, the strategy to not attack immediately, gives the same payoff regardless of whether there is a disaster alert or not. So, the same equilibrium remains, and disaster alert is completely ineffective. However, if $u_1(\tau) < l$, then the disaster alert is valuable. This will incentivize the agents to wait for the alert. The higher the $\lim_{\tau\to 0} u_1(\tau)$, the higher the incentive, and thus, lower the resulting $\hat{\theta}^{\tau}$.

A critical assumption in our model is that, within the time window, agents can save some money by attacking a regime after learning that the regime is doomed to fail. This assumption can be conflicting with the nature of financial panics, or the first mover advantage. Panic in financial market may happen due to the sequential service constraint, i.e., if sufficiently many investors already lined up to make redemptions, then there will be no chance for later withdrawers to get money back. In other words, the fact that this assumption does not hold true might be exactly the reason for financial panic in the first place. This will make the policy maker's job more difficult since she has to make sure that even when investor make withdrawal after the bank is doomed to fail, the payoff of such a late withdrawal will be close to an early withdrawal. Or more precisely, the difference is continuous in the timing of action without discrete jumps. For our information disclosure policy to work, the policy maker needs to provide payoff guarantee for redemption so that a later redemption will be treated close to that for early redemption.[9] According to our

[9]Note that in reality, even with full guarantee, without information disclosure policy, panics can still

theory, the information disclosure policy will be effective as long as the policy maker can guarantee that the redemptions happened shortly after the time of disclosure. There is no need to provide a life-time guarantee as the deposit insurance.

5.4 Reversibility and Panic

One may think that if attack is a reversible action, then it will make the result even stronger. After all, if the agent who had left can come back, it will reassure the agents who stayed. However, if attack is a reversible, an agent can take the following strategy ($\mathcal{W}^c$): attack right away, and reverse the action only if the alert is not triggered. This could trigger the alarm although it was not warranted.

Consider the following simple example: The investors decide whether to stay in (not attack) or stay out (attack). The payoffs depends on the duration of each action. If he stays in for a duration of t, then he gets a flow return r for the duration that he stays out $(T-t)$, and a higher return (g) or lower return (l), depending on the fate of the regime, weighted by the duration of stay, i.e., his payoff is $r(T-t) + t \cdot (g\mathbf{1}(\theta \geq N_T) + l\mathbf{1}(\theta < N_T))$, where $g > r > l$. Then, the strategy $\mathcal{W}^c$ gives expected payoff $\mathbb{P}(d^\tau = 1|s_i)rT + \mathbb{P}(d^\tau = 0|s_i)(r\tau + (T-\tau)g)$. On the other hand, the strategy $\mathcal{W}$ gives expected payoff $\mathbb{P}(d^\tau = 1|s_i)((T-\tau)r + \tau l) + \mathbb{P}(d^\tau = 0|s_i)(Tg)$. Therefore, the net payoff from playing $\mathcal{W}^c$ as opposed to $\mathcal{W}$ is

$$\mathbb{P}(d^\tau = 1|s_i)(\tau(r-l)) + \mathbb{P}(d^\tau = 0|s_i)(\tau(r-g)).$$

Therefore, if an agent believes that the alert is sufficiently likely to be triggered, he will play $\mathcal{W}^c$ rather than $\mathcal{W}$. When attacking is reversible, agents (especially with low signals) may decide to stay out unless the alert says otherwise. This strategy can trigger the alert (especially when θ is low). One can easily construct an equilibrium to show that panic is not eliminated.

Proposition 3 *When attacking is reversible, there is an equilibrium in which the agents with $s_i < \hat{s}$ plays $\mathcal{W}^c$, and the ones with $s_i \geq \hat{s}$ plays $\mathcal{W}$.*

---

happen because the opportunity cost of a later withdrawal. In other words, even full guarantee cannot completely remove the first mover advantage. Such costs includes the time of waiting to get full payment from the government agency, the loss from continuous capital gain from moving the capital to other investment opportunities, etc.

