Timely Persuasion

Deepal Basak, Zhen Zhou*

First Version: Dec 2018

This Version: Apr 2019

Abstract

We consider a regime change game but allow the agents to attack within a short time window. Attack is irreversible and delayed attack is costly. There could be panic-based attacks, i.e., the agents attack thinking others will attack, even though it is not warranted. We propose a simple dynamic information disclosure policy, called "disaster alert", which at a given date publicly discloses whether the regime is doomed to fail. 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 signals, and thus, eliminates panic. If the time window is large, then the agents may receive additional information over time and panic may reappear. However, repeated timely disaster alerts stop agents from panicking. We build an example of FDI outflow from an emerging market, in which the disaster alerts manifest endogenously and dissuade the investors from exiting the market.

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

Key Words: *Information Design, Panics, Global Game*

---

*We thank Douglas Gale, Alessandro Pavan, Debraj Ray, Laura Veldkamp, Xingye Wu and seminar participants at 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 such as political turnover 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 and delayed attack is costly. This delay cost can be thought of as 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 who wants the regime to survive. She does not know the noisy signals each agent has received privately. But she has superior information than the agents in the following sense. Unlike the agents, she observes the actual fundamental and the history of attack perfectly. She commits to a dynamic information disclosure rule: At some date t, she will send a message to the agents based on the fundamental and the history until time t. We propose a simple policy, called the "disaster alert." A disaster alert at some date t is triggered 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. What does it mean when the disaster alert is not triggered? 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, the agents will follow the principal's recommendation. However, this does not mean that the 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. We show that, 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, irrespective of what others do. 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. We extend our model to a general binary action dynamic game involving strategic complementarity where one of the action is irreversible, and delaying this irreversible action is costly. Under fairly reasonable assumptions, we show that our insight is robust.

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 agents 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.

As long as the agents doubt that immediate attack could be a mistake, a timely disaster alert will persuade them to wait for the alert, and thus eliminate panic. 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, a given information structure may violate the *Doubt* assumption and the result may still hold. In Section 3, we provide some examples to discuss this.

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.

In practice, even if there is no principal who deliberately sets a disaster alert, it is possible that timely disaster alerts are naturally in place. For example, going back to FDI outflow, it is possible that the flow payoff reveals to the agents whether the regime is doomed to fail or not. We consider such an example and show that the observability of the flow payoff within the time window of action implies the disaster alert is continuously in place. Thus, even without any deliberate disclosure policy, there is no chance of panic in that flow payoff setting.

Related Literature This paper belongs to the recent growing literature on information design. See Bergemann and Morris (2013), Kamenica and Gentzkow (2011), Ely (2017) for recent development of this literature. Kamenica (2018) is an excellent survey of this literature. In this paper, an information designer (principal) tries to dissuade a mass of agents from attacking a regime. The agents are privately informed about the fundamental state. The principal knows the state but does not know the private signals that the agents have received. Two most closely related works are Inostroza and Pavan (2017) and Basak and Zhou (2018). Both these papers consider the same problem but in Inostroza and Pavan (2017) the agents move simultaneously, while in Basak and Zhou (2018) the agents move at an exogenously specified order. In contrast, in this paper, the agents get a short time window and endogenously decide when to attack. This endogenous delay gives a much stronger result as compared to Inostroza and Pavan (2017) and Basak and Zhou (2018). We discuss this relation in details in Section 3.

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. We extend this canonical regime change game to allow for endogenous delay in attack. This is similar to Gale (1995), Dasgupta (2007) and Dasgupta, Steiner and Stewart (2012). In Section 3, we consider the global game information structure as a special case, and formally establish the correction with the existing result. The global

[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.

game literature shows a unique cutoff equilibrium, and thus uniquely identifies the ex-ante probability of panic-based attacks. We show that a disaster alert reduces this chance of panic. The less the agents have to wait for the alert, the lower the probability of panic, and in the limit as the alert is almost immediate, there is no panic.

The rest of the paper is organised as follows. Section I describes the model. Section II describes how and when Disaster alert in a timely manner leads to no panic. Section III extends the result to a more general payoff structure than the one considered initially in the model. Section IV describes how panic can be avoided using a disaster alert in a timely manner when there is 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.

Payoff The payoffs are realized after time T when all the decisions have been made. The agents are identical. If an agent decided to attack, i..e, $a_{iT} = 1$ and the time of attack is t_i, then he gets

$$u(\theta, t_i, N_T) = e^{-rt_i}.$$

This means that if an agent attacks, he should attack as early as possible. On the other hand, if $a_{iT} = 0$, i.e., the agent does not attack at all, then he gets the payoff $v(\theta, N_T)$ depending on the fundamental state θ and the aggregate attack N_T until the end of the time window. If $\theta \geq N_T$, i.e., the fundamental is large enough to sustain the aggregate attack, then the regime succeeds and the agent who does not attack, receives a payoff g higher than 1. On the other hand, if $\theta < N_T$, then the regime fails and the agent who does not attack gets a payoff l, which is less than the payoff from any successful attack, i.e., $l < e^{-rT}$. Thus,

$$v(\theta, N_T) = \begin{cases} g & \text{if } \theta \geq N_T \\ l & \text{if } \theta < N_T. \end{cases}$$

This payoff specification resembles the one in a canonical regime change game (see Morris and Shin (2003)) with the only modification that the agents do not move simultaneously, rather they have a small time window $[0, T]$ to react. In the discussion section, we consider more general payoff structure.

