Diffusing Coordination Risk

By Deepal Basak and Zhen Zhou*

In a regime change game, privately informed agents sequentially decide whether to attack without observing others' previous actions. To dissuade them from attacking, a principal adopts a dynamic information disclosure policy – frequent viability tests. A viability test publicly discloses whether the regime has survived the previous attacks. When such tests are sufficiently frequent, in the unique cutoff equilibrium, agents never attack if the regime passes the latest test, regardless of their private signals. We apply this theory to demonstrate that a borrower can eliminate panic-based runs by sufficiently diffusing the rollover choices across different maturity dates. (*JEL* C72, D82, D83, G28, G33)

In a coordination game, the strategic uncertainty that agents face concerning the actions and beliefs of others may lead to undesirable outcomes. Consider a borrower who has issued short term debts to finance some illiquid investment. When the debt matures, a creditor may not roll over if he is wary of other creditors withdrawing their funds. This debt structure could cause a debt run solely on the basis of panic and not on underlying fundamental. We model this as a global game of regime change.[1] In a dynamic setting, we propose a simple information disclosure policy that influences the beliefs – and thus, actions – of agents. This policy resolves strategic uncertainty and avoids undesirable outcomes.

Consider a regime, a principal, and a mass 1 of agents. The agents move sequentially: an agent $i \in [0,1]$ moves at time i and decides whether to attack a regime or not, but he does not see other agents' previous actions. The underlying fundamental strength of the regime is θ. If θ is sufficient to withstand the

* We are grateful to Douglas Gale, Laurent Mathevet, David Pearce, and Ennio Stacchetti for their invaluable advice and continued support on this project, and to Jeff Ely (co-editor) and three anonymous referees for their thoughtful and constructive comments. We also thank Viral Acharya, Toni Ahnert, George-Marios Angeletos, Jess Benhabib, William Cong, Amil Dasgupta, Joyee Deb, William Fuchs, Itay Goldstein, Zhiguo He, Chong Huang, Cecilia Parlatore, Thomas Philippon, Edouard Schaal, Michal Szkup, Laura Veldkamp, Ming Yang, Sevgi Yuksel and participants at Penn State, NYU Financial Economics workshop, CUHK(SZ), Fudan, Seoul National University, HKUST, Delhi Theory workshop, RES 2015, SED 2015, Global Games at Ames, CICF 2016, UECE Lisbon 2016, The 9th Shanghai Microeconomics Workshop and Econometric Society meetings (St. Louis, Lisbon and Hong Kong) for their helpful comments and suggestions. We thank Mayur Choudhary 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`

[1] For examples of a global game of regime change, see Morris and Shin (1998) and Angeletos, Hellwig and Pavan (2007) for the currency crisis, Goldstein and Pauzner (2005) for self-fulfilling bank runs, Vives (2014) for financial fragility, Edmond (2013) for riots and political change and Konrad and Stolper (2016) for fight against tax havens. For more recent developments see Szkup and Trevino (2015).

aggregate attack, the regime survives; otherwise, it does not. Suppose there is a threshold p such that an agent will not attack if and only if he believes with probability higher than p that the regime will survive. As is the standard in literature of global games, agents are uncertain about the underlying fundamental θ and receive some noisy private signals s_i.

The principal wants the regime to survive. However, she cannot influence the exogenous fundamental. She can strategically disclose information at different dates regarding the exogenous fundamental and past endogenous attacks to dissuade the agents from attacking.

First, let us consider two simple policies – no disclosure and full disclosure. Under no disclosure, agents have no information about past actions of other agents. Therefore, the game is essentially a simultaneous move game (see Morris and Shin (2003)), and there is a unique equilibrium in which agents attack if and only if they receive a private signal below some cutoff.

On the other hand, if agents know the underlying fundamental and perfectly observe other agents' previous actions, all attacking is a possible equilibrium outcome. If there are only finitely many agents, it follows from backward induction that agents will coordinate on the payoff-dominant action.[2] However, in many examples such as the currency crisis problem (See Morris and Shin (1998) and Angeletos, Hellwig and Pavan (2007)), attacking the fixed exchange rate regime is the payoff-dominant action for agents, while the principal's interest is not aligned with that of the agents because the principal wants to defend it.

Thus, even if full disclosure is feasible, the principal may want to conceal some information. A vast range of partial disclosure policies exist. We consider a simple dynamic information disclosure policy – *frequent viability tests*. As the name suggests, a viability test at some date $t \in [0, 1]$ checks whether the regime continues to be viable – i.e., if it has survived attacks thus far (if any). The result of this test is publicly disclosed. The principal only chooses an integer J denoting the frequency of viability tests. The tests are conducted at a regular interval of $1/J$, starting at 0.

Our main result shows that if the principal runs viability tests with sufficient frequency, then there is a unique cutoff equilibrium in which agents ignore their private information and never attack a viable regime.

In the spirit of Bayesian Persuasion, we can interpret this policy as a recommendation by the principal to the agents not to attack a regime that passes the latest viability test and to attack if it fails. If the regime fails a viability test, it cannot survive even without any further attack. Hence, it is the dominant strategy for all agents in the subsequent groups to follow the principal's recommendation and attack. The challenging case arises when the regime passes the test.

If the regime passes a viability test, then the agents become more optimistic about the fundamental θ and other agents' beliefs about θ and so on. If one agent is less likely to attack, then it follows from strategic complementarity that others

[2]We discuss the case of full disclosure in detail in Section V.

are also less likely to attack. Thus, it is intuitive that positive viability news reduces aggregate attack in the equilibrium.

When the principal runs viability tests for J times, the policy separates the agents into J groups. The $1/J$ mass of agents moving between the jth test and the $(j+1)$th test are referred to as the group j agents. We examine the equilibrium in cutoff strategies – for any j, after learning that the regime has passed the jth viability test, any agent in the group j does not attack if and only if his private signal is higher than some cutoff $\hat{s}_j$. Thus, when the regime is stronger (higher θ), more agents receive signals that exceed the cutoff and fewer agents from each group attack. This induces a non-decreasing sequence of fundamental cutoffs $\{\underline{\theta}_{j-1}\}_{j=1}^{J}$ such that the regime passes the jth test if and only if the underlying fundamental is no lower than the cutoff $\underline{\theta}_{j-1}$.

These cutoffs are endogenous, which makes the policy history-dependent. Moreover, the effectiveness of the present viability news depends on the effectiveness of the future viability news. We develop an inductive argument: if the tests are sufficiently frequent, then agents in the group j follow the principal's recommendation if they believe that agents in the subsequent groups will do so regardless of the cutoff strategy played by others in the past (equivalently for any $\underline{\theta}_{j-1}$).

Given that the regime has passed the jth viability test and the agents in subsequent groups will follow the principal's recommendation, the regime will survive as long as it withstands attacks from the group j. Suppose that the agents in group j follow some cutoff strategy $\hat{s}_j$. This monotone strategy can constitute an equilibrium if the "marginal agent" with the signal $\hat{s}_j$ is indifferent between attacking and not attacking – i.e., he believes that the regime survives with probability p. We argue that there exists a $\hat{J}$ such that if the group size is smaller than $1/\hat{J}$, then no such equilibrium can exist.

This is because, under frequent viability tests, the size of each group is small. If the group size is smaller, the magnitude of attack from this group is less for any given cutoff strategy $\hat{s}_j$. Given that no one from the subsequent groups will attack a viable regime, the regime is more likely to survive. This, in turn, amplifies the influence of positive news from the jth viability test and makes the marginal agent more confident about the survival of the regime. We demonstrate that the marginal agent's belief regarding the chance of survival *uniformly converges* to 1 with group size regardless of $\underline{\theta}_{j-1}$. It follows from uniform convergence that there exists a $\hat{J}$ such that, under sufficiently frequent viability tests (i.e., $J > \hat{J}$), the marginal agent always believes that the regime will survive with a probability strictly higher than p; and thus, he strictly prefers not to attack if agents in the subsequent groups will not attack a regime that passes the latest viability test.

Under sufficiently frequent viability tests, since no one moves after the last group of agents, there cannot be any cutoff equilibrium in which group J agents attack a viable regime (regardless of the cutoff strategies played by others in the past). Given that, so will group $J-1$ agents and so on. Thus, the risk that agents may attack a viable regime unravels from the end.

In practice, borrowers often adopt an asynchronous debt structure – i.e., diversify debt rollovers across dates. This type of debt structure diffuses the rollover risk that would otherwise be concentrated at a single maturity date. Under such a diffused structure, the creditors whose debts mature at a later date can learn whether the borrower is still viable (has not defaulted) or not. Hence, there is essentially a viability test at each maturity date. The news that the borrower has not defaulted yet may sound obvious and not a deliberate attempt to manipulate agents' beliefs. One may believe that the news is not likely to have any substantial influence on creditors' behavior. However, the borrower provides this news J many times when only $1/J$ fraction of debts mature at a time. We extend our model to show that a sufficient asynchronous debt structure (sufficiently large J) makes the borrower immune to panic-based runs.

Related Literature.– Similar to a Pigouvian planner, our principal attempts to achieve the desired outcome that the market fails to deliver. Sakovics and Steiner (2012) and Cong, Grenadier and Hu (2018) find optimal subsidies that, at a given cost, maximize the likelihood of successful coordination.[3] Unlike the planner in the papers mentioned above, we consider a principal who cannot offer monetary incentives. Instead, she acts as an information designer and discloses some relevant information. We consider a canonical global game of regime change, and as is standard in the literature, we assume a private information environment (see Carlsson and Van Damme (1993)). Similar to Bergemann and Morris (2013), the principal does not have access to agents' private information.

The two most closely related papers are Inostroza and Pavan (2017) and Goldstein and Huang (2016). The authors also propose a partial information disclosure policy - a one time "stress test."[4] In this paper, the principal exploits the fact that agents do not move simultaneously and runs multiple viability tests over time. Thus, the paper belongs to the recent literature on dynamic information design. Ely (2017) is the first paper to extend the static Bayesian Persuasion model of Kamenica and Gentzkow (2011) to the dynamic setting, where the disclosure policies are history-independent. Doval and Ely (2019) consider a general extensive form game with incomplete information in which the principal does not know the exact extensive form that governs the play.[5]

Our paper is related to the dynamic coordination game literature as well. Dasgupta (2007) considers a two-period problem in which agents receive noisy private information about the attacks from the first period. Similar to our model, in Angeletos, Hellwig and Pavan (2007) and Huang (2017), agents learn from the

[3]While Sakovics and Steiner (2012) consider heterogeneous agents and demonstrate that subsidizing the more reluctant agents matters more, Cong, Grenadier and Hu (2018) demonstrate that if liquidity injection is equally costly across periods in a dynamic coordination problem, early injection is more helpful.

