Diffusing Coordination Risk

Deepal Basak and Zhen Zhou*

October 17, 2018

Abstract

In a regime change game, agents sequentially decide whether to attack or not, without observing the past actions by others. To dissuade them from attacking, a principal adopts a dynamic information disclosure policy - repeated viability tests. A viability test publicly discloses whether the regime has survived the attacks so far. When such tests are sufficiently frequent, in the unique cutoff equilibrium, regardless of their private signals, agents never attack if the regime passes the latest test. We apply our theory to show that by sufficiently diffusing the rollover choices across different maturity dates, a borrower can eliminate panic-based runs.

JEL Classification Numbers: *C72, D82, D83, G28, G33*

Key Words: *Coordination, Global Game, Information Design, Self-fulfilling Runs*

*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

In a coordination game, the strategic uncertainty that agents face concerning the actions and beliefs of others may lead to an undesirable outcome. Think of 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 could cause a debt run solely based on panic and not on fundamental. We describe this as a global game of regime change.[1] In a dynamic setting, we propose a simple information disclosure policy to influence the beliefs and thus actions of the agents. This policy resolves the strategic uncertainty and avoids the undesirable outcome.

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 they do not see the past actions by other agents. The underlying fundamental strength of the regime is θ. If the underlying fundamental is strong enough to withstand the aggregate attack, the regime succeeds, otherwise it fails. Suppose there a threshold p such that an agent will not attack if, and only if he believes that with probability higher than p the regime will succeed. As is standard in the global game literature, the agents are uncertain about the underlying fundamental θ and receive some noisy private signals s_i about it.

The principal prefers the regime to succeed. She strategically discloses information at different dates regarding the underlying fundamental and the past attacks to manage the agents' beliefs in order to dissuade them from attacking. First, let us consider two simple policies - no disclosure and full disclosure. Under no disclosure, the agents do not have any information about the past actions by other agents. Therefore, the game is essentially a simultaneous move game (see Morris and Shin (2003)) and there is a unique equilibrium in which the agents attack if they receive private signal below some cutoff. On the other hand, if the agents know the underlying fundamental and perfectly observe the past actions, all attacking is a possible equilibrium outcome.[2] Thus, even if full disclosure is feasible, the principal may not want to disclose all the information. Could a partial disclosure policy be

[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).

[2] If there are finitely many agents instead of a continuum of agents, an agent can influence the aggregate attack. When finitely many agents move sequentially, and there is perfect and complete information, it follows from backward induction that the agents will play the payoff dominant action. Hence, the full disclosure policy cannot dissuade the agents from attacking if attacking is the payoff dominant action (see Section 5 for details).

more effective in dissuading the agents from attacking?

A vast range of partial disclosure policies exist. We consider a simple dynamic information disclosure policy - *repeated viability tests*. As the name suggests, a viability test at some date $t \in [0, 1]$ checks if the regime continues to be viable, i.e., if it has survived the attacks so far (if any), and publicly discloses the test result. The principal only chooses an integer J denoting the frequency of viability tests. The tests are conducted at regular interval $1/J$ starting at 0.

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

Viability tests, at any date, have binary outcome. In the spirit of Bayesian Persuasion, we can interpret this policy as a recommendation by the principal to the agents to not attack the regime that passes the viability test, and attack if it fails. If a regime fails a viability test, it cannot succeed even without any further attack. Hence, it is the dominant strategy for all the agents in the subsequent groups to follow the principal's recommendation and attack. Therefore, the question is - when the regime passes the test, will the agents follow the principal's recommendation and not attack? The above result says that when the principal repeats the viability tests frequently enough, all agents following the principal's recommendation is the only cutoff equilibrium.

The positive viability news makes the agents 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 are less likely to attack. Thus, it is intuitive that a positive viability news will reduce the aggregate attack in equilibrium. However, since agents are privately informed about the underlying fundamental, it is possible that some agents receive very low private signals and believe that the regime is unlikely to succeed even if a small mass of agents attack. So, how can this partial disclosure policy be so influential that all the agents, regardless of their private signals, are convinced that the regime will succeed with probability higher than p, and thus completely ignore their private information and follow the principal's recommendation?