Proof. Suppose the agent play such a cutoff strategy. Then, for any θ, the aggregate attack at time 0 is $F((\hat{s}-\theta)/\sigma)$. Clearly this is decreasing in θ. Therefore, there is a $\hat{\theta}$ such that the alert is triggered ($d^\tau = 1$) if and only if $\theta < \hat{\theta}$, where $F((\hat{s}-\hat{\theta})/\sigma) = \hat{\theta}$. An agent with private signal s_i believes that the net expected payoff from playing $\mathcal{W}^c$ compared to $\mathcal{W}$ is

$$\left(1 - F\left(\frac{s_i - \hat{\theta}}{\sigma}\right)\right)(\tau(r-l)) + F\left(\frac{s_i - \hat{\theta}}{\sigma}\right)(\tau(r-g)).$$

Clearly this is decreasing in s_i. Consider the marginal agent with signal $s_i = \hat{s}$. He must be indifferent between playing $\mathcal{W}^c$ and $\mathcal{W}$. Substituting $F((\hat{s}-\hat{\theta})/\sigma) = \hat{\theta}$, we get that the net benefit of the marginal agent is

$$\tau(r-l) - \hat{\theta}\tau((g-r)+(r-l)) = 0 \implies \hat{\theta} = \frac{1}{1+\frac{g-r}{r-l}}.$$

Thus, $\hat{s} = \hat{\theta} + \frac{1}{\sigma}F^{-1}(\hat{\theta})$ constitutes an equilibrium, and in this equilibrium a regime with fundamental below $\hat{\theta}$ but above 0 will not survive because of panic-based runs. ■

6 Concluding Remarks

The U.S. government countered the recent financial panic in the great recession in 2008 with various measures, including liquidity injection and debt guarantees. Stress testing, as the only measure which involves information production and disclosure, was introduced during the crisis to avert the financial panics. Panic in financial market happen and evolve in a dynamic manner. Investors and financial institutions decide on when to make withdrawal or redemption. For that reason, the policy maker who wants to quell the financial panic by disclosing information should take the dynamic feature of decision making into consideration. However, the current literature on understanding the optimal information disclosure on averting coordination failure investigate such question in a static setting. In contrast, we build a dynamic regime change game where attack is irreversible, delay is costly and the cost of delay is continuous. We show that a disaster alert (weakest stress test) can eliminate panic if it is set in a timely manner.

Indeed, the first practice of the stress testing, i.e., Supervisory Capital Assessment Program (SCAP), is conducted in a timely manner. The plan for stress testing was announced on Feb 10, 2009. The white paper describing the procedures employed in SCAP was released on Apr 24, 2009 and the results of SCAP were disclosed on May 7, 2009. This timely disclosure of stress test results also proved to be successful. Peristiani, Morgan and Savino (2010) document evidence to show that stress tests helped quell the financial panic by producing vital information about banks. Bernanke (2013) states *"Supervisors' public disclosure of the stress tests results helped restore confidence in the banking system..."* Gorton (2015) states that the tests results were viewed as credible and the stress tests are widely viewed as a success.

References

Angeletos, George-Marios, and Chen Lian. 2017. "Incomplete Information in Macroeconomics: Accommodating Frictions in Coordination." *Handbook of Macroeconomics*, 2: 1065–1240.

Angeletos, George-Marios, Christian Hellwig, and Alessandro Pavan. 2007. "Dynamic Global Games of Regime Change: Learning, Multiplicity, and the Timing of Attacks." *Econometrica*, 75(3): 711–756.

Basak, Deepal, and Zhen Zhou. 2019. "Diffusing Coordination Risk." *American Economic Review (forthcoming)*.

Ben-Porath, Elchanan, and Eddie Dekel. 1992. "Signaling Future Actions and Potential for Sacrifice." *Journal of Economic Theory*, 57: 36–51.

Bergemann, Dirk, and Stephen Morris. 2016. "Bayes correlated equilibrium and the comparison of information structures in games." *Theoretical Economics*, 11.2: 487–522.

Bernanke, Ben S. 2013. "Stress Testing Banks: What Have We Learned?"

Chamley, Christophe, and Douglas Gale. 1994. "Information revelation and strategic delay in a model of investment." *Econometrica: Journal of the Econometric Society*, 1065–1085.

Dasgupta, Amil. 2007. "Coordination and Delay in Global Games." *Journal of Economic Theory*, 134(1): 195–225.