[4]Information design has been studied in other strategic contexts, such as voting and auctions. See Bergemann and Morris (2019) for a recent survey of this literature.

[5]Significant work has been conducted on dynamic information feedback in the context of strategic experimentation. See Hörner and Skrzypacz (2016) for a survey of this literature.

past through viability news, which results in multiple cutoff equilibria. Unlike the papers mentioned above, we optimally choose the information transmission over time. Frankel and Pauzner (2000) introduce asynchronousity by allowing the agents to revise their decisions following a Poisson process. In their model, agents know the current state but are uncertain about the volatile future.[6] He and Xiong (2012) extend this framework to study the role of volatile fundamental under asynchronous debt structure. While the two papers mentioned above take the asynchronous structure as a given, we take a step back and address what necessitates the asynchronous debt structure in the first place. We find the justification in terms of information transmission and identify the optimal design of the asynchronous structure from the borrower's perspective.

The rest of the paper is organized as follows. Section I describes the model, the diffused policy, and the solution concept. Section II illustrates the role of a one time viability test. In Section III, we leverage this idea to develop the argument for frequent viability tests and establish our main result. In Section IV, we extend the model to study its application in debt rollover. In Section V, we discuss full disclosure policy and the factors that may make persuasion harder or even impossible. Section VI concludes this paper. Proofs that are not presented in the paper are provided in the appendix.

I. A Simple Model of Regime Change

There is a regime, a principal, and a mass 1 of agents. The agents are indexed by $i \in [0, 1]$. Agent i moves at time i and takes an action $a_i \in \{0, 1\}$, where 1 (0) indicates attacking (not attacking) the regime.

The agents cannot observe other agents' previous actions. Let θ be the underlying fundamental strength of the regime and $w = \int_i a_i di$ be the aggregate attack. The regime survives if and only if its fundamental strength θ is strong enough to withstand the aggregate attack against it – i.e., $\theta \geq w$.

Agents are assumed to be ex-ante identical and risk-neutral. Given that the fundamental strength (θ) and the aggregate attack (w), let $u(a_i, w, \theta)$ be the payoff for agent i if he takes action a_i, where

$$(1) \qquad u(0, w, \theta) = \begin{cases} b_1 & \text{if } \theta \geq w \\ c_0 & \text{if } \theta < w \end{cases}, \qquad u(1, w, \theta) = \begin{cases} c_1 & \text{if } \theta \geq w \\ b_0 & \text{if } \theta < w. \end{cases}$$

We assume that (1) $b_1 > c_1$ – i.e., if an agent knows that the regime will survive, not attacking is the desirable action – and (2) $b_0 > c_0$ – i.e., if an agent knows that the regime will not survive, attacking is the desirable action.[7]

[6]Sequential move has also been studied in the context where agents' payoffs are independent of others' actions but there is incomplete information about a common fundamental (see Banerjee (1992) and Bikhchandani, Hirshleifer and Welch (1992)). For games of strategic substitutes such as public good contribution, Varian (1994) shows that sequential move usually reduces the total contribution.

[7]This payoff specification is standard in the literature; however, as the agents move sequentially, in

However, agents are uncertain about θ and hence do not know whether the regime will survive. Each agent i receives a noisy private signal about θ, denoted by $s_i = \theta + \sigma\epsilon_i$, where ϵ_i is a random noise with zero mean, and $\sigma > 0$ scales the random noise ϵ_i. It is common knowledge that θ follows a distribution Π, and the error terms ϵ_i are drawn independently of θ and follow the independent and identical distribution F.[8] We assume that F has a support $[-\frac{1}{2}, \frac{1}{2}]$ and that it admits a continuous density f that is strictly positive and bounded in this support – i.e., $0 < \underline{f} \leq f(\epsilon) \leq \overline{f} < \infty$. Moreover, f is *log-concave*. We also assume that Π has a support $[\underline{\theta}, \overline{\theta}]$ that is sufficiently wide such that $\underline{\theta} < -\sigma$ and $\overline{\theta} > 1 + \sigma$. Π admits a continuous density π that is strictly positive and bounded in $[-\sigma, 1+\sigma]$ – i.e., $0 < \underline{\pi} \leq \pi(\theta) \leq \overline{\pi} < \infty$.

If $\theta \geq 1$, the regime will surely survive regardless of w, and if $\theta < 0$, the regime cannot survive regardless of w. Only when $\theta \in [0, 1)$, the regime's survival is dependent upon the aggregate attack w. Note that for any $\theta \in [0, 1)$, the proportion of agents who receive a signal lower than s is $F(\sqrt{\tau}(s - \theta))$. For an agent with private signal $s \geq 1 + \sigma/2$, the dominant strategy is not to attack, and for an agent with private signal $s < -\sigma/2$, the dominant strategy is to attack. An agent who receives a private signal $s \in [-\sigma/2, 1+\sigma/2)$, does not have a dominant strategy. He updates his belief that $\theta \geq A$ (for some A) as[9]

$$P(\theta \geq A|s) = \int_A^{\overline{\theta}} \left(\frac{\pi(\theta) f(\sqrt{\tau}(s - \theta))}{\int_{\underline{\theta}}^{\overline{\theta}} \pi(\theta) f(\sqrt{\tau}(s - \theta)) d\theta} \right) d\theta.$$

The log-concavity of f is equivalent to the *monotone likelihood ratio property (MLRP)*. That is, for $s_1 > s_2$,

$$\frac{f(\sqrt{\tau}(s_1 - \theta))}{f(\sqrt{\tau}(s_2 - \theta))}$$

is increasing in θ. This implies that, for any $A > B$, $P(\theta \geq A|s, \theta \geq B)$ is increasing in s (see the online appendix for the formal argument).

Suppose an agent believes that the regime survives with probability P. The expected payoff from attacking ($a_i = 1$) is $Pc_1 + (1 - P)b_0$, while the expected payoff from not attacking ($a_i = 0$) is $Pb_1 + (1 - P)c_0$. Therefore, the agent does not attack if and only if he believes that the probability of the regime surviving

practice, the payoff may depend on the history of attacks. In Section IV, we show that our result is robust to more general payoff structures.

[8]See Judd (1985) for the existence of a continuum of independent random variables.

[9]Note that although we use the $\underline{\theta}$ and $\overline{\theta}$ in the limit of the integration, for a given s, if $\theta > s + \sigma/2$ or $\theta < s - \sigma/2$, $f(\sqrt{\tau}(s - \theta)) = 0$.

(or P) is greater than

$$p := \frac{1}{1 + \frac{b_1 - c_1}{b_0 - c_0}}. \tag{2}$$

The Principal's Objective.– The principal obtains a payoff of 1 if the regime survives and 0 otherwise. If $\theta < 0$, the regime cannot survive. Otherwise, the regime is viable, which means the regime survives if no one attacks it. However, a viable regime may not survive if agents attack. We call the ex-ante probability that a viable regime may not survive - the *coordination risk*. The principal seeks to minimize this coordination risk. If agents never attack a viable regime, then this risk is eliminated. However, an agent will not attack only if he believes that the regime will survive with a probability of at least p. We denote p the reluctance of the agents. Note that when attacking is the payoff-dominant action ($b_1 < b_0$), the principal's interest is not aligned with that of the agents (and the agents are more reluctant). We discuss this in Section V.

A diffused policy J.– The principal, who can be thought of as an information designer, can disclose some information over time based on the (exogenous) fundamental as well as the (endogenous) history of actions. We focus on a particular dynamic information disclosure policy. Consider a test that checks whether the regime continues to be viable at any date t – i.e., if it can survive in the absence of further attacks – and disclose the test result publicly. We call such tests - viability tests, and a positive test result - the public news of continued viability - **PNV**. Before agents move and the fundamental θ is realized, the principal chooses the frequency at which such viability tests should run. The policy is announced publicly. Without loss of generality, we only consider regular intervals between tests.

Definition 1 *A diffused policy J conducts the viability test at a regular interval of $\alpha = 1/J$ starting at 0 – i.e., at $(j-1)\alpha$ for $j = 1, 2, \ldots, J$.*

The jth test is conducted at time $(j-1)\alpha$ for any $j = 1, 2, \ldots, J$. The agents in $[(j-1)\alpha, j\alpha)$ move between the jth test and the next test. We refer to these α mass of agents as group j. The policy J separates the agents into J different groups based on the latest public news from the viability tests they can receive. A more diffused policy indicates the increased frequency of viability tests (higher J), or equivalently the group size α is smaller.

Cutoff Equilibrium.– Under a diffused policy J, the group j agents obtain public information on whether the regime has passed the jth viability test before taking action. Furthermore, each agent has a private signal regarding the underlying fundamental. If the regime fails a viability test, then the regime will not survive regardless of the remaining agents' actions. Hence, attacking is the dominant strategy for the group j agents. The strategic decision is non-trivial only when the regime passes the test, which is the focus of our analysis. In what follows, we limit our attention to the Perfect Bayesian Equilibrium in monotone strategies such as the case after agents learn that the regime has passed the viability test, the probability that an agent i attacks $\rho_i(s_i)$ is non-increasing in the private signal s_i. When the agents follow monotone strategies, a larger mass of agents attack against a weaker regime. Therefore, a stronger regime is more likely to survive. Under log-concavity of f, agent i believes the regime is strictly more likely to survive when s_i is higher. Thus, in equilibrium, all agents in the same group j (for any j) must follow a symmetric cutoff strategy: attack if and only if $s_i < \hat{s}_j$.[10] We refer to such equilibrium as cutoff equilibrium.

As the game continues, more agents attack. Hence, there exist a non-decreasing sequence of cutoff $\{\underline{\theta}_{j-1}\}_{j=1}^{J}$ such that the regime passes the jth viability tests if and only if $\theta \geq \underline{\theta}_{j-1}$. Let us define $\hat{\theta}$ such that if $\theta \geq \hat{\theta}$, then the regime passes all viability tests and survives. If there were a $(J+1)$th viability test after group J, then $\underline{\theta}_J$ would be equal to $\hat{\theta}$.

Persuasion.– In the spirit of Bayesian Persuasion, one can interpret this policy as a recommendation by the principal to the agents to "not attack" when the regime passes the test and "attack" when it does not.

Definition 2 *A policy is* ***persuasive*** *if there is no cutoff equilibrium, in which an agent attacks a regime that passes the latest test.*

If a diffused policy is persuasive, then the only cutoff equilibrium is that in which agents ignore their private information and follow the principal's recommendation. This means that no agent attacks a viable regime and thus, any regime with underlying fundamental $\theta \geq 0$ survives. In other words, under a persuasive policy, the only equilibrium cutoff fundamental is $\hat{\theta} = 0$, and the coordination risk is eliminated.