Dominance Regions In this regime change game, if the underlying fundamental θ is very strong (or weak), or the shock is very mild (or severe), then the regime succeeds (or fail) regardless of the size of the attack. We refer to Θ^U (or Θ^L) as the upper (or lower) dominance region where not attacking (or attacking) is the dominant strategy. According to the above payoff specification, $\Theta^U = \Theta \cap [1, +\infty)$ and $\Theta^L = \Theta \cap (-\infty, 0)$. 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. 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.

Principal The principal has superior information than the agents. She can learn the underlying fundamental θ as well as the past actions N_t perfectly. However, the principal may not have access to the noisy information s_i each agent has. The principal's payoff only depends on whether the regime succeeds or not. She gets 1 if the regime succeeds and 0 if it fails.

Disclosure Policy The feature of endogenous move enables the principal to select the time of disclosure, and to make this information disclosure policy history dependent. Formally, for any $\tau \in [0, T]$, the principal can disclose some information to the agents based on θ and the 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 $\Gamma = (\pi, \mathcal{S})$ consists of the set of disclosed messages $\mathcal{S}$ and the disclosure rule $\pi : [0, T] \times \Theta \times [0, 1] \to \Delta(\mathcal{S}^{[0,1]})$.

Robust Design We use rationalizability as our solution concept. Given a disclosure policy Γ, let $R(\Gamma)$ be the set of all possible rationalizable strategy profiles $a \equiv (a_i(s_i))$. Define

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

Given any disclosure policy Γ, $\Theta^F(\Gamma)$ contains all states under which the regime could possibily fail. The principal's objective is

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

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 Γ to minimize the ex-ante chance that the regime fails.

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

2 Disaster Alert and No Panic

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

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

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., $\theta < N_T$, since N_t is weakly increasing in t. Hence, for agents who have not attacked, it is the dominant strategy to attack at time $\tau + dt$. That is why we call $d^\tau = 1$ a disaster alert. On the other hand, if the alert is not triggered, or $d^\tau = 0$, agents understand that $\theta \geq N_\tau$ 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.

Screening Property

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 *(Screening) 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 of all, it is not rational for agents to attack at any time other than 0 and $\tau + dt$. Secondly, given $d^\tau = 1$, agents would take the dominant action to 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 $e^{-(\tau+dt)} < 1$, and thus it is strictly dominated by $\mathcal{A}$, which guarantees a payoff of 1. ■

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. This argument is similar to Chamley and Gale (1994) and Gul and Lundholm (1995). Unlike this paper, the authors consider a social learning problem in which an agent can learn from others' actions, but such actions do not affect his payoff. [2]

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 he will not attack when the alert is not triggered. This explains the intuition behind Lemma 1.[3]

[2]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.

[3]It is worth pointing out that Lemma 1 is not restricted to a coordination game. In any endogenous move game, as long as there is costly delay in taking the irreversible action, a one-time public disclosure that partitions the state of the world into a dominant region of the irreversible action and a non-dominant region would induce the same screening property.

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 ($\theta \geq N_\tau$) implies the survival of the regime, i.e., $\theta \geq N_T$. 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.

In contrast in a static regime change game, in which the agents move simultaneously, the strategic complementarity cannot be removed when they know publicly that attacking is not the dominant strategy. Imagine that there is 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. However, there are many other possible equilibria (See Angeletos, Hellwig and Pavan (2007) and Basak and Zhou (2018)) in which the regime could fail because of panic-based attacks. The crucial difference is that under endogenous timing early attack works like a better outside option. Ben-Porath and Dekel (1992) (also see Kohlberg and Mertens (1986) and Van Damme (1989)) make a forward induction argument that if an agent has such an 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 signaling need not select the mutually preferred outcome."

It follows from the above Lemma that a disaster alert screens the agents into two categories: (1) They attack right away, or (2) they wait for the disaster alert and then follow the principal's recommendation. We refer to this as the *screening property*. Although a disaster alert policy Γ^τ takes away the strategic uncertainty once the game reaches time τ, there is still strategic uncertainty before time τ, i.e., the agents are uncertain about whether others would wait for the alert or attack right away.

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. [4] To better understand this assumption, 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. For most of this paper, we assume that the information structure satisfies assumption 1. We will relax this assumption in Section 3.

[4] 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 $\theta \in \Theta^U$.

2.1 No Panic

With the doubt assumption in place, the following Lemma shows that attacking at time 0 without waiting for the disclosure, or $\mathcal{A}$, is a strictly dominated strategy if the principal sets the disaster alert at a timely manner.