When the principal repeats the viability tests for J times, the policy endogenously 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 look into equilibrium in cutoff strategies - for any j, after learning that the regime has passed the jth viability test, any agent in 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 higher than the cutoff and thus fewer agents attack from each group. This induces a non-decreasing sequence of fundamental cutoffs ($\underline{\theta}_{j-1}$) such that the regime passes the jth test if and only if the underlying fundamental is no lower than such cutoffs. Hence, the viability tests publicly disclose whether $\theta \geq \underline{\theta}_{j-1}$ or not. The policy is history dependent because of such endogenous cutoffs. The regime succeeds in the end when the fundamental is no lower than $\underline{\theta}_J$. The main result claims that when the viability tests are sufficiently frequent, $\underline{\theta}_J = 0$, i.e., any viable regime succeeds in the end.

When the viability tests are repeated, the effectiveness of viability news depends on how effective it will be in the future. We build an inductive argument - if the tests are conducted with sufficient frequency, then regardless of the cutoff strategy played by others in the past (equivalently for any $\underline{\theta}_{j-1}$), the agents in group j follow the principal's recommendation if they believe that the agents in the subsequent groups will do so.

Given the regime has passed the jth viability test and the agents in subsequent groups will follow the principal's recommendation, the regime will succeed as long as it withstands attacks from group j. Under frequent viability tests, the group size is small. Thus, from an ex-ante perspective, it is very likely that the regime will succeed even if all agents in group j attack. However, this does not mean that the regime will always succeed in this case. Think of the agent who receives a very low private signal and thus knows that the regime cannot succeed if other agents in group j attack. Therefore, whether this agent will attack or not depends on his belief about others, i.e., the equilibrium strategy.

Suppose that group j agents play some cutoff strategy $\hat{s}_j$. If the group size is smaller, then there will be less attack from group j and thus it gets easier to succeed. On the other hand, the "marginal agent" with the signal $\hat{s}_j$ becomes more optimistic after learning the positive viability news. The combined effect of the positive viability news and small group size resulting from frequent tests is that there exists J^* such that when $J > J^*$, regardless of $\underline{\theta}_{j-1}$, if agents in group j follow some cutoff strategy $\hat{s}_j$, the marginal agent believes that the regime will succeed with probability strictly higher than p. Hence, he strictly prefers to not attack. This shows that although an agent in group j may not think the regime will succeed if others in group j attack, there is no cutoff equilibrium in which the group j agents attack the regime that passes the jth viability test.

Under sufficiently frequent viability tests, since no one moves after the last group of agents, group J agents will follow the principal's recommendation, and 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.[3] This type of debt structure diffuses the rollover risk that would otherwise be concentrated at one point. 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 yet) 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 the agents' beliefs. One may think that the news is not likely to have any substantial influence on the creditors' behavior. However, the borrower can repeatedly provide this news for J many times when only $1/J$ fraction of debts matures 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 tries to achieve the desired outcome that the market fails to achieve. Sakovics and Steiner (2012) and Cong, Grenadier and Hu (2018) find the optimal subsidies that, at a given cost, maximize the likelihood of successful coordination.[4] Unlike the planner in the above mentioned papers, we consider a principal who cannot offer monetary incentives. However, she can disclose some information to persuade the agents to follow her recommendations as in Bergemann and Morris (2013).

We consider a canonical global game of regime change and as is standard in the literature (see Carlsson and Van Damme (1993)), we assume a private information environment. Our principal repeatedly runs a simple test that publicly discloses partial information - whether the regime continues to be viable or not. A positive viability news removes the lower dominance region. Under such one-sided dominance, Angeletos, Hellwig and Pavan (2007) show that there can be multiple equilibria (also see Bueno de Mesquita (2014) and Huang (2017)). Acharya and Ramsay (2013) provide a sufficient condition - a restriction on the information structure - that uniquely selects the no attack equilibrium under one-sided dominance. In the online appendix, we show that this sufficient condition may be violated under our policy.

Broadly speaking, this paper belongs to the Bayesian Persuasion and information de-

[3] For empirical evidence see He and Xiong (2012) and Choi, Hackbarth and Zechner (2017).