II. One-time Viability Test

Consider a degenerate policy - in which no tests are conducted. Since agents do not have any information about past actions of other agents, the game is essentially a simultaneous move regime change game. It follows from Morris and Shin (2003) that if the private signals are sufficiently precise, then there is a unique cutoff equilibrium. Under the uniform prior, one can explicitly solve for the unique cutoff signal as $s^* = p + \frac{1}{\sqrt{\tau}}F^{-1}(p)$. This means that the regime survives the attacks if and only if $\theta \geq p$. Recall that an agent can be dissuaded

[10] Suppose an agent in group j randomizes upon receiving a signal $\hat{s}_j$. Then, the agent is indifferent between attacking and not attacking. Therefore, any agent i in group j must play $\rho_i(s_i) = 0$ for $s_i > \hat{s}_j$ and $\rho_i(s_i) = 1$ for $s_i < \hat{s}_j$. We follow the convention that $\rho_i(\hat{s}_j) = 0$ (This does not play a significant role in our results).

from attacking if he believes the regime will survive with probability higher than p. Thus, it is intuitive that when p is higher, the ex-ante chance of survival is lower.

To understand the effects of the viability test, let us first consider the policy $J = 1$ – i.e., a one-time viability test. The regime passes the test when $\theta \geq 0$. The agents have different beliefs over θ based on their private signals. Thus, PNV will not affect all agents in the same fashion. The agents with private signal $s_i \in [-\frac{\sigma}{2}, \frac{\sigma}{2})$ once believed that θ could be less than 0, but the public news tells them otherwise. This makes them more optimistic about the strength of the regime, and thus less likely to attack. Owing to the strategic complementarity, other agents – including those who already know that $\theta \geq 0$ from their private information, become more optimistic about the success of the regime; hence, they are also less likely to attack.[11] It is intuitive that after learning that the regime is viable, in equilibrium, the agents will attack less aggressively. However, the news may not be strong enough, and there can be an equilibrium in which an agent attacks the regime that has passed the viability test (as in Goldstein and Huang (2016)).

Consider a hypothetical game that exactly reflects that presented above with one exception – it is played between some $\alpha < 1$ mass of agents instead of 1 mass of agents. The smaller the α, the more likely it is that $\theta \geq \alpha$; in other words, even if all the α mass of agents attack, the regime will survive. However, agents may receive private signal $s_i < \alpha - \sigma/2$ and believe that the regime will not survive if others attack. To understand whether such agents can be persuaded not to attack, we need to look into their beliefs about others in the equilibrium. Below, we investigate how the equilibrium strategy is affected when a small mass of agents learns that the regime has passed the viability test.

Persuasive Viability Test

Although the positive news $\theta \geq 0$ may not be strong enough to persuade a mass 1 of agents, the following Lemma shows that a one-time viability test can be persuasive when the mass of agents is sufficiently small. We emphasize this result because it plays a crucial role in deriving our main result.

Lemma 1 *Suppose that there is only one group with size α and they learn PNV. There exists $\alpha^* > 0$ such that there is no cutoff equilibrium in which an agent attacks a viable regime if and only if $\alpha < \alpha^*$.*

PROOF:

Suppose that agents follow some cutoff strategy - $\hat{s}$. Then, for any θ, the aggregate attack is $\alpha\mathbb{P}(s < \hat{s}|\theta) = \alpha F(\sqrt{\tau}(\hat{s} - \theta))$. Define $A(\hat{s}, \alpha) \in [0, \alpha]$ such

[11]Note that PNV makes the agents more optimistic about the success of the regime regardless of whether attacking is the payoff- dominated action. We revisit this issue later in Section V.

that

$$(A^\alpha) \qquad \alpha F(\sqrt{\tau}(\hat{s} - A(\hat{s},\alpha))) = A(\hat{s},\alpha).$$

By definition, the regime survives if and only if $\theta \geq A(\hat{s},\alpha)$. We refer to this as the Aggregate Condition (A^α).

After learning PNV ($\theta \geq 0$), an agent with private signal s_i believes that the regime will survive with probability

$$(3) \qquad \mathbb{P}(\theta \geq A(\hat{s},\alpha)|s_i, \theta \geq 0) = \frac{\int_{A(\hat{s},\alpha)}^{\overline{\theta}} \pi(\theta) f(\sqrt{\tau}(s_i - \theta))d\theta}{\int_0^{\overline{\theta}} \pi(\theta) f(\sqrt{\tau}(s_i - \theta))d\theta}.$$

It follows from log-concavity of f that this probability is higher for the agents with higher private information s_i (See the online appendix for the formal proof). Define $I^v(\hat{\theta})$ such that if the survival criteria is $\theta \geq \hat{\theta}$, the agent who receives private signal $s_i = I^v(\hat{\theta})$ is indifferent between attacking and not attacking.

$$(I^v) \qquad \frac{\int_{\hat{\theta}}^{\overline{\theta}} \pi(\theta) f(\sqrt{\tau}(I^v(\hat{\theta}) - \theta))d\theta}{\int_0^{\overline{\theta}} \pi(\theta) f(\sqrt{\tau}(I^v(\hat{\theta}) - \theta))d\theta} = p.$$

We refer to this as the Indifference Condition (I^v).

Suppose that all agents are following the cutoff strategy $\hat{s}$ and consequently, the survival criterion is $\hat{\theta} = A(\hat{s},\alpha)$. By definition, any agent with $s_i = \beta(\hat{s}) := I^v(A(\hat{s},\alpha))$ is indifferent between attacking and not attacking. Hence, a cutoff strategy $\hat{s}$ can constitute an equilibrium if $\hat{s} = \beta(\hat{s})$. If that is the case, we have

$$\frac{\int_{A(\hat{s},\alpha)}^{\overline{\theta}} \pi(\theta) f(\sqrt{\tau}(\hat{s} - \theta))d\theta}{\int_0^{\overline{\theta}} \pi(\theta) f(\sqrt{\tau}(\hat{s} - \theta))d\theta} = p.$$

Finally, substituting $\hat{s}$ from the Aggregate Condition (A^α) in the above, we get

$$\frac{\int_{A(\hat{s},\alpha)}^{\overline{\theta}} \pi(\theta) f(\sqrt{\tau}(A(\hat{s},\alpha) - \theta) + F^{-1}(\frac{A(\hat{s},\alpha)}{\alpha}))d\theta}{\int_0^{\overline{\theta}} \pi(\theta) f(\sqrt{\tau}(A(\hat{s},\alpha) - \theta) + F^{-1}(\frac{A(\hat{s},\alpha)}{\alpha}))d\theta} = p.$$

The left hand side captures the belief of the marginal agent, given he learns PNV, that - "given that α fraction of agents play the cutoff strategy $\hat{s}$, the regime has a per capita fundamental strength that surpasses the required cutoff $\frac{A(\hat{s},\alpha)}{\alpha}$, and thus survives." We define $G : [0,1] \times [0,1] \to [0,1]$ as follows

$$\text{(G)} \qquad G(x, \alpha) := \frac{\int_{\alpha x}^{\overline{\theta}} \pi(\theta) f(\sqrt{\tau}(\alpha x - \theta) + F^{-1}(x)) d\theta}{\int_{0}^{\overline{\theta}} \pi(\theta) f(\sqrt{\tau}(\alpha x - \theta) + F^{-1}(x)) d\theta}.$$

Note that $G(x = \frac{A(\hat{s},\alpha)}{\alpha}, \alpha)$ represents the belief of the marginal agent. By definition, $G(x, \alpha = 0) = 1$ for any $x \in [0,1]$, and $G(x = 0, \alpha) = 1$ for any $\alpha \in [0,1]$. For any $x \in (0,1]$, $\lim_{\alpha \to 0} G(x, \alpha) = 1$, i.e., the belief of the marginal agent point-wise converges to 1 as $\alpha \to 0$. Since G is a continuous function defined over a compact set, it is uniformly continuous (Heine-Cantor Theorem). Therefore, for any $\varepsilon > 0$, there exists $\delta > 0$, such that for any $(x, \alpha), (x', \alpha') \in [0,1] \times [0,1]$ with $||(x, \alpha) - (x', \alpha')|| < \delta$,

$$|G(x, \alpha) - G(x', \alpha')| < \varepsilon.$$

Consider any $x \in [0,1]$, fix $x' = x$ and $\alpha' = 0$, we have $G(x, \alpha) > 1 - \varepsilon$ for all $\alpha < \delta$. Thus, $G(x, \alpha)$ uniformly converges to 1 (not only point-wise) when α converges to 0. In other words, there exists $\tilde{\alpha} > 0$ such that whenever $\alpha < \tilde{\alpha}$, $G(x, \alpha) > p$ for all $x \in [0,1]$.

Define $\mathcal{P} := \{\tilde{\alpha} \in (0,1] | \forall x \in [0,1], \forall \alpha \in [0, \tilde{\alpha}), \ G(x, \alpha; \tau) > p\}$ and $\alpha^* := \sup \mathcal{P}$. It follows from the above arguments that $\mathcal{P} \neq \emptyset$ and $\alpha^* > 0$, and by construction, the set $\mathcal{P}$ is bounded above by 1, so the supremum exists.

Claim 1 *For any $\alpha \geq \alpha^*$, there exists $x \in [0,1]$ such that $G(x, \alpha) \leq p$*

The proof of the above claim is relegated to the appendix. From the above argument, it follows that there exists α^* such that whenever the mass of agents playing the game is less than α^*, given any cutoff strategy $\hat{s}$, the marginal agent with signal $\hat{s}$ believes that the regime will survive with probability strictly higher than p. Thus, the agent strictly prefers not to attack – i.e., $\beta(\hat{s}) < \hat{s}$. This implies that for any possible $\hat{s}$, the cutoff strategy $\hat{s}$ cannot constitute an equilibrium. However, if $\alpha \geq \alpha^*$, then it follows from continuity of $G(x, \alpha)$ that there exists a cutoff strategy for which $\beta(\hat{s}) = \hat{s}$. $\square$

There are two forces influencing this result: (1) the public news of viability and (2) a small mass of agents.

The news effect.– As we have already highlighted, PNV makes agents more optimistic about the survival of the regime. For any survival criteria $\hat{\theta}$, the private signal realization $I^v(\hat{\theta})$ that makes agents indifferent between attacking and not attacking, is lower than that in the absence of any viability test. To observe this, consider the signal $I(\hat{\theta})$ that makes an agent indifferent between attacking and not attacking when there is no viability test. $I(\hat{\theta})$ must satisfy

$$(I) \qquad \frac{\int_{\hat{\theta}}^{\overline{\theta}} \pi(\theta) f(\sqrt{\tau}(I(\hat{\theta}) - \theta)) d\theta}{\int_{\underline{\theta}}^{\overline{\theta}} \pi(\theta) f(\sqrt{\tau}(I(\hat{\theta}) - \theta)) d\theta} = p.$$

It follows from log-concavity of f that the LHS of condition (I) (or (I^v)) is increasing in the required cutoff signal $I(\hat{\theta})$ (or $I^v(\hat{\theta})$). Compared with condition (I^v), the denominator on the LHS of condition (I) is larger. Thus, it is easy to see that the required cutoff signal $I^v(\hat{\theta}) \leq I(\hat{\theta})$. By the same logic, one can easily check that the required cutoff signal $I^v(\hat{\theta})$ and $I(\hat{\theta})$ increase with the required fundamental strength $\hat{\theta}$.