Lemma 2 *(Timely Persuasion) Under the disclosure policy Γ^τ, where*

$$\tau < \hat{\tau} = \frac{1}{r} ln \left(\frac{1-\varepsilon}{1-g\varepsilon} \right),$$

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.

Proof. For an agent who decided to wait for the disclosure, if $d^\tau = 1$, then he will attack at time $\tau + dt$ and will get $e^{-r(\tau+dt)}$. Otherwise, if $d^\tau = 0$, the agent will not attack. Given the regime survives in this case, he gets g. Hence, the expected payoff from playing $\mathcal{W}$ is

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

While, the expected payoff from playing $\mathcal{A}$, i.e, attacking immediately, is 1. Hence, the expected payoff difference from strategy $\mathcal{W}$ as compared to $\mathcal{A}$ is

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

Thus, $D(\Gamma^\tau, s_i) > 0$ whenever

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

Note that if $\theta \geq 1$, then the disaster alert cannot be triggered. Hence,

$$\mathbb{P}(d^\tau = 0|s_i) \geq \mathbb{P}(\theta \geq 1|s_i) > \varepsilon.$$

The last inequality comes from Assumption 1. Therefore, $D(\Gamma^\tau, s_i) > 0$ for any s_i if

$$\tau < -\frac{1}{r} ln \left(1 - \frac{\varepsilon(g-1)}{1-\varepsilon} \right) = \frac{1}{r} ln \left(\frac{1-\varepsilon}{1-g\varepsilon} \right).$$

■

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 $1 - e^{-r(\tau+dt)}$ 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 in the end. The benefit from this more informed choice is $g - 1$.

Assumption 1 guarantees that regardless of the signal s_i, an agent assigns positive probability to states under which the disaster alert will definitely not be triggered (or equivalently, 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.

Theorem 1 *(No Panic) Under the disclosure policy Γ^τ with $\tau < \hat{\tau}$, there is no panic, i.e.,*

$$\Theta^P(\Gamma^\tau) = \emptyset.$$

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

Theorem 1 follows immediately from Lemma 2. Under a timely disclosure policy, all agents, regardless of their noisy signal, would wait for the disclosure and follow the recommendation of the principal afterwards. Any regime that can survive without any attack will survive for sure. However, if $\theta \in \Theta^L$, the regime will fail and the principal cannot do anything about it. The timely disaster alert ($\tau < \hat{\tau}$) dissuades the agents from attacking whenever $\theta \notin \Theta^L$ and thus eliminates the chance of panic.

Hence, this simple public disclosure policy enables the principal to achieve her first best. Note that in designing such disclosure policy, the principal does not need to know the signal realization for each agent. This means that (1) even if the principal can observe the signals, she cannot do any better by conditioning the information disclosure on the realized signals, (2) the principal cannot increase her payoff by discriminatory disclosure, e.g., by sending additional noisy private signals to agents. In addition, the principal does not need ex ante commitment power to implement this policy.

The above result shows that "panic" is a fragile concept. A very simple information disclosure policy can eliminate panic. Is this surprising? In a game of strategic complementarity, an agent may attack if he believes that the other agents will attack, and this may lead to a panic even when it is not warranted. One may conjecture that there should a way to disclose information that will take away the strategic uncertainty and eliminate panic. However, the existing results prove otherwise. When the agents move simultaneously, Inostroza and Pavan (2017) find the optimal information disclosure policy, and this policy does not eliminate panic. But if we consider a small time window in which the agents can attack, while delayed attack is costly, then panic can be eliminated by a very simple public disclosure policy.

2.2 Arrival of New Information

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.

2.3 Endogenous Disaster Alert: An example

Consider a continuum of foreign investors who have invested in an emerging economy. Similar to Mathevet and Steiner (2013), one can model FDI outflow as a regime change game, in which attacking is equivalent to exiting the market. However, unlike our benchmark setting, the investors may receive flow payoff which could depend on the underlying fundamental θ and aggregate exit so far N_t. Thus, the flow payoff may act as endogenous disaster alert. We build a simple example to show that such endogenous disaster alert can dissuade the investors from panicking.

We will refer to the following game as FDI outflow game. The game is as before except the following modification. If an investor does not exit then at date t he receives a flow payoff

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

where $w(\theta, N)$ is a continuous and differentiable function with $\frac{dw}{d\theta} > 0$ and $\frac{dw}{dN} < 0$. For example, $w(\theta, N) = \theta - N$. Thus, if the investor decided not to exit, then his payoff is

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

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, if the investor i withdraws, the government will tax the flow payoff realized before the time of exiting at a rate $\mu \in (0, 1)$. Thus, his payoff from withdrawing at time $t_i \in [0, T]$ is