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

sign literature for heterogeneously informed multiple receivers as defined by Bergemann and Morris (2016). We focus on public information disclosure. The two most closely related papers in this literature are Inostroza and Pavan (2017) and Goldstein and Huang (2016). The authors also propose a partial information disclosure policy - a one time "stress test." If the agents are not moving simultaneously, then repeated viability tests outperform such one-time stress test. In Section 3, we discuss this comparison in detail. Optimal information disclosure has also been studied in other strategic contexts such as voting (see Alonso and Câmara (2016*a*), Alonso and Câmara (2016*b*) and Bardhi and Guo (2016)) and auctions (see Eső and Szentes (2007)). Mathevet, Perego and Taneva (2017) and Taneva (2016) build a general persuasion setting with multiple receivers. There is a large literature on dynamic information feedback in the context of strategic experimentation. See Hörner and Skrzypacz (2016) for a survey of this literature. Ely (2017) extends the static Bayesian Persuasion model of Kamenica and Gentzkow (2011) to dynamic setting but restricts to history independent policies.

Our paper is also related to dynamic coordination game literature. Dasgupta (2007) considers a two-period problem where the agents receive noisy private information about the attacks from the first period. However, the agents endogenously decide when to attack. In our model, the timing of action is exogenous.[5] In Frankel and Pauzner (2000), the agents get chances to revise their decisions following a poisson process. The agents know the current state but are uncertain about the volatile future. This also introduces asynchronous moves but in a different way.[6] He and Xiong (2012) extend this framework to study the role of volatile fundamental under asynchronous debt structure. While these above mentioned papers take the asynchronous structure as given, we take a step back and address the question - why should there be asynchronous debt structure in the first place? We find the justification in terms of information transmission and find the optimal design of asynchronous structure.

The rest of the paper is organized as follows. Section 1 describes the model, the diffused policy and the solution concept. Section 2 shows the role of a one time viability test. In Section 3, we leverage this idea to build the argument for repeated viability tests and establishes our main result. In Section 4, we extend the model to study the debt rollover

[5]We discuss how the disclosure policy works with endogenous move in the online appendix.

[6]Sequential moves 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.

application. In Section 5, we discuss the full disclosure policy and what could make persuasion harder, or even impossible. Section 6 concludes. The proofs that are not in the paper can be found in the appendix.

1 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 0 means attacking the regime and 1 means not attacking the regime. The agents do not observe past actions by other agents. Let θ be the underlying fundamental strength of the regime and $w = \int_i \mathbb{1}(a_i = 0)di$ be the aggregate attack. The regime succeeds in the end 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 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

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

We assume that (1) $b_1 > c_1$, i.e., if an agent knows that the regime will succeed, he is better off by not attacking it, and (2) $b_0 > c_0$, i.e., if an agent knows that the regime will fail, he is better off by attacking it.[7]

However, agents are uncertain about θ and hence do not know if the regime will succeed or not. Nature selects the fundamental strength θ from the interval $[\underline{\theta}, \bar{\theta}]$ with uniform probability. Before the game begins, any agent i receives noisy private information about θ, denoted by $s_i = \theta + \sigma\epsilon_i$, where the error terms ϵ_i are conditionally independent[8] 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 with $f(-\frac{1}{2}) < \infty$. σ scales the random noise ϵ_i and $\tau \equiv 1/\sigma^2$ is defined as the precision of the private signal. We assume that $\underline{\theta} \leq -\sigma$ and $\bar{\theta} \geq 1 + \sigma$.[9]

[7] This payoff specification is standard in the literature, but since the agents are moving sequentially, in practice, the payoff may depend on the history of attacks. In Section 4, 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] This implies that for any $\theta \in [0,1]$, the probability of receiving private signal s is distributed via $F((s-\theta)/\sigma)$. Given the uniform prior, for any $s \in [-\sigma/2, 1+\sigma/2]$, θ is distributed via $1 - F((s-\theta)/\sigma)$.

The distribution F is continuously differentiable and log-concave, which implies that for any $a > 0$, $\frac{\partial}{\partial x}\left(F(x-a)/F(x)\right) > 0$.[10] For numerical examples or illustrations, we use a triangle probability density of error - $\hat{f}(x) := (2+4x)\mathbf{1}(-0.5 \le x < 0) + (2-4x)\mathbf{1}(0 \le x \le 0.5)$.

After receiving his private signal an agent updates his belief about θ. Suppose he believes that the regime will succeed with probability P. Then his expected payoff from attacking ($a_i = 0$) is $Pc_1 + (1-P)b_0$, while his expected payoff from not attacking ($a_i = 1$) is $Pb_1 + (1-P)c_0$. Therefore, he does not attack if and only if he believes that the regime will succeed with probability P higher than

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

Thus, p measures the predisposition to attacking or reluctance to not attacking the regime.