The group size effect.– It directly follows from the Aggregate Condition (A^α) that a small group size α would make the success criterion $A(\hat{s}, \alpha)$ lower, given any $\hat{s}$.

The combined effect.– Combining the two forces, we can say that for any $\hat{s}$, a smaller group size α translates into a lower $A(\hat{s}, \alpha)$ and in that case, PNV lowers $\beta(\hat{s}) = I^v(A(\hat{s}, \alpha))$, which is the cutoff signal that makes an agent indifferent between attacking and not attacking. If the group size α is sufficiently small, this combined effect is so significant that for any $\hat{s}$, $\beta(\hat{s}) < \hat{s}$.

Figure 1 explains this combined effect graphically. We use the uniform prior and triangle probability density of error - $\hat{f}(x) := (2+4x)\mathbf{1}(-0.5 \leq x < 0) + (2-4x)\mathbf{1}(0 \leq x \leq 0.5)$.[12] Consider any cutoff per capita fundamental x. As defined in Equation (G), $G(x, \alpha)$ is the belief of the marginal agent that the regime will survive. Figure 1 plots this belief against any candidate cutoff $x \in [0, 1]$. Under a degenerate policy $G(x, \alpha)$ is the 45° line and hence at the cutoff per capita fundamental $x = p$, the agent with the cutoff signal is indifferent. Thus, under a uniform prior, in the absence of PNV, a small group size does not affect the equilibrium per capita fundamental cutoff.

We can see that PNV pushes this belief upward. As α decreases, $G(x, \alpha)$ increases. However, under the general prior, this monotonicity may not hold. Nevertheless, since $G(x, \alpha)$ uniformly converges to 1, there is an α^* such that for $\alpha < \alpha^*$, $G(x, \alpha) > p$ for all x. This implies that the only possible cutoff solution is $x = 0$ and the implied unique equilibrium is that in which no one attacks a viable regime.

A Lower Bound.– In the following section, we consider frequent viability tests that induce a dynamic setting. Constructing the belief of the marginal agent in this dynamic setting will be more involved. Below, we construct a lower bound $\underline{G}(x, \alpha)$ on the belief $G(x, \alpha)$ of the marginal agent. This lower bound will be

[12] Note that the error distribution $\hat{F}$ is log-concave, and under a uniform prior, this is enough to guarantee that for any $A > B$, the posterior belief $\mathbb{P}(\theta \geq A | s, \theta \geq B)$ is increasing in s.

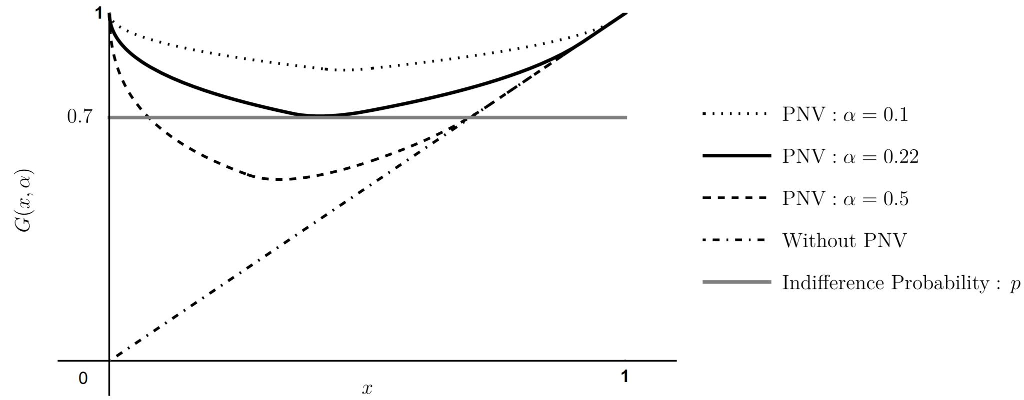

Figure 1. : $G(x, \alpha)$ when $\{\underline{\theta}, \bar{\theta}\} = \{-1, 2\}$, $\tau = 1$, $\pi(\theta)$ is Uniform and $f = \hat{f}$.

useful in the dynamic setting. Let us define

$$(\underline{G}) \qquad \underline{G}(x,\alpha) := 1 - \left(\frac{\overline{\pi}\overline{f}}{\underline{\pi}\underline{f}}\right)\left(\frac{\alpha x}{\alpha x + \frac{1}{\sqrt{\tau}}\left(F^{-1}(x) + \frac{1}{2}\right)}\right)$$

for $x > 0$ and $\underline{G}(0, \alpha) = 1$. It is clear that for any $x > 0$, [13]

$$G(x,\alpha) \geq 1 - \frac{\int_0^{\alpha x} \overline{\pi}\overline{f} d\theta}{\int_0^{\alpha x + \frac{1}{\sqrt{\tau}} F^{-1}(x) + \frac{1}{2\sqrt{\tau}}} \underline{\pi}\underline{f} d\theta} = \underline{G}(x,\alpha)$$

and $G(x = 0, \alpha) = \underline{G}(x = 0, \alpha) = 1$.

Corollary 1 *There exists $\hat{\alpha} \in (0, \alpha^*]$ such that whenever $\alpha < \hat{\alpha}$, $G(x, \alpha) \geq \underline{G}(x, \alpha) > p$ for all $x \in [0, 1]$*

The Corollary 1 directly follows from the same argument as in Lemma 1 since $\underline{G}(x, \alpha)$ also uniformly converges to 1 as $\alpha \to 0$. $\hat{\alpha} \leq \alpha^*$ because $\alpha < \alpha^*$ serves as the necessary and sufficient condition for Lemma 1 while $\alpha < \hat{\alpha}$ is only the sufficient condition. In other words, $\hat{\alpha}$ qualifies as a group size that is small enough to be dissuaded from attacking the regime by a viability test.

III. Main Result: Frequent Viability Tests

We understand that positive viability news can dissuade a sufficiently small mass of agents from attacking. Now, let us return to the frequent viability tests

[13]Note that, for any cutoff strategy $\hat{s} \in [-\frac{1}{2\sqrt{\tau}}, 1 + \frac{1}{2\sqrt{\tau}})$, the effective upper bound for the intergral is $\hat{s} + \frac{1}{2\sqrt{\tau}}$ instead of $\bar{\theta}$.

- a diffused policy J. There can be multiple cutoff equilibria. We relegate the full characterization of cutoff equilibria for any diffused policy J to the online appendix. The following numerical example shows all possible cutoff equilibria for $J = 2$.

Numerical Example 1 *In Table 1, the first equilibrium is the one in which no agent attacks a viable regime* ($\hat{s}_1 = \hat{s}_2 = -\frac{\sigma}{2}$)*. Thus,* $\hat{\theta} = 0$*. Similar to the one time viability test, this remains a possible equilibrium outcome. There are two equilibria (*2 *and* 3*), in which no agent in group* 2 *attacks a regime if the regime passes the* 2*nd viability test* ($\hat{s}_2 = \underline{\theta}_1 - \frac{\sigma}{2}$)*. Thus, if the regime passes the second viability test, it survives* ($\underline{\theta}_1 = \hat{\theta}$)*. It is also possible (equilibrium* 4 *and* 5*) that agents in both group* 1 *and* 2 *attack a viable regime.*

Table 1—: Cutoff Equilibria for $J = 2$

Equilibrium	$\hat{s}_1$	$\hat{s}_2$	$\underline{\theta}_1$	$\hat{\theta}$
1	−0.50	−0.50	0	0
2	−0.26	−0.46	0.04	0.04
3	0.46	−0.15	0.35	0.35
4	0.77	0.72	0.46	0.66
5	0.57	0.29	0.39	0.46

Note: Parameter Values: $\tau = 1$, $p = 0.7$, π is Uniform in $[-1, 2]$ and $f = \hat{f}$

It is easy to see that for any diffused policy J, one can always construct a cutoff equilibrium in which for some $j = 1, 2, \ldots, J$, no agent attacks the regime after the regime passes the jth viability test. Thus, the regime that passes the jth viability test continues to be viable until the end. However, if an agent believes that others may attack a regime even after it passes the viability tests, then he may do the same – particularly if this agent has received a very low private signal; hence, multiple cutoff equilibria may arise. Nevertheless, we argue that when the principal adopts a sufficiently diffused policy, the only possible cutoff equilibrium is the one in which all agents follow the principal's recommendation and do not attack a viable regime.

Theorem 1 *A diffused policy* J *with* $J > \hat{J} = \frac{1}{\hat{\alpha}}$ *is persuasive.*

Consider the case in which group j agents learn that the regime has passed the jth viability test and decide whether to attack the regime. When $J > \hat{J}$, there are less than $\hat{\alpha}$ mass of agents in group j. Recall from Corollary 1 that a viability test can be persuasive if the mass of agents is smaller than $\hat{\alpha}$. However, there are two crucial differences in this dynamic case. First, whether the regime can pass the jth viability test depends on θ as well as past attacks. Thus, the interpretation of the positive viability news is history-dependent. To observe this, note that the viability test discloses whether $\theta \geq \underline{\theta}_{j-1}$, where $\underline{\theta}_{j-1}$ is endogenous. Second, although only α mass of agents are moving, $(1 - j\alpha)$ mass of agents will move later, and the regime may fail if agents in the subsequent groups attack. Thus, the effectiveness of jth viability test is dependent upon the effectiveness of future viability tests.

However, we build on Lemma 1 and Corollary 1, and argue that the following statement is true for any j – in particularly for $j = 0$.

M_j: When the agents in any group $j' > j$ learn that the regime has passed the j'th viability test (regardless of whichever cutoff strategy was played by others in the past), any cutoff strategy in which an agent in group j' may attack, violates sequential rationality.

As we only consider cutoff strategies, we sometimes refer to the above statement as – no agent in group $j' > j$ attacks a regime that passes the j'th viability test. The theorem claims that when $J > \hat{J}$, M_0 is true. We prove this using the induction argument - regardless of history, the agents in any group j follow the principal's recommendation, if they believe that the agents in the subsequent groups will do the same. Formally speaking,

N_j: If M_j is true, then M_{j-1} is true.