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

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

Proof. See Appendix. ■

Although there is no external disaster alert, the agents see the flow payoff $R(\theta, N_t)$. If $R(\theta, N_t) = \overline{R}$, the investor knows that $w(\theta, N_t) \geq 0$, and if $R(\theta, N_t) = \underline{R}$, the investor knows that $w(\theta, N_t) < 0$. Since $w(\theta, N)$ is decreasing in N, if $w(\theta, N_t) < 0$, then $w(\theta, N) < 0$ for any $N \geq N_t$. Since T is sufficiently small (Assumption 2), this implies that for $w(\theta, N_t) < 0$, it is dominant strategy for the investor to exit, i.e.,

$$u(\theta, t, N) > v(\theta, N), \forall N \geq N_t.$$

Thus, at any date t, the flow payoff sends a disaster alert $d^t = 1$, if it has become the dominant strategy to exit, and $d^t = 0$, otherwise. Formally,

$$d^t = \mathbf{1}\{(\theta, N_t) \in \{(\theta, N) | u(\theta, t, N') \geq v(\theta, N'), \forall N' \geq N\}\}.$$

In this sense the flow payoff acts as a *continuous endogenous disaster alert*.

After seeing $R(\theta, N_t) = \underline{R}$, 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 $R(\theta, N_t) = \underline{R}$. In other words, the screening property 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 $R = \underline{R}$ 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.

3 Discussion

3.1 (Un)necessary Doubt

In Section 2, we consider a regime change game and we show that a timely disaster alert stops the agents from panicking under the doubt assumption on the information structure. In our setting, the agents can have homogeneous or heterogeneous and arbitrarily correlated beliefs regarding the fundamental.

Recall that we assume that there exists some $\epsilon > 0$ such that $P(\theta \geq 1|s_i) > \epsilon$ for any s_i (Assumption 1). This is a sufficient condition which guarantees that the probability of making a mistake by attacking at time 0 is positive regardless of what other agents do. But is this a necessary condition? What if there is some s_i such that an agent who receives the signal believes that $P(\theta \geq 1|s_i) = 0$?

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. In standard global game information structure, as will be discussed in the next subsection, assumption 1 is not needed for our main result.

3.2 Relation to Global Game literature

In Section 2, we investigate the effectiveness of the disaster alert policy in a general information environment. In particular, under some reasonable assumptions on the information structure, we show that setting such policy in a timely manner can completely eliminate coordination failure. However, without a specific information structure, we cannot discuss some important question such as, what will happen when there is no information disclosure, or how effective the proposed policy is when timely disclosure is not possible.

In this section, to understand these questions, we will consider a game with a specific information structure while keeping payoff specification exactly the same as in Section

2. Being more specific, we will start with the standard global game of regime change Morris and Shin (2003) and extend the static game to an endogenous move game with the disclosure policy of disaster alert in place.

The information structure is as follows.[5] Agents have common prior about the underlying fundamental θ is $U[\underline{\theta}, \bar{\theta}]$. In addition, agent $i \in [0, 1]$ receives noisy private information about θ, 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. σ scales the random noise ϵ_i. We assume that $\underline{\theta} \leq -\sigma$ and $\bar{\theta} > 1 + \sigma$.[6]

Under this information environment, Assumption 1 is violated. For example, for an agent with $s_i < 1 - \frac{1}{2}\sigma$, he understands that the fundamental is not in the upper dominance region for sure, i.e., $\mathbb{P}(\theta \geq 1|s_i) = 0$.

Before we move forward to examine the effectiveness of the disclosure policy Γ^τ, under this specific information structure, we will first remind the readers of the equilibrium outcome when there is no further information disclosure.

Proposition 1 *When there is no information disclosure, in the unique rationalizable equilibrium, agents attack the regime if and only if $s < s^* = \theta^* + \sigma F^{-1}(\theta^*)$. Hence, a regime fails if and only if its fundamental*

$$\theta < \theta^* \equiv \frac{1}{1 + \frac{g-1}{1-l}}.$$

When there is no information disclosure, any delayed attack is not rational. For that reason, the game is reduced to a static game, which is exactly the same as Morris and Shin (2003). Proposition 1 is a standard result in the global game literature and thus the proof is omitted here.

Recall that the screening property in Lemma 1. Under the disclosure policy Γ^τ, any rational agent will either attack immediately (or $\mathcal{A}$) or wait for the disclosure and attack when only the disaster alert is triggered (or $\mathcal{W}^1$). That is why $N_\tau = N_0$ and the regime fails only when the disaster alert is triggered ($\theta < N_0$).

This screening property builds a nice connection between the game with the disclosure

[5] Here, we work with the a common prior with bounded support and bounded noise. One can check that everything we discuss here can be easily extended to the case with normal prior and normal distributed noise.

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

policy Γ^τ and the one without. When there is no disclosure, agents choose whether to attack the regime or not. The regime fails only when the mass of agents who decide to attack is large enough to fail the regime. While in the game under the disclosue policy Γ_τ, agents are also making a binary choice between attacking immediately ($\mathcal{A}$) or the strategy of wait and see ($\mathcal{W}^1$). In this case, the regime fails only when the mass of agents who attack immediately is large enough to fail the regime. Hence, based on the standard iterated elimination argument in the global game literature, we can show that in equilibrium agents play the threshold strategy. They choose to not to wait for the disclosure and attack the regime immediately only when their private information is lower than some cutoff. The following proposition summarizes this result.