The Principal's Objective The principal gets a payoff of 1 if the regime succeeds and 0 otherwise. If $\theta < 0$, the regime cannot succeed. But if $\theta \ge 0$, the regime is viable, which means the regime succeeds if no one attacks the regime. However, a viable regime may fail if agents do not coordinate on not attacking. We call the ex-ante probability that a viable regime may fail - the *coordination risk*. The principal wants to minimize this coordination risk. If the agents never attack a viable regime, then this risk is eliminated. 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. We discuss this in Section 5.

A diffused policy J The agents are heterogeneously informed about the fundamental because of the private signals they receive. The principal, who can be thought of as an information designer, is able to disclose some information over time based on the fundamental as well as the history of actions so far. 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 attack, and then disclose the test result publicly. We call such tests - viability tests, and a positive test result - the public news of continued viability - **PNV**. At time 0 before agents move, the principal chooses how frequently to run such viability tests. The policy is announced publicly. Without loss

[10] Inostroza and Pavan (2017) also make a similar assumption regarding the noise distribution. For tractability, the global game literature often uses Gaussian error distribution, which satisfies log-concavity.

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 after the jth test and before the next one. 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. Thus, the group identity of an agent is endogenously determined by the policy. A more diffused policy means repeating the viability tests more frequently (higher J), or equivalently the group size α is smaller.

Cutoff Equilibrium Under a diffused policy J, when the group j agents decide whether to attack or not, they get the public information whether the regime has passed the jth viability test, or not. In addition, each agent has a private signal regarding the underlying fundamental. If the regime ever fails a viability test, then the regime will fail regardless of what the rest of the agents do. Hence, attacking is the the dominant strategy for the rest of the agents. The strategic decision is non-trivial only when the regime passes the test, which would be the focus of our analysis. In what follows, we limit our attention to Perfect Bayesian Equilibrium in monotone strategies, i.e., after learning 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 agents follow monotone strategies, the aggregate attack is non-increasing in the fundamental strength θ. Therefore, agent i believes the regime is strictly more likely to succeed when s_i is higher. That means in equilibrium all agents in the same group j (for any j) must follow a symmetric cutoff strategy - attack if and only if $s_i < \hat{s}_j$.[11] We refer to such equilibrium as cutoff equilibrium.

When the agents follow cutoff strategies, a larger mass of agents attack against a weaker regime, and as the game continues (weakly) more agents will attack. Hence, there exist a non-decreasing sequence of cutoff $\underline{\theta}_{j-1}$ such that the regime passes the jth viability tests if and only if $\theta \geq \underline{\theta}_{j-1}$. The regime passes all the viability tests, and succeeds in the end if and only if $\theta \geq \underline{\theta}_J =: \hat{\theta}$. There can be multiple cutoff equilibria.

[11] Suppose an agent in group j randomizes when he receives a signal $\hat{s}_j$. Then, he 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 plays no significant role for our results).

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. If the regime passes a viability test, the agents learn that it will succeed in the end without any further attack. Hence, if an agent believes that others will follow the principal's recommendation, then he should do the same. Therefore, all agents follow the principal's recommendation is a cutoff equilibrium. However, there may exist some cutoff equilibrium in which some agents do not follow the principal's recommendation and attack a regime even when it passes the latest test.

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 the one in which the 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 $\theta \geq 0$ succeeds in the end. In other words, under a persuasive policy, the only equilibrium cutoff fundamental is $\hat{\theta} = 0$ and the coordination risk is eliminated.

2 One-time Viability Test

Consider a degenerate policy - no test. Since the agents do not have any information about the past actions by other agents, the game is essentially a simultaneous move regime change game.

Proposition 1 *Under the degenerate policy, there is a unique equilibrium in which the agents follow a cutoff strategy : do not attack if and only if $s_i \geq s^* = p + \frac{1}{\sqrt{\tau}} F^{-1}(p)$. Consequently, the regime succeeds if and only if $\theta \geq p$.*

The result follows directly from iterated elimination of never best responses as in Morris and Shin (2003). This shows that if the agents are more reluctant to not attacking, it is harder to succeed.