Step A

Lemma 2 *If $\alpha < \hat{\alpha}$, then N_1 is true.*

It directly follows from Corollary 1 that when $\alpha < \hat{\alpha}$, for any possible cutoff strategy $\hat{s}_1$, the marginal agent with private signal $s_i = \hat{s}_1$ believes that the regime will withstand the current attacks with a probability strictly higher than p. If the regime withstands the current attacks, then it will pass the next test. Given that M_1 is true, this means that the agent believes the regime will survive with a probability strictly higher than p. Hence, the agent prefers not attacking. Thus, any cutoff strategy in which an agent in group 1 attacks the regime that passes the first viability test, violates sequential rationality – i.e., M_0 is true.

Step B

The difference between the agents in the first group and those in the later groups is that the result of the viability test for group $j > 1$ is history-dependent. However, the following lemma shows that $\underline{G}(x, \alpha)$ (as defined in the equation ($\underline{G}$)) serves as a lower bound to the belief of the marginal agent in group j as well regardless of $\underline{\theta}_{j-1}$. Thus, given that M_j is true, independent of history, i.e., the cutoff strategies played by the agents before group j, the positive viability news is also persuasive for agents in group j.

Lemma 3 *If $\alpha < \hat{\alpha}$, then N_j is true for any $j > 1$.*

PROOF:

Suppose that an agent in group j believes that the agents in group $l < j$ have played some cutoff strategy $\hat{s}_l$ for $l = 1, 2, \ldots, (j-1)$. Then, the regime will pass the jth viability test if

$$\sum_{l=1}^{j-1} \alpha F(\sqrt{\tau}(\hat{s}_l - \theta)) \leq \theta. \tag{4}$$

For any cutoff strategies the agents may have played in the past, PNV means $\theta \geq \underline{\theta}_{j-1}$, where $\underline{\theta}_{j-1}$ solves the above condition (4) with equality.[14]

Suppose that the agents in group j, follow a cutoff strategy $\hat{s}_j$. Let us define $A_j(\hat{s}_j, \alpha, \underline{\theta}_{j-1})$ such that

$$A_j(\hat{s}_j, \alpha, \underline{\theta}_{j-1}) = \underline{\theta}_{j-1} + \alpha F(\sqrt{\tau}(\hat{s}_j - A_j(\hat{s}_j, \alpha, \underline{\theta}_{j-1}))). \tag{A_j^α}$$

Unless $\theta \geq \underline{\theta}_{j-1}$, the regime will not pass the jth viability test. Given the cutoff strategy $\hat{s}_j$, the second term on the RHS captures the aggregate attack from group j when $\theta = A_j$. Based on the condition (A_j^α), if $\theta \geq A_j$, the regime will pass the jth viability test and then sustain the attack from group j.[15]

Given M_j is true, this means whenever $\theta \geq A_j$, the regime survives.

The marginal agent in group j with signal $\hat{s}_j$ believes that the regime will survive with probability at least

$$\mathbb{P}(\theta \geq A_j | \hat{s}_j, \theta \geq \underline{\theta}_{j-1}) = \frac{\int_{A_j}^{\overline{\theta}} \pi(\theta) f(\sqrt{\tau}(\hat{s}_j - \theta)) d\theta}{\int_{\underline{\theta}_{j-1}}^{\overline{\theta}} \pi(\theta) f(\sqrt{\tau}(\hat{s}_j - \theta)) d\theta}.$$

Substituting the Aggregate Condition (A_j^α), we can write the marginal agent's belief as

$$\frac{\int_{A_j}^{\overline{\theta}} \pi(\theta) f(\sqrt{\tau}(A_j - \theta) + F^{-1}(\frac{A_j - \underline{\theta}_{j-1}}{\alpha})) d\theta}{\int_{\underline{\theta}_{j-1}}^{\overline{\theta}} \pi(\theta) f(\sqrt{\tau}(A_j - \theta) + F^{-1}(\frac{A_j - \underline{\theta}_{j-1}}{\alpha})) d\theta}.$$

Substituting the per capita required cutoff strength for agents in group j to pass

[14] Given the strategic complementarity, it follows from Milgrom and Roberts (1990) argument that the unique cutoff equilibrium in Lemma 1 means it is the unique rationalizable strategy. However, under the endogenous information disclosure, if the agents in the past were attacking more aggressively, then surviving those attacks means the regime is likely to be strong. This, in turn, gives an incentive to the agents to not attack, and thus, the strategic complementarity may be violated. Therefore, the same generalization to rationalizable strategy may not hold for Lemma 3, and the existence of non-monotone equilibria is an open question for future research.

[15] Note that $\theta \geq A_j$ is sufficient for the regime to pass the $(j+1)$th viability test. It is not necessary. If $\theta = A_j > \underline{\theta}_{j-1}$, then the aggregate past attack is strictly lower than $\underline{\theta}_{j-1}$ (unless $\underline{\theta}_{j-1} = 0$). Thus, at $\theta = A_j$, the aggregate attack is less than the right hand side of the equation (A_j^α).

the next test as $x = \frac{A_j - \underline{\theta}_{j-1}}{\alpha}$, we get that belief of the marginal agent with signal $\hat{s}_j = \frac{1}{\sqrt{\tau}} F^{-1}(x) + \alpha x + \underline{\theta}_{j-1}$ is

$$\frac{\int_{\underline{\theta}_{j-1}+\alpha x}^{\overline{\theta}} \pi(\theta) f(\sqrt{\tau}(\alpha x - \theta) + F^{-1}(x) + \sqrt{\tau}\underline{\theta}_{j-1}) d\theta}{\int_{\underline{\theta}_{j-1}}^{\overline{\theta}} \pi(\theta) f(\sqrt{\tau}(\alpha x - \theta) + F^{-1}(x) + \sqrt{\tau}\underline{\theta}_{j-1}) d\theta}. \tag{5}$$

Note that if $\underline{\theta}_{j-1} = 0$ (as in Lemma 2), then this belief is equal to $G(x, \alpha)$ as defined in Equation (G). Since π and f are bounded away from 0, for $x > 0$, the belief of the marginal agent $\hat{s}_j$ is at least

$$1 - \frac{\int_{\underline{\theta}_{j-1}}^{\underline{\theta}_{j-1}+\alpha x} \overline{\pi}\overline{f} d\theta}{\int_{\underline{\theta}_{j-1}}^{\underline{\theta}_{j-1}+\alpha x + \frac{1}{\sqrt{\tau}} F^{-1}(x) + \frac{1}{2\sqrt{\tau}}} \underline{\pi}\underline{f} d\theta} = 1 - \left(\frac{\overline{\pi}\overline{f}}{\underline{\pi}\underline{f}}\right)\left(\frac{\alpha x}{\alpha x + \frac{1}{\sqrt{\tau}}\left(F^{-1}(x) + \frac{1}{2}\right)}\right) = \underline{G}(x, \alpha),$$

while for $x = 0$, the belief of the marginal agent $\hat{s}_j$ is 1.

Therefore, for $\alpha < \hat{\alpha}$ (as in Corollary 1), given M_j is true and given any possible fundamental cutoff $\underline{\theta}_{j-1}$, for any cutoff strategy $\hat{s}_j$, the marginal agent with cutoff signal $\hat{s}_j$ believes that the regime survives (or $\theta \geq A_j(\hat{s}_j, \alpha, \underline{\theta}_{j-1})$) with probability strictly higher than p. This implies M_{j-1} is true. □

Step C

M_J is trivially true since there is no attack after the last group. If the mass of agents in each group satisfies $\alpha < \hat{\alpha}$, by induction, we can show that in any cutoff equilibrium, no agent will ever attack a viable regime. When the principal runs the viability tests with sufficient frequency, the agents are assured that no agent will attack when the regime passes the next viability test (M_j is true). This makes the current positive viability news persuasive (N_j is true). Thus, under a sufficiently diffused policy ($J > \hat{J}$), the coordination risk unravels from the end. This proves our main result in Theorem 1.

Frequent Viability Test vs Stress Test.– Inostroza and Pavan (2017) and Goldstein and Huang (2016) propose a one time "stress test" policy: the regime passes the stress test only if $\theta \geq k$ for some k. The authors demonstrate that if the stress test is sufficiently tough $k > \hat{k}$, then no agent will attack the regime that passes the stress test even in the cutoff equilibrium where agents attack most aggressively. In contrast, our paper considers frequent tests with the weakest possible strength or frequent viability tests. Although a one time viability test may not be sufficiently "tough," when the principal runs the tests with sufficient frequency, it can dissuade agents from attacking a viable regime.

Any viable regime survives when viability tests are repeated with sufficient frequency. In contrast, under a one-time stress test with $k > \hat{k}$, a viable regime whose strength is below k fails. Thus, from an ex-ante perspective, the principal benefits from adopting a sufficiently frequent viability test policy compared with a sufficiently tough one-time test policy when the agents move sequentially. Furthermore, unlike the viability test, the stress test policy violates ex post incentive compatibility; if the principal misreports that the regime has passed the test when $\theta \in [0, k)$, then the agents will not attack and the regime will survive. This is not so for viability tests; if the regime fails a viability test, then the principal cannot benefit from misreporting it.[16]

IV. Application : Asynchronous Debt Structure

Consider a standard maturity mismatch problem, in which a borrower finances a long-term illiquid investment project by short-term debt contracts. Creditors may not roll over the short-term debts if they believe others will not do so. Hence, borrowers with reasonable sound fundamental may fail due to coordination failure among creditors in rolling over their debt. Any financial institution that performs liquidity transformation – e.g., commercial banks, hedge funds, and mutual funds – could be exposed to such panic-based runs. Brunnermeier (2009) argues that this coordination risk resulting from the maturity mismatch is a primary cause of the recent financial crisis.

Asynchronous debt structures that prevent creditors from withdrawing simultaneously are fairly common (see He and Xiong (2012)). Choi, Hackbarth and Zechner (2017) provide empirical evidence that corporate bond issuers diversify debt rollovers across dates; they find that there is a significant increase in maturity dispersion for firms facing high rollover risk. Hedge funds and mutual funds are also allowed to lift redemption gates to limit momentary liquidity outflows.[17] We are seeking a theoretical rationale behind the asynchronous debt structure.

Let us consider a stylized model in which there is a unit mass of creditors, and each of them has lent 1 unit to the borrower. The borrower uses this funding to finance some illiquid investment. An asynchronous debt structure J ensures only $\alpha = \frac{1}{J}$ proportion of debt matures at time $(j-1)\alpha$ for $j = 1, 2, \ldots, J$. If the borrower fails to service the debt at any of these maturity dates, she must go through the bankruptcy process, which is public information to all creditors. This enforces the disclosure of borrower's viability at each maturity date. Hence, the debt structure J is equivalent to the policy J we introduced in Section I.[18]

[16]Thus, implementing a stress test policy requires commitment power as is standard in the Bayesian persuasion literature. Implementing viability tests do not require such commitment power. However, if repeating the tests are costly, then this may not be the case. It is important that the agents do not doubt that the principal will stop conducting the tests in the future.