Proposition 2 *Under the disclosure policy Γ_τ, agents take the strategy $\mathcal{A}$ if and only if $s_i < \hat{s}^\tau = \hat{\theta}^\tau + \sigma F^{-1}(\hat{\theta}^\tau)$. The regime fails if and only if*

$$\theta < \hat{\theta}^\tau \equiv \frac{1}{1 + \frac{g-1}{1-e^{-r\tau}}}$$

and $\Theta^P(\Gamma^\tau) = [0, \hat{\theta}^\tau]$.

Proof. See Appendix. ∎

Irrespective of the disclosure policy, agents can always attack the regime at time 0 without waiting for the new information. In the case without information disclosure, the other choice is committing to not attacking at time 0. However, under the disclosure policy Γ_τ, agents can take the strategy of wait and see. This strategy is indeed a better one because agents still have chance to attack the regime later when the disaster alert is triggered. Under this strategy, as discussed before, they will only attack the regime that fails in the end and panics will be eliminated completely. Hence, the disclosure policy Γ_τ makes agents less likely to attack the regime, i.e., $s^* > \hat{s}^\tau$, and makes the regime more likely to survive, i.e., $\theta^* > \hat{\theta}^\tau$ for any $\tau \in (0, T)$.

Recall that we did not solve for the equilibrium for $\tau > \hat{\tau}$ under the general information structure in Section 2. Adopting the information structure in the standard global game enables us to have a full characterization of the equilibrium. It is clear that when the principal has some flexibility in selecting the time of disclosure, she would set the disaster alert as early as possible, since $\hat{\theta}^\tau$ is decreasing in the disclosure time τ. The following corollary summarizes the results discussed above.

Corollary 2 *In a global game of regime change,*

1. *the disaster policy helps to reduce the panic, i.e., for any* $\tau \in (0, T)$*,* $\hat{\theta}^\tau < \theta^*$ *and* $\hat{s}^\tau < s^*$*;*

2. *the principal would set a disaster alert as early as possible, i.e., for any* $\tau < \tau'$*,* $\Theta^P(\Gamma^\tau) \subset \Theta^P(\Gamma^{\tau'})$*.*

3. *(limiting case) if the principal can set the disclosure at* $\tau \to 0$*,* $\Theta^P(\Gamma^\tau) \to \emptyset$*.*

Proof. Result in (1) is obvious since any delayed attack generates a higher payoff than no attack, i.e., $e^{-rT} > l$. (2) and (3) follow immediately from Proposition 2. ■

The limiting result in Corollary 2 resembles our main result in Theorem 1. This confirms that when agents have heterogenous private signal, the doubt assumption is not necessary for a timely disaster alert policy to be effective.

3.3 Relation to Information Design Literature

Similar to this paper, Inostroza and Pavan (2017) also consider a regime change game where agents receive some noisy private signals, where the noises can be correlated. However, unlike this paper, the agents have to move simultaneously. The authors consider an information designer who commits to an information disclosure rule. Since the agents move simultaneously, it is a static information design problem. The authors characterize the optimal disclosure policy.

Under some condition, this optimal disclosure policy is a public disclosure similar to a stress test, i.e., the designer discloses whether $\theta \geq k$ or not, for some fixed k. This is not the first best outcome. The regimes with $\theta \in (0, k)$ will fail even when it could have passed 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 principal cannot eliminate panic. In sharp contrast, when the agents endogenously choose when to attack, we show that the principal can eliminate panic using a simple dynamic information disclosure policy - a timely disaster alert.

The disaster alert can be thought of as the weakest stress test ($k = 0$). When agents move simultaneously Inostroza and Pavan (2017) and Goldstein and Huang (2016) show that if the principal discloses $\theta \geq k$, then there can be multiple equilibria. But when k

is sufficiently large, even in the worst equilibrium, i.e., the one in which the agents attack most aggressively, the agents will not attack a regime that passes the stress test. Under endogenous timing, as long as delay is costly, an agent who has waited for the test, will never attack after the regime passes the test. Thus, unlike the case where agents move simultaneously, even the weakest stress test takes away all the strategic uncertainty after the test results are disclosed. We then argue that, if the alert is set in a timely manner, then the unique rationalizable strategy is to wait for the alert and follow the principal's recommendation thereafter.

Finally, 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 (2018) also consider the same problem with the only difference that the agents move sequentially in an exogenous order. When timing is not endogenous, the weakest stress test cannot eliminate the strategic uncertainty. It is possible that in equilibrium the agents do not attack a regime that passes the test, but there are other equilibria in which they do attack even when the regime passes the test. Basak and Zhou (2018) show that if the principal runs the tests sufficiently frequently, then there is a unique monotone equilibrium in which the agents do not attack the regime that passes the test.