Now, to understand the effect of 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 the agents in the same fashion. Agents who receive private signal $s_i \in [-\frac{\sigma}{2}, \frac{\sigma}{2}]$ used to believe that θ can be less than 0, but this public news tells them otherwise. This makes them more optimistic about the strength of the regime, and thus less likely to attack. Because of strategic complementarity, other agents, including the ones who already know that $\theta \geq 0$ from their private information, become more optimistic about the success of the regime, and hence they are also less likely to attack.[12] It is intuitive that after learning that the regime is viable, in equilibrium, the agents will attack less aggressively. But when will the news become so strong that there cannot be an equilibrium in which an agent attacks the regime that has passed the viability test?

Consider a hypothetical game which is exactly as above with one exception - it is played between some $\alpha < 1$ mass of agents in stead of 1 mass of agents. The smaller the α, the more likely it is that $\theta \geq \alpha$, which means that even if all the α mass of agents attack the regime will succeed. However, note that an agent can receive a private signal $s_i < \alpha - \sigma/2$ and when she does, she believes that the regime will not succeed 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 look into how the equilibrium strategy is affected when a small mass of agents learn 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 playing the game is sufficiently small. We emphasize this result here 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 $\alpha^*(p, \tau) > 0$ such that for $\alpha < \alpha^*$, there is no cutoff equilbrium in which an agent attacks a viable regime.*

Proof. Suppose that the agents are following 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 that

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

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

[12] Note that PNV makes the agents more optimistic about the success of the regime regardless of whether attacking is the payoff dominated action or not. We will revisit this issue later in Section 5.

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

$$\mathbb{P}(\theta \geq A(\hat{s}, \alpha)|s_i, \theta \geq 0) = \frac{\mathbb{P}(\theta \geq A(\hat{s}, \alpha)|s_i)}{\mathbb{P}(\theta \geq 0|s_i, \theta \geq 0)} = \frac{F(\sqrt{\tau}(s_i - A(\hat{s}, \alpha)))}{F(\sqrt{\tau}s_i)}. \tag{3}$$

It follows from log-concavity of F that this probability is higher for agents with higher private information s_i. Define $I^v(\hat{\theta})$ such that if the success 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.

$$\frac{F(\sqrt{\tau}(I^v(\hat{\theta}) - \hat{\theta}))}{F(\sqrt{\tau}I^v(\hat{\theta}))} = p. \tag{I^v}$$

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

Suppose that all the agents are following the cutoff strategy $\hat{s}$ and consequently, the success 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

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

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

$$\frac{\left(\frac{A(\hat{s},\alpha)}{\alpha}\right)}{F\left(F^{-1}\left(\frac{A(\hat{s},\alpha)}{\alpha}\right) + \alpha\sqrt{\tau}\left(\frac{A(\hat{s},\alpha)}{\alpha}\right)\right)} = p.$$

The left hand side captures the belief of the marginal agent that - "given that all α fraction of agents play the cutoff strategy $\hat{s}$, the regime has a per capital fundamental strength that surpasses the required cutoff $\frac{A(\hat{s},\alpha)}{\alpha}$, and thus succeeds." We define

$$G\left(x, \tau, \alpha\right) := \frac{x}{F\left(F^{-1}(x) + \alpha\sqrt{\tau}x\right)} \tag{4}$$

and $\alpha^*(p, \tau)$ is such that

$$\min_{x \in [0,1]} G(x, \tau, \alpha^*) = p, \tag{5}$$

and if $\min\limits_{x \in [0,1]} G(x, \tau, \alpha) > p$ for all $\alpha \leq 1$, then $\alpha^* = 1$.

Claim 1 $\alpha^*(p,\tau) > 0$. *For any* $\alpha < \alpha^*$, $G(x,\tau,\alpha) > p$ *for all* $x \in [0,1]$.

The result follows from $G(x,\tau,\alpha)$ being decreasing in α for any x. The formal proof of this claim is relegated to the appendix.

The above claim implies that when $\alpha < \alpha^*$, given any cutoff strategy $\hat{s}$, the marginal agent with signal $\hat{s}$ believes that the regime will succeed with probability strictly higher than p. Thus, he strictly prefers not to attack, i.e., $\beta(\hat{s}) < \hat{s}$. This implies that for any $\hat{s}$, the cutoff strategy $\hat{s}$ cannot constitute an equilibrium. □ ■

Recall that there are two forces at work here: (1) the public news of viability and (2) a small mass of agents. Together they are responsible for Claim 1. Let us first consider them separately.