[17]The redemption gates policy forces investors to make withdrawals asynchronously. Hedge fund managers can lift investor-level gates to limit investors' redemptions within a certain period. A common investor-level gate limits redemptions to 25% of an investor's money each quarter over four quarters (see Barr (2010)). On October 14, 2014, SEC Rule 2a-7 was amended to enable managers of Money Market Mutual Funds to set redemption gates within a certain period during which a fund's liquidity position is unfavorable.

[18]We do not model the ex-ante lending decision; rather, we focus on analyzing the rollover problem.

The asynchronous debt structure separates creditors into groups according to dispersed maturity dates. Group j creditors refers to the creditors who decide whether to withdraw ($a_i = 1$), or roll over the debt ($a_i = 0$) at time $(j-1)\alpha$. The return from the illiquid investment R is realized at time 1. θ indicates the liquidity position of the borrower. The borrower can use the liquid assets and the access to the extra funding source to sustain total withdrawals up to θ. [19]

The information environment is assumed to be the same as that described in Section I. The borrower chooses the debt maturity structure J prior to θ being realized. Thus, the debt structure itself does not provide information about θ.[20] All creditors know the maturity date of their debt as well as the overall debt structure. Let w_j be the proportion of group j creditors who decide to withdraw. Given the debt structure J, group j creditors know whether the borrower is still viable at time $(j-1)\alpha$, i.e., $\theta \geq \alpha \sum_{l=1}^{j-1} w_l$.

If the borrower cannot service the withdrawal before time $(j-1)\alpha$, she defaults and all creditors whose debt matures after that time have no chance to make the withdrawal and obtain 0. Below, we describe the payoff for group j creditors when the borrower is still viable at time $(j-1)\alpha$.

By rolling over the debt, a group j creditor gets $1 + r < R$ if the borrower can service all the withdrawals up to time 1; otherwise, the borrower obtains 0.

$$u(0, \theta, \{w_l\}_{l=1}^{J}) = \begin{cases} 1 + r & \text{if } \theta \geq \alpha \sum\limits_{l=1}^{J} w_l \\ 0 & \text{if } \theta < \alpha \sum\limits_{l=1}^{J} w_l. \end{cases}$$

Upon withdrawal, the creditor gets his principal back at time $(j-1)\alpha$ if the borrower can sustain the withdrawal from group j, i.e., $\theta \geq \alpha \sum_{l=1}^{j} w_l$. Otherwise, the borrower defaults instantly, and the withdrawing creditors in group j split the remaining liquid assets θ_j.[21] Thus, the payoff for group j creditors from withdrawal is

$$u(1, \theta, \{w_l\}_{l=1}^{j}) = \begin{cases} 1 & \text{if } \theta \geq \alpha \sum\limits_{l=1}^{j} w_l \\ \frac{\theta - \alpha \sum\limits_{l=1}^{j-1} w_l}{\alpha w_j} & \text{if } \theta < \alpha \sum\limits_{l=1}^{j} w_l. \end{cases}$$

However, ceteris paribus, if the borrower is less likely to default, then the creditors will be more willing to lend in the first place. Thus, such modification will only reinforce our result.

[19] The determination of the borrower's liquidity position θ is outside of the model. θ may be negative if the liquidity outflow is high, due to the borrower's business operations and/or derivative positions, for example.

[20] If the borrower only undertakes such a debt structure in times of distress, then this may defeat the purpose of the policy.

[21] For simplicity, we assume that the long-term illiquid investment has no liquidation value before it matures and the borrower cannot borrow against the return R.

This payoff specification is different from that which is introduced in Section I. The payoff from rolling over depends on whether the borrower can sustain all withdrawals to time 1, while that from withdrawing only depends on whether the borrower can withstand the withdrawal from the current group. This asymmetry in payoff may increase the incentive for creditors to withdraw. Moreover, unlike in Section I, the payoff from withdrawing is not constant when the borrower defaults. This payoff specification is similar to that presented by Goldstein and Pauzner (2005); their specification indicates that the payoff from withdrawing is negatively dependent on aggregate withdrawal.[22]

The following proposition shows that, despite these differences, when the borrower adopts a sufficiently asynchronous debt structure, the public information of borrower's continued viability – i.e., success in rolling over debt maturing early – can avert panic-based runs altogether.

Proposition 1 *A sufficiently asynchronous debt structure eliminates the risk of panic-based debt runs.*

Our result is robust to the case when payoff from attacking (or withdrawal) is independent of future attacks because we prove our main result using the statement M_j inductively – i.e., no one withdraws if the borrower remains viable after group j creditors had made their decisions. Thus, the asymmetry disappears since under M_j; so long as the borrower can service the current withdrawal, no future withdrawal will occur ($w_l = 0$ for all $l > j$). Accommodating a general payoff structure broadens the scope of applying our proposed policy. Although the payoffs depend on both the fundamental and withdrawals, we can always find a threshold $\overline{p} < 1$ such that a creditor never withdraws if he believes that probability of no default is higher than $\overline{p}$ so long as the payoffs are bounded. Thus, a sufficiently asynchronous debt structure with $J > 1/\hat{\alpha}(\overline{p}, \tau)$ eliminates the chance of panic-based runs.

V. Discussion

Applying our theory to debt run provides a rationale for why borrowers often adopt asynchronous debt structures in practice. The insight also applies to regime change game in general (under more general payoff structure than that specified in Section I). In this section, we consider some variations of our set up to understand what could make it more or less difficult to dissuade agents from attacking. Moreover, to identify the essential features of the set up without which such persuasion may not be possible. First, let us consider the policy in which the principal discloses all information to the agents.

[22]We adopt this payoff specification for simplicity. We prove our result for a more general payoff structure. One can easily accommodate any reasonable payoff specification for withdrawing and rolling over when the borrower defaults.

Suppose that the principal knows the underlying fundamental and discloses the true θ at the beginning. Furthermore, she discloses the information regarding past attacks with high frequency. Under such environment, when $\theta < 1$, it is easy to see that there exists an equilibrium in which all agents attack regardless of the history. Hence, any regime with $\theta < 1$ cannot survive.

Note that in our setting, an agent cannot make a difference in the aggregate attack by unilateral deviation. If this is not the case, full disclosure will lead to a very different result. Let us assume that attacking is the payoff-dominant action – i.e., $b_1 < b_0$. This means the agents' incentives are not aligned with the principal, and this is possible in practice. Consider, for example, a currency attack game as that presented in Morris and Shin (1998): attacking is costly but would be profitable if agents can coordinate ($b_0 > 0, c_1 < 0$), while not attacking has zero payoff ($b_1 = c_0 = 0$). In this case, attacking the regime of fixed exchange rate is the payoff-dominant action, but the policy designer may have incentive to defend it.

Consider the following simple example. Two agents are moving sequentially, one agent is equivalent to half mass, and the principal discloses all the information about θ and the past attack, regardless of whether the first agent has attacked. Suppose that $\theta \in (1/2, 1)$.[23] Then, the second agent will attack if the first agent attacks ($b_0 > b_1$) and will not attack if the first agent does not attack ($b_1 > c_1$). Since attacking is the payoff-dominant action ($b_0 > b_1$), the first agent attacks. Thus, it follows from backward induction that both agents will take the payoff-dominant action - attack. This simple example demonstrates that full disclosure may fail to dissuade agents from attacking. Thus, even if full disclosure is a feasible policy, the principal may not want to adopt such a policy.

The result that agents play the payoff-dominant action even generalizes to repeated games. In contrast with the folk theorem results, Dutta (2012) (see also Lagunoff and Matsui (1997)) shows that if this coordination game is repeated (finitely but sufficiently many times), and agents alternately get chances to revise their decisions, then the agents will soon move to playing the payoff-dominant action regardless of the current state.

In contrast, our selective partial disclosure policy is persuasive regardless of whether attacking is the payoff-dominant action. If attacking is the payoff-dominant action, then it may be more difficult to dissuade the agents from attacking. Let us now move to the comparative statics and examine the factors that makes persuasion more difficult.

[23]The argument for $\theta < 1/2$ is straight forward because attacking is the dominant action for the first agent.

Comparative Statics

We have shown that a sufficiently diffused policy is persuasive. However, whether a diffused policy J qualifies as a sufficiently diffused policy or not depends on the model parameters. In our simple model, two parameters in which we are interested are p (reluctance) and τ (precision). We may expect that if the agents are more reluctant to follow when the principal recommends that they do not attack, it will be more difficult to persuade them. Furthermore, if the agents have more precise signals, it will be more difficult to persuade them to ignore their private signals.

It is easy to see that $\underline{G}(x, \alpha; \tau)$ is decreasing in α and τ. This in turn implies that $\hat{\alpha}(p, \tau)$ is decreasing in p and τ.[24] However, this does not exactly capture that the principal must diffuse more when p or τ is higher. It is because $\hat{J}$ is only a sufficient threshold – i.e., even if $J < \hat{J}(p, \tau)$, the diffused policy can be persuasive. However, under the uniform prior, we can provide a tight bound, i.e., a necessary and sufficient threshold for J. The following proposition describes the comparative statics results under the uniform prior.

Proposition 2 *Under the uniform prior, there exists $J^*(p, \tau)$ such that a diffused policy is persuasive if and only if $J > J^*(p, \tau)$, where*

1) for $p' \geq p$, $J^*(p', \tau) \geq J^*(p, \tau)$, and

2) for $\tau' \geq \tau$, $J^*(p, \tau') \geq J^*(p, \tau)$.[25]

The result that $J^*(p, \tau)$ is not only sufficient but also necessary, follows from the fact that under the uniform prior, the belief of the marginal agent in group j regarding the survival of the regime (as defined in (G)) is decreasing in the group size α for any given $\underline{\theta}_{j-1}$.

This implies that when p is higher, a smaller group size – or equivalently, higher diffusion is required to make $G(x, \alpha; \tau) > p$ for all $x \in [0, 1]$. If $b_1 < b_0$, then an agent gets a higher payoff from attacking a regime that does not survive, as opposed to not attacking a regime that does survive. Thus, attacking is the payoff-dominant action, and the principal's interest is not aligned with the agents (recall the currency attack example). If attacking is the payoff-dominant action ($b_1 < b_0$), then p is higher (see Equation 2). From the above proposition, it follows that it will be harder to dissuade the agents from attacking. Moreover, under the uniform prior, $G(x, \alpha; \tau)$ is also increasing in τ. Therefore, persuading the agents with more precise private signals to ignore their private information and follow the principal's recommendation requires higher diffusion.