Unlike the exogenous timing case, the principal does not need to repeat the tests. A one time weakest test is enough to eliminate panic as long as it is done in a timely manner. Also, this paper uses Rationalizability as the solution concept, while Basak and Zhou (2018) is limited to Monotone equilibrium only. In this sense the endogenous timing makes it easy for the principal to dissuade the agents from attacking.

3.4 Selling Information and (Ir)reversibility

The dynamic setup we consider a three crucial features - the agents endogenously choose when to attack, attack is irreversible, and delayed attack is costly. The disclosure rule is simple but it not only depends on the fundamental but also the endogenous attack so far. In a static setup such as Inostroza and Pavan (2017), these features are missing for obvious reason. But the principal may be able to introduce endogenous information disclosure by selling information rather than just disclosing them as in Bergemann, Bonatti and Smolin

(2018). For example, suppose before the game begins, the principal sells a supplementary news d at a very small price c. An agent receives his private signal and decides whether to buy this supplementary news d or not. Suppose that b proportion of the agents buy d. An agent who buys d, gets the message

$$d = \begin{cases} 1 \text{ if } \theta < 1 - b \\ 0 \text{ if } \ \theta \geq 1 - b. \end{cases}$$

Based on his private signals and the supplementary news, if he buys it, an agent decides whether to attack the regime or not. The regime change game is as before. If N proportion of agents attack and $\theta \geq N$, then the regime survives, otherwise it fails. An agent gets 1 if he attacks, but if he does not, then he gets $g(> 1)$ if the regime survives and $l(< 1)$ if the regime fails.

Thus, by selling the supplementary news, the principal introduces some of the dynamic features we need in a static game. Buying the supplementary news is similar to not attacking early. Also, it is costly. Thus, similar to our dynamic setup, by buying the supplementary news, an agent could signal his future action. However, since an agent who does not buy the supplementary news, can attack later, this policy does not generate the irreversibility of attack feature. Below, we argue that although this policy convinces an agent to never attack without buying the supplementary news, this may not achieve the first best outcome.

If an agent buys d, then it must be that he will decide differently based on whether $d = 1$ or $d = 0$, otherwise, there is no positive option value of this supplementary news d and he should not have bought d in the first place. It is natural to think that an agent attacks after learning $d = 1$ and not attack after learning $d = 0$. Thus, when $d = 0$, the agent understands that no agent who has bought d will attack, and even if all the agents who have not bought d decide to attack, the regime will survive. Therefore, if an agent buys d, his expected payoff is

$$P(d = 1|s) \cdot 1 + P(d = 0|s) \cdot g - c.$$

In contrast if an agent does not buy d and attacks, he gets 1. Since $P(d = 0|s) = P(\theta \geq 1 - b|s) \geq P(\theta > 1|s) > 0$, then for $c \to 0$, buying d dominates not buying d and attack. However, note that an agent with high signal may not buy the supplementary news and hence b could be less than 1. Hence, even if $\theta > 0$, the supplementary news may reveal $d = 1$, and those who buy the news will attack, and regime may fail. Thus, panic may happen in equilibrium. It will be interesting to find out what is optimal way to sell supplementary news to minimize panic. We leave this for future research.

3.5 General Payoff

Our payoff specification can be easily extended to a general regime change game. Suppose that the regime survives if and only if $w(\theta, N) \geq 0$ where the partial derivatives $w_\theta > 0$ and $w_N < 0$. Agent i gets $u(\theta, t_i, N_T)$ if he decides to attack at $t_i \in [0, T]$, and $v(\theta, N_T)$ if he never attacks.

Assumption 3 *The payoff specification is such that*

1. *(Costly Delay) For all $\theta \in \Theta$, $N \in [0, 1]$, u is strictly decreasing and Lipschitz continuous in t.*

2. *(Complementarity) For all $\theta \in \Theta$, $t \in [0, T]$, the net payoff from staying as opposed to attacking $(v - u)$ is strictly decreasing in N, i.e.,*

$$\forall\ N, N' \in [0, 1] \text{ such that } N > N', v(\theta, N) - u(\theta, t, N) \leq v(\theta, N') - u(\theta, t, N').$$

3. *For any $\theta \in \Theta$, $N \in [0, 1]$ and $t \in [0, T]$, there exists $\delta > 0$ such that $v - u > \delta$ whenever $w(\theta, N) \geq 0$.*

The first two assumptions are natural, while Assumption 3.3 ensures some strictly positive gain from not making a mistake by attacking immediately without waiting for the disclosure. This assumption tells us what qualifies as "timely" disaster alert. If $\delta \to 0$, then the timely alert also needs to be close to 0.