The news effect As we have already pointed out, PNV makes the agents more optimistic about the success of the regime. For any success criteria $\hat{\theta}$, the private information $I^v(\hat{\theta})$ that makes agents indifferent between attacking and not attacking, when there is PNV, is lower than that in the absence of any viability test. To see this, consider the signal $I(\hat{\theta})$ that makes an agent indifferent between attacking and not attacking. $I(\hat{\theta})$ must satisfy

$$F(\sqrt{\tau}(I(\hat{\theta}) - \hat{\theta})) = p. \tag{I}$$

Comparing with condition (I^v), it is easy to see that

$$I^v(\hat{\theta}) \leq I(\hat{\theta}).$$

The log-concavity of F guarantees that $I^v(\hat{\theta})$ is a strictly increasing function of success criterion $\hat{\theta}$. Therefore, if success criterion $\hat{\theta}$ is low, then the denominator $F(\sqrt{\tau}I^v(\hat{\theta}))$ in Condition (I^v) is low and hence $I^v(\hat{\theta})$ will be further away from $I(\hat{\theta})$. This implies that the impact of PNV is more significant for a lower success criteria $\hat{\theta}$.

The group size effect It directly follows from 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 translate into a lower $A(\hat{s},\alpha)$ and in that case, the PNV is able to significantly lower $\beta(\hat{s}) = I^v(A(\hat{s}, \alpha))$, which is the cutoff signal that makes an agent indifferent between attacking and not attacking. If the group size α is sufficiently small, this impact is so significant that for any $\hat{s}$, $\beta(\hat{s}) < \hat{s}$. This gives us Claim 1.

Figure 1 explains this combined effect graphically. Consider any cutoff per capita fundamental x. As defined in Equation 4, $G(x, \tau, \alpha)$ is the belief of the marginal agent that the regime will succeed. Figure 1 plots this belief against any candidate cutoff $x \in [0, 1]$. For a degenerate policy, since there is no PNV, the denominator in $G(x, \tau, \alpha)$ is 1. As in Morris and Shin (2003), $G(x, \tau, \alpha)$ is the $45°$ line and hence at the cutoff per capita fundamental $x = p$, the agent with the cutoff signal is indifferent (as in Proposition 1). Thus, in the absence of PNV, a small group size does not affect the equilibrium per capita fundamental cutoff.

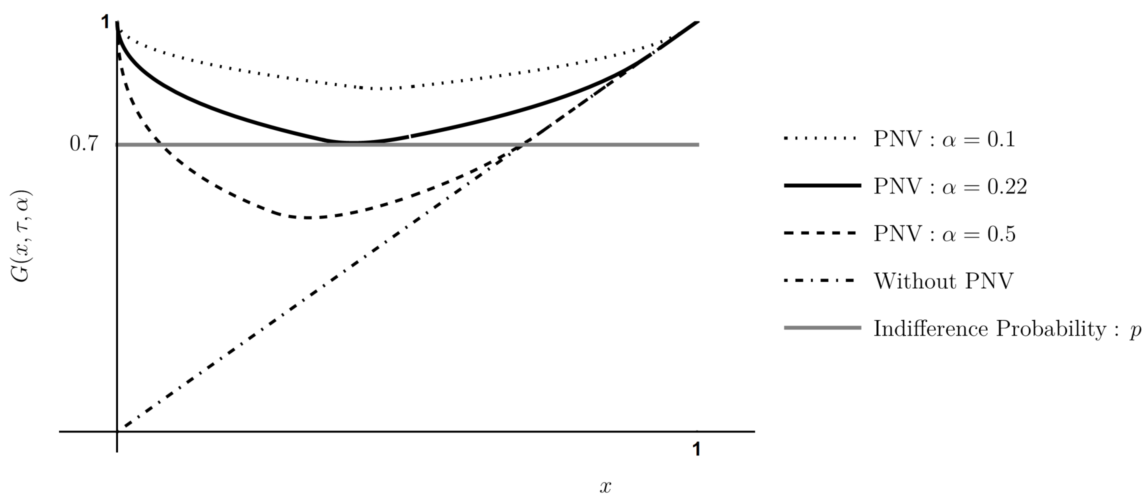

Figure 1: $G(x, \tau, \alpha)$ when $\tau = 1$, $\{\underline{\theta}, \bar{\theta}\} = \{-1, 2\}$ and $f = \hat{f}$.