Dasgupta, Amil, Jakub Steiner, and Colin Stewart. 2012. "Dynamic Coordination with Individual Learning." *Games and Economic Behavior*, 74(1): 83–101.

Doval, Laura, and Jeffrey Ely. 2019. "Sequential information design." *Working Paper*.

Ely, Jeffrey C. 2017. "Beeps." *The American Economic Review*, 107(1): 31–53.

Gale, Douglas. 1995. "Dynamic Coordination Games." *Economic theory*, 5(1): 1–18.

Goldstein, Itay, and Chong Huang. 2016. "Bayesian Persuasion in Coordination Games." *The American Economic Review*, 106(5): 592–596.

Gorton, Gary. 2015. "Stress for success: A review of Timothy Geithner's financial crisis memoir." *Journal of Economic Literature*, 53(4): 975–95.

Gul, Faruk, and Russell Lundholm. 1995. "Endogenous timing and the clustering of agents' decisions." *Journal of political Economy*, , (103.5): 1039–1066.

Inostroza, Nicolas, and Alessandro Pavan. 2017. "Persuasion in Global Games with Application to Stress Testing." http://faculty.wcas.northwestern.edu/~apa522/persuasion-GG.pdf.

Kamenica, Emir. 2018. "Bayesian persuasion and information design." *Annual Review of Economics*.

Kamenica, Emir, and Matthew Gentzkow. 2011. "Bayesian Persuasion." *The American Economic Review*, 101(6): 2590–2615.

Kohlberg, Elon, and Jean-Francois Mertens. 1986. "On the Strategic Stability of Equilibria." *Econometrica*, 54(5): 1003–1037.

Makris, Miltiadis, and Ludovic Renou. 2018. "Information design in multi-stage games." *Queen Mary, University of London School of Economics and Finance Working Paper 861*.

Mathevet, Laurent, and Jakub Steiner. 2013. "Tractable dynamic global games and applications." *Journal of Economic Theory*, 148(6): 2583–2619.

Morris, Stephen, and Hyun Song Shin. 2003. "Global Games: Theory and Applications." In *Advances in Economics and Econometrics (Proceeding of the Eighth World Congress of the Econometric Society)*. , ed. Dewatripont, Hansen and Turnovsky. Cambridge University Press.

Pearce, David G. 1984. "Rationalizable strategic behavior and the problem of perfection." *Econometrica: Journal of the Econometric Society*, 1029–1050.

Peristiani, Stavros, Donald P Morgan, and Vanessa Savino. 2010. "The information value of the stress test and bank opacity." *FRB of New York Staff Report*, , (460).

Salcedo, Bruno. 2017. "Interdependent choices." *Working Paper*.

Szkup, Michal. 2017. "Multiplier effect and comparative statics in global games of regime change." https://econ.sites.olt.ubc.ca/files/2017/11/pdf_szkup_comparative_statics.pdf.

Van Damme, Eric. 1989. "Stable Equilbria and Forward Induction." *Journal of Economic Theory*, 48: 476–496.

Appendix

Proof of Theorem 2

Lemma 3 *Under the disclosure policy Γ^τ, for any noisy signal s_i^0, s_i^1, the only rationalizable strategies are*

$\mathcal{A}$: *attack at time* 0*, i.e.,* $a_{it}(s_i^0) = 1$ *for all* $t \in [0, T]$*,*

$\mathcal{W}^1$: *wait for the first disaster alert but not for the next one, i.e.,* $a_{it}(s_i^0) = 0$ *for all* $t \in [0, \tau + dt)$*, then play* $a_{it}(s_i^0, d^\tau = 1) = 1$ *for* $t \in [\tau + dt, T]$*,* $a_{it}(s_i^0, d^\tau = 0) = 0$ *for* $t \in [\tau + dt, t_1 + dt)$*, and* $a_{it}(s_i^0, d^\tau = 0, s_i^1) = 1$ *for* $t \in [t_1 + dt, T]$*.*