Proposition 3 *When the payoff specification satisfies assumption 3, there exists $\tilde{\tau}_\delta > 0$ such that setting the disaster alert policy Γ^τ at $\tau < \tilde{\tau}_\delta$, $\Theta^P(\Gamma^\tau) = \emptyset$.*

Proof. See Appendix ■

3.6 Timely Stress Testing of Banks

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 such as panic, including Goldstein and Huang (2016), Inostroza and Pavan (2017), investigate such question in a static setting.

Taking the endogenous move of withdrawal decision (or attacking in our model) into consideration, we show that disclosing the information about whether a bank is fundamentally flawed and doomed to fail in a timely manner is the optimal policy which completely quells the financial panic. Moreover, this binary disclosure rule is similar to what we have seen in the stress testing of banks. The information released will be credible and thus if the disaster alert is not triggered, all investors should restore their confidence about the bank's fundamental and not worry about other investors' panicking. Our main result shows from the time the policy designer commits to conducting such a test, there will be no panic.

In a dynamic setting, apart from the disclosure rule, or stress testing in this example, disclosing information in a timely manner is critical to our result. 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. In the guidance on Stress Testing disclosed by FED, FDIC and OCC, [7], it was also 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"*.

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) *"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.

Although the above evidence (including the adopted information disclosure policy as well as the impact of such policy) seem to support our theoretical results, however, in this discussion, we want to take one step back and think about whether our model captures all

[7] see SR Letter 12-7

important features of a financial panic. A critical assumption in our model is that, within a short time window, agents are able to attack a regime even when the regime is doomed to fail. Although there is some cost for delayed attack, the payoff is still continuous in the timing of attacks.

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.

That 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.[8] According to our 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.

4 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* $t \in [t_1 + dt, T]$*.*

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

$\mathcal{W}^2$: *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 Note that $R(\theta, N_t)$ is weakly decreasing in t. The disaster alert is triggered, or $d^\tau = 1$, if and only if

$$\begin{aligned} & u(\theta, \tau, N_\tau) = \int_{t=0}^{\tau} e^{-\beta t}(1-\mu)R(\theta, N_t)dt + \frac{r}{\beta}e^{-\beta\tau} \\ > \quad & v(\theta, (N_t)_{t<\tau}) = \int_{t=0}^{\tau} e^{-\beta t}R(\theta, N_t)dt + \frac{R(\theta, N_\tau)}{\beta}e^{-\beta\tau}. \end{aligned}$$

At any time τ, if $R(\theta, N_\tau) = \overline{R}$, then for any $t \in [0, \tau)$, $R(\theta, N_t) = \overline{R}$. Hence,

$$v(\theta, (N_t)_{t<\tau}) = \frac{\overline{R}}{\beta} > u(\theta, \tau, N_\tau) = \frac{\overline{R}}{\beta}(1-\mu)(1-e^{-\beta\tau}) + \frac{r}{\beta}e^{-\beta\tau}$$

and thus $d^\tau = 0$. Otherwise, if $R(\theta, N_\tau) = \underline{R}$, then $\tau < \hat{T} \equiv \frac{1}{\beta}\log(1 + \frac{r-\underline{R}}{\mu\overline{R}})$ guarantees exiting at τ is the strictly dominant action ($d^\tau = 1$), i.e.,

$$\begin{aligned} v(\theta, (N_t)_{t<\tau}) - u(\theta, \tau, N_\tau) &= \int_0^\tau e^{-\beta t}\mu R(\theta, N_t)dt - e^{-\beta\tau}\frac{r-\underline{R}}{\beta} \\ &\leq \frac{1}{\beta}\left[(1-e^{-\beta\tau})\mu\overline{R} - e^{-\beta\tau}\frac{r-\underline{R}}{\beta}\right] \\ &< 0 \quad (\text{if } \tau < \hat{T}) \end{aligned}$$

Hence, when $T < \hat{T}$, $R = \underline{R}$ at any time $\tau \in (0, T)$ is equivalent to the disaster alert $d^\tau = 1$.

The screening property holds because $u(\theta, t, N)$ is decreasing in t (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 $\mathcal{W}^1$ is g when $d^\tau = 1$ (or equivalently $\theta \geq N_0$) and $e^{-r\tau}$ 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 $\mathcal{W}^1$ is $[1 + \frac{1}{2}\sigma, \bar{\theta} + \frac{1}{2}\sigma]$. The rest of the proof follows the standard iterated elimination arguments as the one in Proposition 1. □

Proof of Corollary 2

1. This is obvious since any delayed attack generates a higher payoff than no attack, i.e., $e^{-rT} > l$.

2. This follows immediately from Proposition 2.

3. This result resembles our main result in Theorem 1. Recall that Assumption 1 is violated here for this information structure. The following proofs show how the iterative arguments work here for this heterogeneous information structure.