We can see that PNV pushes this belief upward. As α decreases, $G(x, \tau, \alpha)$ increases. There is an α^* such that for $\alpha < \alpha^*$, $G(x, \tau, \alpha) > p$ for all x. This implies that the only possible cutoff solution is $x = 0$ and the implied unique equilibrium is that no one attacks a viable regime. Also, note that if $\alpha \geq \alpha^*$, then there exits a cutoff equilibrium in which the agents with sufficiently low private signal will attack a viable regime. Hence, $\alpha < \alpha^*$ is a necessary and sufficient condition for the policy to be persuasive.[13]

[13] Acharya and Ramsay (2013) provide a sufficient condition (A1(ii) in their paper) - a restriction on the information structure - for unique equilibrium selection in such an environment. It is worth mentioning that this sufficient condition may be violated under a sufficiently diffused policy (see the online appendix).

3 Main Result: Repeated Viability Tests

We understand that a positive viability news can dissuade a sufficiently small mass of agents from attacking. Now let us go back to the repeated viability tests - 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 $J = 2$.

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 2nd viability test ($\hat{s}_2 = \underline{\theta}_1 - \frac{\sigma}{2}$). Thus, if the regime passes the second viability test, it succeeds in the end ($\underline{\theta}_1 = \hat{\theta}$). It is also possible (Equilibrium 4 and 5) that agents in both group 1 and 2 may 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

Talenotes 1 *Parameter Values: $\tau = 1$, $p = 0.7$, $\{\underline{\theta}, \bar{\theta}\} = \{-1, 2\}$ and $f = \hat{f}$*

It is easy to see that for any diffused policy J, one can always construct a cutoff equilibrium in which after the regime passes the jth viability test, for some $j = 1, 2, \ldots, J$, no agent attacks the regime. This means the regime that passes the jth viability test, continues to be viable thereafter 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, especially if he has received a very low private signal, and hence, the multiplicity of cutoff equilibria arise. Nevertheless, we argue that when the principal adopts a sufficiently diffused policy, the only possible cutoff equilibrium is the one in which all the agents follow the principal's recommendation and not attack a viable regime.

Theorem 1 *A diffused policy J with $J > J^* = \frac{1}{\alpha^*(p,\tau)}$ is persuasive.*

As we can see from Figure 1, which shares the same parametric specification with above numerical example, $\alpha^* = 0.22$. Thus, based on above theorem, if $J = 5$ rather than 2, the only cutoff equilibrium is such that $\hat{\theta} = 0$.

Consider the case where group j agents learn that the regime has passed the jth viability test and decide whether to attack the regime or not. When $J > J^*$, there are less than α^* mass of agents in group j. Recall from Lemma 1 that a viability test can be persuasive if the mass of agents is less than α^*. However, there are two crucial differences here. First, whether the regime can pass the jth viability test depends on θ as well as the past attacks. Thus, the interpretation of the positive viability news is history dependent. To see 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 now, there are $(1 - j\alpha)$ mass of agents who will move later and the regime may fail if the agents in the subsequent groups attack. Thus, effectiveness of jth viability test depends on how effective the future viability tests are going to be.

However, we build on Lemma 1 and argue that the following statement is true for any j, in particular 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 whatever cutoff strategy has been played in the past), any cutoff strategy in which an agent in group j' may attack, violates sequential rationality.

Since we are only considering cutoff strategies, we sometime 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 $J > J^*$, 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 $\alpha < \alpha^*$, then N_1 is true.*

It directly follows from Lemma 1 that when $\alpha < \alpha^*$, if the agents in group 1 follow some cutoff strategy $\hat{s}_1$, then the agent who receives the private signal $s_i = \hat{s}_1$, believes that the regime will withstand the current attacks with probability strictly higher than p. If the regime withstands the current attack, then it will pass the next test. Given M_1 is true, this means that the agent believes that the regime will succeed with probability strictly higher than p. Hence, he is strictly better off by 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 the later groups is that the result of the viability test for group $j > 1$ is history dependent. Interestingly, the following lemma shows that given M_j is true, independent of history, i.e., the cutoff strategies played by the agents before group j, the positive viability news is at least as effective for group $j > 1$ as it is for group $j = 1$.