[24]Consider, for example, $p' \geq p$. If $\alpha < \hat{\alpha}(p', \tau)$, then $\underline{G}(x, \alpha; \tau) > p' > p$ for all $x \in [0, 1]$. Hence, $\hat{\alpha}(p', \tau) \leq \hat{\alpha}(p, \tau)$.

[25]Given the support of the prior, we can only compare τ and τ', which are not too small, i.e., $\frac{1}{\sqrt{\tau'}} < \frac{1}{\sqrt{\tau}} < \min\{\overline{\theta} - 1, -\underline{\theta}\}$.

In our basic set up, we consider a regime change game in which agents have noisy private information as is standard in global game literature. Suppose that the noisy information was public; thus, agents share a homogeneous belief over θ. Below we argue that in such an environment, it is possible that no diffused policy can dissuade the agents from attacking.

Homogeneous Belief

Consider the same set up as in Section I with the following modification. The agents receive a noisy public signal $s = \theta + \sigma\epsilon$ instead of private signals. For simplicity, let us assume the prior is uniform. The principal observes the public signal and announces the policy after learning s. Conditional on the realized signal s, the agents share a homogeneous belief over θ. Under a homogeneous belief, the only way to dissuade agents from attacking is to convince them "even if others attack, the regime is very likely to survive." Otherwise, there is always a possible equilibrium in which all the agents attack.

If s is sufficiently high such that $\mathbb{P}(\theta \geq 1|s) > 0$, then the agents believe that θ can be in the upper dominance region. The positive viability news is able to increase this probability to $\mathbb{P}(\theta \geq 1|\theta \geq 0, s)$. The following proposition shows that if the tests are repeated with sufficient frequency (or $J > J^*(s)$), the probability that "θ is in the upper dominance region" may be high enough to dissuade the agents from attacking a viable regime. It is clear that the lower the s, the lower the $\mathbb{P}(\theta \geq 1|s)$; consequently, it becomes more difficult to dissuade the agents from attacking. Thus, $J^*(s)$ increases as s falls.

If s is sufficiently low, then $\mathbb{P}(\theta \geq 1|s) = 0$ – i.e., the agents have a common belief that $\theta < 1$. This implies that the agents share the common belief that if all the agents attack, then the regime cannot survive. Thus, regardless of the frequency of the viability tests, all attack is a possible equilibrium.

In contrast, under heterogeneous beliefs, there is a uniform J^* such that no agent attacks regardless of the private signals. Even the agent who receives a private signal $-\sigma/2$ and knows that $\theta = 0$ can be persuaded not to attack. The crucial difference with the heterogeneous belief case is as follows. Under heterogenous beliefs, when θ is low, all the agents may privately learn that $\theta < 1$, but it is not a commonly held belief that $\theta < 1$. Agent i may entertain the possibility that other agents believe $\theta \geq 1$, or even if others do not believe so, they may entertain the possibility that others believe $\theta \geq 1$ and so on. Thus, the principal may not be able to convince an agent who receives a low private signal such as $s_i = -\sigma/2$ that "even if others attack the regime is very likely to survive." However, the principal may still be able to convince the agent that "others are not likely to attack." Thus, the regime is likely to survive the attacks.

Proposition 3 *Under the information environment with the uniform prior and noisy public signal,* *(I) A diffused policy can dissuade the agents from attacking only if the noisy public signal s is sufficiently high such that $\mathbb{P}(\theta \geq 1|s) > 0$, and the required diffusion $J^*(s)$ increases when s decreases.*

(II) If s is sufficiently small such that $\mathbb{P}(\theta \geq 1|s) = 0$, then no diffused policy can dissuade the agents from attacking.

(III) From an ex-ante perspective, for any $\theta < 1$, there is a positive probability that (II) will happen.[26]

The result (I) and (II) follow from the argument preceding the proposition. For (III), note that when $\theta < 1$, there is always a positive probability that the signal is so low that it becomes a common belief that $\theta < 1$; thus, no diffused policy can dissuade the agents from attacking. There is always a possible equilibrium in which all agents attack. In this sense, a diffused policy cannot eliminate the coordination risk when agents have noisy public information. This shows that private information environment is essential for our result.

VI. Conclusion

This paper proposes a simple policy called frequent viability tests that may be used to eliminate the strategic uncertainty in a global game of regime change. The frequent viability tests diffuse the coordination risk that would otherwise be concentrated at one point in time. We show that when the principal sufficiently diffuses the coordination risk, agents ignore their private information and follow the principal's recommendation. The underlying mechanism is as follows. When the principal recommends the agent not to attack a regime that passes the viability test, and when the agents are assured that the agents in the subsequent groups will follow the principal's recommendation (statement M_j), then a sufficiently small group of agents will follow the recommendation as well (argument N_j).

This result contributes to dynamic information design literature. From a methodological perspective, our paper develops an inductive argument to show that the coordination risk unravels from the end. From an applied perspective, we show that a sufficiently asynchronous debt structure eliminates the possibility of panic.

Readers may wonder that, since the agents move sequentially, some exogenous shock may hit the fundamental or new information may arrive over time. In the online appendix, we detail the robustness of our main result to such perturbations in the basic model. Within the scope of this paper, we refrain from discussing the effectiveness of limited diffusion (when sufficient diffusion is not feasible) or optimal information disclosure under endogenous timing of attack. We believe that these are promising directions for future research.

[26] Under unbounded noise (for example, the Gaussian distribution), the impossibility in (II) and (III) go away. However, for any J (however large), we can always find s such that $J^*(s) > J$. In this sense, although a diffused policy can work for any possible realization of public signal, the required diffusion can be unrealistically high.

REFERENCES

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.

Banerjee, Abhijit V. 1992. "A Simple Model of Herd Behavior." *The Quarterly Journal of Economics*, 107: 797–817.

Barr, Alistair. 2010. "Hedge funds try new way to avoid big redemptions." `https://www.marketwatch.com/story/hedge-funds-try-new-way-to-avoid-big-redemptions-2010-06-10`.

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

Bergemann, Dirk, and Stephen Morris. 2019. "Information design: A unified perspective." *Journal of Economic Literature*, 57.1: 44–95.

Bikhchandani, Sushil, David Hirshleifer, and Ivo Welch. 1992. "A Theory of Fads, Fashion, Custom, and Cultural Change as Informational Cascades." *Journal of Political Economy*, 992–1026.

Brunnermeier, Markus K. 2009. "Deciphering the Liquidity and Credit Crunch 2007-2008." *The Journal of Economic Perspectives*, 23(1): 77–100.

Carlsson, Hans, and Eric Van Damme. 1993. "Global Games and Equilibrium Eelection." *Econometrica*, 61(5): 989–1018.

Choi, Jaewon, Dirk Hackbarth, and Josef Zechner. 2017. "Corporate Debt Maturity Profiles." `http://people.bu.edu/dhackbar/CHZ-2017.pdf`.

Cong, Lin William, Steven R Grenadier, and Yunzhi Hu. 2018. "Dynamic Interventions and Informational Linkages." *Journal of Financial Economics*, Forthcoming.

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

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

Dutta, Prajit K. 2012. "Coordination Need Not Be A Problem." *Games and Economic Behavior*, 76(2): 519–534.

Edmond, Chris. 2013. "Information Manipulation, Coordination, and Regime Change." *Review of Economic Studies*, 80(4): 1422–1458.

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

Frankel, David, and Ady Pauzner. 2000. "Resolving indeterminacy in dynamic settings: the role of shocks." *The Quarterly Journal of Economics*, 115(1): 285–304.

Goldstein, Itay, and Ady Pauzner. 2005. "Demand–Deposit Contracts and the Probability of Bank Runs." *The Journal of Finance*, 60(3): 1293–1327.

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

He, Zhiguo, and Wei Xiong. 2012. "Dynamic Debt Runs." *The Review of Financial Studies*, 25(6): 1799–1843.

Hörner, Johannes, and Andrzej Skrzypacz. 2016. "Learning, experimentation and information design." *Survey Prepared for the 2015 econometric summer meetings in Montreal.*

Huang, Chong. 2017. "Defending against Speculative Attacks: The Policy Maker's Reputation." *Journal of Economic Theory*, Forthcoming.

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

Judd, Kenneth L. 1985. "The Law of Large Numbers with a Continuum of iid Random Variables." *Journal of Economic Theory*, 35(1): 19–25.

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

Konrad, Kai A, and Tim BM Stolper. 2016. "Coordination and the Fight Against Tax Havens." *Journal of International Economics*, 103: 96–107.

Lagunoff, Roger, and Akihiko Matsui. 1997. "Asynchronous Choice in Repeated Coordination Games." *Econometrica*, 65(6): 1467–1477.

Milgrom, Paul, and John Roberts. 1990. "Rationalizability, Learning, and Equilibrium in Games with Strategic Complementarities." *Econometrica*, 58(6): 1255–1277.

Morris, Stephen, and Hyun Song Shin. 1998. "Unique Equilibrium in A Model of Self-fulfilling Currency Attacks." *The American Economic Review*, 88(3): 587–597.

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.

Sakovics, Jozsef, and Jakub Steiner. 2012. "Who Matters in Coordination Problems?" *The American Economic Review*, 102(7): 3439–3461.

Szkup, Michal, and Isabel Trevino. 2015. "Information Acquisition in Global Games of Regime Change." *Journal of Economic Theory*, 160: 387–428.

Varian, Hal R. 1994. "Sequential contributions to public goods." *Journal of Public Economics*, 53(2): 165–186.

Vives, Xavier. 2014. "Strategic Complementarity, Fragility, and Regulation." *Review of Financial Studies*, 27(12): 3547–3592.

Appendix

Proof of Claim 1

First, we show that $\exists x \in [0,1]$ such that $G(x, \alpha^*) \leq p$. Otherwise, it follows from uniform continuity that $\exists \delta > 0$ such that $\forall \alpha \in [\alpha^*, \alpha^* + \delta)$, $\forall x \in [0,1], G(x,\alpha) > p$. This implies $\alpha^* + \delta \in \mathcal{P}$, which contradicts the fact that $\alpha^* = \sup \mathcal{P}$.