$\mathcal{W}^2$: *wait for both disaster alerts, i.e.,* $a_{it}(s_i^0) = 0$ *for all* $t \in [0, \tau + dt)$*, then play* $a_{it}(s_i^0, d^\tau = 1) = 1$ *for* $t \in [\tau + dt, T]$*,* $a_{it}(s_i^0, d^\tau = 0) = 0$ *for* $t \in [\tau + dt, t_1 + \tau + dt)$*, and* $a_{it}(s_i^0, d^\tau = 0, s_i^1, d^{t_1+\tau}) = d^{t_1+\tau}$ *for* $t \in [t_1 + \tau + dt, T]$*.*

Proof. Any exit happening at a time when there is no new arrival of information is strictly dominated by exiting at an earlier time with the same information set. Then, applying the same logic as in Lemma 1, when $d^\tau = 0$ (or $d^{t_1+\tau} = 0$), regardless of s_i^0 (and s_i^1), an agent would not exit. Otherwise, the new information disclosed at τ (or $t_1 + \tau$) does not have any value and he should exit early without waiting for the new information. ■

Lemma 4 *Given Assumption 1, under the disclosure policy Γ^τ, where $\tau < \hat{\tau}$, for any signal realization s_i^0, s_i^1, the only rationalizable strategy for an agent with signal s_i^0, s_i^1 is $\mathcal{W}^2$, i.e., wait and follow the principal's recommendation.*

Proof. It follows from Lemma 2 that under the disclosure policy Γ^τ with $\tau < \hat{\tau}$, an agent will rather play $\mathcal{W}^2$ than $\mathcal{W}^1$. The expected payoff from playing $\mathcal{A}$ is 1. While the expected payoff from playing $\mathcal{W}^2$ is

$$\mathbb{P}(d^\tau = 1 | s_i^0) e^{-r\tau} + \mathbb{P}(d^\tau = 0, d^{t_1+\tau} = 1 | s_i^0) e^{-r(t+\tau)} + \mathbb{P}(d^\tau = 0, d^{t_1+\tau} = 0 | s_i^0) g.$$

Since strategy $\mathcal{W}^1$ is dominated by $\mathcal{W}^2$, if the disaster alert is not triggered at τ, it will not be triggered at $t_1 + \tau$, i.e., $\mathbb{P}(d^\tau = 0, d^{t_1+\tau} = 1 | s_i) = 0$. Indeed, $d^\tau = d^{t_1+\tau}$ no matter agents play $\mathcal{A}$ or $\mathcal{W}^2$. Therefore, using the same argument from Lemma 2, we can say that $\mathcal{W}^2$ strictly dominates $\mathcal{A}$ when $\tau < \hat{\tau}$. ■

Therefore, under the disclosure policy Γ^{τ} with $\tau < \hat{\tau}$, regardless of their signals, agents never attacks a regime with $\theta \geq 0$. Following the same argument as in the proof of Theorem 1, $\Theta^P(\Gamma^{\tau}) = \emptyset$. □

Proof of Corollary 1 First we show that if $\tilde{r}(\theta, N_\tau) = \underline{r}$, then exiting right away is the dominant strategy. Let $(N_t)_{t<\tau}$ be the history of attack until time τ. The net payoff from waiting as compared to exiting right away is

$$\begin{aligned} v(\theta, (N_t)_{t<\tau}) - u(\tau, \theta, N_\tau) &= \int_0^\tau e^{-\beta t}\mu\tilde{r}(\theta, N_t)dt - e^{-\beta\tau}\frac{r-\underline{r}}{\beta} \\ &\leq \frac{1}{\beta}\Bigg[(1-e^{-\beta\tau})\mu\overline{r} - e^{-\beta\tau}(r-\underline{r})\Bigg] \end{aligned}$$

The first part is the amount the investor saves because he does not have to pay the capital flight tax. The second part is the loss of future profit. This simplifies to

$$v(\theta, (N_t)_{t<\tau}) - u(\tau, \theta, N_\tau) = \frac{1}{\beta}\Bigg[\mu\overline{r} - e^{-\beta\tau}(\mu\overline{r} + r - \underline{r})\Bigg].$$