   Fix any small positive number $\epsilon > 0$. Define $\hat{s}_1$ such that

   $$\mathbb{P}(\theta \geq \hat{\theta}_1 \equiv 1 | \hat{s}_1) = \epsilon.$$

   So, for any $s > \hat{S}_1$, the doubt assumption is satisfied. We know that

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

   Then, we know that any agent with $s > \hat{s}_1$ would not attack immediately. We can 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 fail the regime and trigger 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 $\epsilon < \hat{\theta}_2 < \hat{\theta}_1 = 1$. 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 ϵ independent of what others do, i.e.,

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

   Hence, $\hat{s}_2 < \hat{s}_1$ since $\hat{\theta}_2 > \epsilon$. 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

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

   Note that when $\epsilon \to 0$, $s^*_\epsilon \to -\frac{\sigma}{2}$. Following Proposition 2, one can understand that under the disclosure policy Γ^0 , any agent who does not believe the regime fails for sure (i.e., $s < -\frac{\sigma}{2}$), would never attack the regime and thus the panics are completely eliminated.

□

Proof of Proposition 3 First of all, the screening property holds true because of the delay cost (see the proof in Lemma 1). Hence, the only rationalizable strategies are $\mathcal{A}$ and $\mathcal{W}^1$. The expected payoff from taking $\mathcal{A}$ is

$$\mathbb{E}\left(u(\theta, 0, N_T)|s_i\right),$$

while the expected payoff from taking $\mathcal{W}^1$ is

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

Under the screening property, $N_\tau = N_T$ when $d^\tau = 0$, and $N_T = 1$ when $d^\tau = 1$. Thus, the expected payoff difference is

$$\begin{aligned} D(\Gamma_\tau, s_i) = &\mathbb{P}(d^\tau = 1|s_i)\mathbb{E}\left(u(\theta, \tau, 1) - u(\theta, 0, 1)|s_i, d^\tau = 1\right) \\ &+ \mathbb{P}(d^\tau = 0|s_i)\mathbb{E}\left(v(\theta, N_\tau) - u(\theta, 0, N_\tau)|s_i, d^\tau = 0\right) \end{aligned}$$

Under Assumption 3.2, $d^\tau = 0$ implies that $v(\theta, N_\tau) - u(\theta, \tau, N_\tau) > 0$. Then, when Assumption 3.3 is satisfied, $v(\theta, N_\tau) - u(\theta, \tau, N_\tau) > \delta$. Moreover, $\theta \in \Theta^U$ guarantees that $d^\tau = 0$ and thus under assumptions 1, $\mathbb{P}(d^\tau = 0|s_i) > \epsilon$. Given u is Lipschitz continuous in t (Assumption 3.1), there exists K such that for any $\theta \in \Theta$,

$$u(\theta, 0, 1) - u(\theta, \tau, 1) < K\tau$$

Therefore,

$$D(\Gamma_\tau, s_i) > -(1-\epsilon)K\tau + \epsilon\delta$$

Thus, whenever $\tau \leq \tau_\delta \equiv \frac{\epsilon}{1-\epsilon}\frac{\delta}{K}$, $\mathcal{W}^1$ strictly dominates $\mathcal{A}$ for any s_i.

Now we will discuss the effectiveness of the policy Γ^τ with $\tau < \tau_\delta$ focusing on this unique rationalizable strategy $R(\Gamma_\tau) = \mathcal{W}^1$. It is easy to see that when there is no disaster alert (or $d^\tau = 0$), $N_\tau = N_T = 0$. Under such a policy, let us define the set of fundamental θ that triggers no alert as

$$\Theta^{d^\tau=0}(\Gamma^\tau) \equiv \{\theta | v(\theta, N_\tau = 0) > u(\theta, \tau, N_\tau = 0)\}.$$

Hence, under the Assumption 3.3, $v(\theta, N_\tau = 0) - u(\theta, \tau, N_\tau = 0) > \delta$ for all $\theta \in \Theta^{d^\tau=0}$. Based on the continuity of u on τ, there exists $\tau'_\delta > 0$ such that for any policy Γ^τ with

$\tau < \tilde{\tau}_\delta \equiv \min\{\tau_\delta, \tau'_\delta\}$,

$$\Theta^{d^\tau=0}(\Gamma^\tau) = \{\theta | v(\theta, 0) > u(\theta, 0, 0)\}$$

Hence, for $\tau < \tilde{\tau}_\delta$, $\Theta^{d^\tau=0}(\Gamma^\tau) = \Theta \setminus \Theta^L$ and $\Theta^{d^\tau=0}(\Gamma^\tau) \cap \Theta^F(\Gamma^\tau) = \emptyset$. That implies $\Theta^p = \emptyset$.

When $\delta \to 0$ (Assumption 3.3 is violated), both τ_δ and τ'_δ converges to 0 and the result only holds true in the limit when $\tau \to 0$. □

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. 2018. "Diffusing Coordination Risk."

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

Bergemann, Dirk, Alessandro Bonatti, and Alex Smolin. 2018. "The Design and Price of Information." *American Economic Review*, 108(1): 1–48.

Bergemann, Dirk, and Stephen Morris. 2013. "Robust Predictions in Games with Incomplete Information." *Econometrica*, 81(4): 1251–1308.

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.

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.

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.

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).

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.