Consider the non-trivial case in which $\alpha^* < 1$. For any $\alpha^{**} > \alpha^*$, by definition of α^*, $\alpha^{**} \notin \mathcal{P}$. Therefore, there exist $x^0 \in [0,1]$ and $\alpha^0 \in (0, \alpha^{**}]$ such that $G(x^0, \alpha^0) \leq p$. Let $A^0 = x^0 \cdot \alpha^0 \in [0, \alpha^0]$. Consider any $\alpha^1 \geq \alpha^0$ and take $x^1 = \frac{A^0}{\alpha^1}$. Then, $x^1 \leq x^0$. Let us define the cutoff strategies $\hat{s}^1$ and $\hat{s}^0$ corresponding to x^1 and x^0 as in the aggregate condition (A^α). Then,

$$\hat{s}^1 = \frac{1}{\sqrt{\tau}} F^{-1}(x^1) + A^0 \leq \frac{1}{\sqrt{\tau}} F^{-1}(x^0) + A^0 = \hat{s}^0.$$

Therefore, it follows from log-concavity of f that

$$G(\frac{A^0}{\alpha^1}, \alpha^1) = P(\theta \geq A^0 \mid \hat{s}^1, \theta \geq 0) \leq P(\theta \geq A^0 \mid \hat{s}^0, \theta \geq 0) = G(\frac{A^0}{\alpha^0}, \alpha^0).$$

Thus, for any $\alpha^1 \geq \alpha^0$, there exists $x^1 \in [0,1]$ such that $G(x^1, \alpha^1) \leq G(x^0, \alpha^0) \leq p$. In particular, for $\alpha^1 = \alpha^{**}$, there is some $x \in [0,1]$ such that $G(x, \alpha^{**}) \leq p$. $\square$

Proof of Proposition 1

Let us consider a general payoff structure for group j agents given the borrower has not failed yet, i.e., $\theta \geq \alpha \sum_{l=1}^{j-1} w_l$, as follows.

$$u(0, \theta, \{w_l\}_{l=1}^J) = \begin{cases} b_1(\theta, \{w_l\}_{l=1}^J) & \text{if } \theta \geq \alpha \sum_{l=1}^{J} w_l \\ c_0(\theta, \{w_l\}_{l=1}^J) & \text{if } \theta < \alpha \sum_{l=1}^{J} w_l. \end{cases}$$

$$u(1, \theta, \{w_l\}_{l=1}^J) = \begin{cases} c_1(\theta, \{w_l\}_{l=1}^J) & \text{if } \theta \geq \alpha \sum_{l=1}^{j} w_l \\ b_0(\theta, \{w_l\}_{l=1}^J) & \text{if } \theta < \alpha \sum_{l=1}^{j} w_l. \end{cases}$$

Let us define the net payoff from rolling over as opposed to withdrawal as $\overline{u}(\theta, \{w_l\}_{l=1}^J) := b_1(\theta, \{w_l\}_{l=1}^J) - c_1(\theta, \{w_l\}_{l=1}^J)$ when the borrower remains viable until the end, and as $\underline{u}(\theta, \{w_l\}_{l=1}^J) := c_0(\theta, \{w_l\}_{l=1}^J) - b_0(\theta, \{w_l\}_{l=1}^J)$ when the borrower fails at time $j\alpha$ or after it.

Assumption 1 *The payoff has the following properties:*

1) *(Complementarity)* $\overline{u}(.) \geq 0$ *and* $\underline{u}(.) \leq 0$*.*

2) *(Boundedness) There exist some finite numbers* $\underline{m}$, $\overline{m}$, $\underline{n}$ *and* $\overline{n}$ *such that* $0 < \underline{n} \leq \overline{u}(.) \leq \overline{n}$ *and* $0 < \underline{m} \leq -\underline{u}(.) \leq \overline{m}$,

The first part of the assumptions says that if the borrower is going to remain viable till time 1, then the agent is better off by rolling over, and if the borrower cannot withstand the withdrawal from group j, then the agent is better off by withdrawing. This captures the strategic complementarity. The second part of the assumption says that these net payoffs are bounded.

The payoff specification in the debt run application is a special case of this general payoff structure with $\overline{u} = r > 0$ and $-\underline{u} \in [0, 1]$. Below, we show that under the general payoff structure that satisfy Assumption 1, the argument N_j holds true for any $j = 1, 2, \ldots, J$.

In this application, statement M_j can be interpreted as the following- no creditor from later groups will withdraw if the borrower can service the withdrawal from group j, i.e., $w_l = 0$ for all $l > j$. That means $\theta \gtreqless \alpha \sum\limits_{l=1}^{J} w_l$ is equivalent to $\theta \gtreqless \alpha \sum\limits_{l=1}^{j} w_l$. Hence, we will follow the proof of Lemma 3 to prove the argument N_j.

Consider the marginal agent who receives the cutoff signal $\hat{s}_j$. Upon receiving the public news that the borrower is still viable at $(j-1)\alpha$, he believes $\theta \geq \underline{\theta}_{j-1}$ for some $\underline{\theta}_{j-1}$. He also understands that the borrower will not default if $\theta \geq A_j(\hat{s}_j, \alpha, \underline{\theta}_{j-1})$ (as defined in (A_j^α)). Therefore, the marginal agent will roll over if he believes that the probability of no default

$$\mathbb{P}(\theta \geq A_j | \hat{s}_j, \theta \geq \underline{\theta}_{j-1}) > \frac{1}{1 + \frac{E(\overline{u} | \theta \geq A_j, \hat{s}_j)}{E(-\underline{u} | \underline{\theta}_{j-1} \leq \theta < A_j, \hat{s}_j)}}.$$

It follows from Assumption 1 that, regardless of $\hat{s}_j$ and A_j, the RHS of the above inequality is lower than $\bar{p} \equiv \frac{1}{1 + \frac{\underline{n}}{\overline{m}}}$. Therefore, $\alpha < \hat{\alpha}(\overline{p}, \tau)$ is a sufficient condition to guarantee that the marginal agent strictly prefers to roll over. Since this is true for any $\hat{s}_j$ and $\underline{\theta}_{j-1}$, the inductive argument N_j holds true. Then, following the same inductive argument as in Theorem 1, the statement M_j is true for all $j = 0, 1, 2, \ldots, J$. $\square$

Proof of Proposition 2

Under uniform prior, the belief of the marginal agent in any group j (see Equation (5)) simplifies to $G^u(x,\alpha;\tau) := \frac{x}{F(\alpha\sqrt{\tau}x+F^{-1}(x))}$, where $x = \frac{A_j-\underline{\theta}_{j-1}}{\alpha}$. It follows from Lemma 1 that there exists a necessary and sufficient threshold $\alpha^*(p,\tau)$.

Hence, whenever $\alpha < \alpha^*(p,\tau)$, $G^u(x,\alpha;\tau) > p$ for all $x \in [0,1]$ (regardless of $\underline{\theta}_{j-1}$). That means argument N_j is true for all $j \geq 1$ when $\alpha < \alpha^*(p,\tau)$.

On the other hand, for $\alpha \geq \alpha^*(p,\tau)$, there exits $x^* \in (0,1]$ that solves $G^u(x^*,\alpha,\tau) = p$. Thus, it is always an equilibrium if the agents in the first group follow a cutoff strategy $\hat{s}_1 = \alpha x^* + \sigma F^{-1}(x^*)$ and no agent attacks after the regime passes the second viability test. The regime survives iff $\theta \geq \alpha x^* > 0$. Therefore, a diffused policy is persuasive if and only if $J > J^* \equiv 1/\alpha^*(p,\tau)$.

Note that $G^u(x,\alpha;\tau)$ is decreasing in α and τ. Let us first consider $p' \geq p$. If $\alpha < \alpha^*(p',\tau)$, then $G^u(x,\alpha;\tau) > p' \geq p$ for all $x \in [0,1]$. Hence, $\alpha^*(p',\tau) \leq \alpha^*(p,\tau)$. Next, consider $\tau' \geq \tau$. If $\alpha < \alpha^*(p,\tau')$, then $G^u(x,\alpha;\tau') > p$ for all $x \in [0,1]$. Since $G^u(x,\alpha;\tau)$ is decreasing in τ, and $\tau' \geq \tau$, we have $G^u(x,\alpha;\tau) > p$ for all $x \in [0,1]$ whenever $\alpha < \alpha^*(p,\tau')$. Hence, $\alpha^*(p,\tau') \leq \alpha^*(p,\tau)$. □

Proof of Proposition 3

We will prove the argument N_j for all $j = 1,2,\ldots,J$ and the inductive steps are the same as in the proof of Theorem 1. Consider any group j. Suppose some $\tilde{w} \leq (j-1)\alpha$ proportion of agents have attacked so far. Then, PNV implies that $\theta \geq \tilde{w}$. Under statement M_j, the regime will survive if it can sustain the attacks from the current group. Hence, if $\theta \geq \tilde{w}+\alpha$, it survives for sure since it is strong enough to withstand all possible attacks from group j. Therefore, no one in any group j attacks a regime that passes the jth viability test if for any $\tilde{w} \leq (j-1)\alpha$,

$$\mathbb{P}(\theta \geq \tilde{w}+\alpha | \theta \geq \tilde{w}, s) = \frac{F(\sqrt{\tau}(s-\tilde{w}-\alpha))}{F(\sqrt{\tau}(s-\tilde{w}))} > p.$$

Given log-concavity of f, this probability is decreasing in $\tilde{w}$. Therefore, the minimum possible value for the above probability is when $\tilde{w} = (j-1)\alpha$ and $j = J$. Hence, if

$$\frac{F(\sqrt{\tau}(s-1))}{F(\sqrt{\tau}(s-1+\alpha))} > p, \tag{6}$$

no agent in any group j will attack a viable regime regardless of their belief about the past attacks $\tilde{w}$. This gives us N_j, i.e., if M_j is true then so is M_{j-1}.

(I) Note that if $s > 1 - \frac{\sigma}{2}$, then $\mathbb{P}(\theta \geq 1|s) = F(\sqrt{\tau}(s-1)) > 0$. The LHS of Inequality (6) is continuous and decreasing in α, and converges to 1 as $\alpha \to 0$. Therefore, there exists $\alpha^*(s)$ such that if $\alpha \leq \alpha^*(s)$, the Inequality

(6) and thus the inductive argument N_j hold true. Moreover, since the LHS of Inequality (6) is continuously increasing in s (given log-concave f) and continuously decreasing in α, $\alpha^*(s)$ is continuously increasing in s. In other words, $J^*(s)$ increases continuously when s decreases.

(II) If $s \leq 1 - \frac{\sigma}{2}$, $\mathbb{P}(\theta \geq 1|s) = F(\sqrt{\tau}(s-1)) = 0$. So, the above Inequality (6) is violated regardless of α. Therefore, if the agents in the last group believe that all their predecessors have attacked, they will attack as well. Thus, agents in group $J-1$ can also attack if they think the agents in group J will and so on. Thus, all attack is a possible equilibrium outcome.

(III) For any $\theta < 1$, there is a positive probability that a public signal $s \leq 1 - \frac{\sigma}{2}$ will be realized. Upon receiving such signal, it is possible that all the agents attack, regardless of the policy J. □