It follows from the first part of Assumption 2 that $(1-\mu)\overline{r} < r$, i.e., $\mu\overline{r} + r > \overline{r} > \underline{r}$, i.e.,

$$\mu\overline{r} + r - \underline{r} > 0.$$

This means that if T is large enough, then this net payoff can be positive. The intuition is simple. If T is large, then the agent may accumulate a lot in flow payoff over time. Thus, the capital flight tax may become higher than the future loss the investor has to incur if he does not exit. Then, an agent may want to wait anyway even if the alert is triggered. The second part of Assumption 2 ensures that this is not the case. If $T < \frac{1}{\beta}\log(1 + \frac{r-\underline{r}}{\overline{r}})$, then for any $\tau \leq T$, the net payoff from waiting as compared to exiting becomes negative, i.e., $v(\theta, (N_t)_{t<\tau}) - u(\tau, \theta, N_\tau) < 0$.

It follows from the first part of Assumption 2 that $u(\theta, t, N)$ is decreasing in t. Therefore, the option value argument holds true (Lemma 1). Since the disaster alert is continuously in place, the rest of the proof is essentially the same as in Theorem 2. □

Proof of Proposition 2 The regime survives only when the fundamental θ is no lower than the mass of agents who choose $\mathcal{A}$, or N_0. The payoff for taking $\mathcal{A}$ is 1, while the payoff for taking $\mathcal{W}$ is g when $d^\tau = 1$ (or equivalently $\theta \geq N_0$) and $u_1(\tau)$ when $d^\tau = 0$ (or equivalently $\theta < N_0$).

For agent i who has received private information $s_i < -\frac{1}{2}\sigma$, he knows that $\theta < 0$ and thus $d^\tau = 1$ for sure and thus he will take $\mathcal{A}$. Hence, the dominance region of $\mathcal{A}$ is $[\underline{\theta} - \frac{1}{2}\sigma, -\frac{1}{2}\sigma)$. Similarly, the dominance region of $\mathcal{W}$ is $[1 + \frac{1}{2}\sigma, \bar{\theta} + \frac{1}{2}\sigma]$. The rest of the proof follows the iterated elimination arguments as in global games.

Consider an agent with signal s_i. If he believes that

$$P(d^\tau = 0|s_i) \geq \frac{1}{1 + \frac{g-1}{1-u_1(\tau)}},$$

then he will play $\mathcal{W}$ rather than $\mathcal{A}$. Let us define

$$p := \frac{1}{1 + \frac{g-1}{1-u_1(\tau)}}.$$

Define $\hat{s}_1$ such that $\mathbb{P}(\theta \geq \hat{\theta}_1 \equiv 1|\hat{s}_1) = p$. This implies

$$\hat{s}_1 = 1 + \sigma F^{-1}(p). \tag{2}$$

Then, for any $s > \hat{s}_1$, an agent will play $\mathcal{W}$ rather than $\mathcal{A}$. Define $\hat{\theta}_2$ such that when $\theta \geq \hat{\theta}_2$, the maximum attack at time 0 (from all agents who receive information $s \leq \hat{s}_1$) would not trigger the disaster alert. Hence,

$$\mathbb{P}(s \leq \hat{s}_1|\hat{\theta}_2) = \hat{\theta}_2 \Longrightarrow \hat{\theta}_2 + \sigma F^{-1}(\hat{\theta}_2) = \hat{s}_1$$

Compared with Equation 2, we know that $p < \hat{\theta}_2 < \hat{\theta}_1 = 1$. Then, we can define $\hat{s}_2$ such that for all $s > \hat{s}_2$, agents know that the disaster alert will not be triggered with probability at least p independent of what others do, i.e.,

$$\mathbb{P}(\theta \geq \hat{\theta}_2|\hat{s}_2) = \epsilon \Longrightarrow \sigma F^{-1}(p) + \hat{\theta}_2 = \hat{s}_2$$

Hence, $\hat{s}_2 < \hat{s}_1$ since $\hat{\theta}_2 > p$. Following this iterative arguments, we can find two decreasing sequences $\{\hat{s}_n\}_{n=1}^{\infty}$ and $\{\hat{\theta}_n\}_{n=1}^{\infty}$ with limits

$$s^* = \lim_{n\to\infty} \hat{s}_n = p + \sigma F^{-1}(p) \text{ and } \theta^* = \lim_{n\to\infty} \hat{\theta}_n = p. \square$$