Lemma 3 *If $\alpha < \alpha^*$, 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 a cutoff strategy $\hat{s}_l$ for $l = 1, 2, \ldots, (j-1)$. Then, the regime will pass the jth viability test if

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

Suppose that $\theta = \underline{\theta}_{j-1}$ solves the above with equality. Thus, for any cutoff strategies the agents may have played in the past, PNV means $\theta \geq \underline{\theta}_{j-1}$ for some $\underline{\theta}_{j-1}$.[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

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

If $\theta \geq A_j$, the past attack is at most $\underline{\theta}_{j-1}$ (since $A_j \geq \underline{\theta}_{j-1}$), and the current attack is at most $\alpha F(\sqrt{\tau}(\hat{s}_j - A_j))$. Therefore, when $\theta \geq A_j$, the regime is strong enough to withstand the attacks up until group j.[15] Given M_j is true, this means whenever $\theta \geq A_j$, the regime

[14] If agents were not following cutoff strategies, then the implication of PNV is not necessarily that the underlying fundamental is greater than such cutoffs. Although the result in Lemma 1 can be generalized to rationalizability (as in standard in the static global game literature), we cannot do the same for Lemma 3, unless we restrict the strategy space to monotone strategies only.

[15] If $\underline{\theta}_{j-1} > 0$, then even for $\theta = A_j > \underline{\theta}_{j-1}$, the past attack is strictly lower than $\underline{\theta}_{j-1}$. Therefore, when $\underline{\theta}_{j-1} > 0$, $\theta = A_j$ is sufficient but not necessary to withstand the attacks up until group j. However, when

succeeds.

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

$$\frac{\mathbb{P}(\theta \geq A_j|\hat{s}_j)}{\mathbb{P}(\theta \geq \underline{\theta}_{j-1}|\hat{s}_j)} = \frac{F\left(\sqrt{\tau}(\hat{s}_j - A_j)\right)}{F\left(\sqrt{\tau}(\hat{s}_j - \underline{\theta}_{j-1})\right)}.$$

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

$$\frac{\left(\frac{A_j - \underline{\theta}_{j-1}}{\alpha}\right)}{F\left(F^{-1}\left(\frac{A_j - \underline{\theta}_{j-1}}{\alpha}\right) + \alpha\sqrt{\tau}\left(\frac{A_j - \underline{\theta}_{j-1}}{\alpha}\right)\right)} = G\left(\left(\frac{A_j - \underline{\theta}_{j-1}}{\alpha}\right), \tau, \alpha\right).$$

By definition, A_j is increasing in $\hat{s}_j$ and $A_j \in [\underline{\theta}_{j-1}, \underline{\theta}_{j-1} + \alpha]$. We want to show that for all possible cutoff strategy $\hat{s}_j$, the above probability is strictly greater than p, regardless of $\underline{\theta}_{j-1}$. Note that $\hat{s}_j$ enters the $G(.)$ function through the first argument.

Consider $\alpha < \alpha^*(p, \tau)$. It follows from Claim 1 that $G(x, \tau, \alpha) > p$ for all $x \in [0, 1]$. Therefore, 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 $\theta \geq A_j(\hat{s}_j, \alpha, \underline{\theta}_{j-1})$ and thus the regime succeeds with probability strictly higher than p. Therefore, M_{j-1} is true.

■

Step C

M_j is true for the last group $j = J$ since there is no agent who can attack after group J. By induction, we can show that in any cutoff equilibrium, no agent will ever attack a viable regime. Thus, if the principal adopts a sufficiently diffused policy ($J > J^*$), the coordination risk unravels from the end. This gives us our main result theorem 1.

When the principal repeats the viability tests frequently enough, the agents are assured that no agent will attack when the regime passes the next viability test (M_j is true), and this makes the current positive viability news persuasive (N_j is true). Note that if $J \leq J^*$, then one can always construct a cutoff equilibrium in which some agents may attack a viable regime. For $\alpha \geq \alpha^*$, there exits $x^* \in (0, 1]$ that solves $G(x^*, \alpha, \tau) = p$. Thus, it is always an equilibrium if the agents in the first group follow a cutoff strategy $\hat{s}_1 = \alpha x^* + \sigma F^{-1}(x^*)$

$\underline{\theta}_{j-1} = 0$ (as in Lemma 2), then $A_j(\hat{s}_j, \alpha, \underline{\theta}_{j-1} = 0) = A(\hat{s}_j, \alpha)$, and for $\theta = A_j > \underline{\theta}_{j-1}$, the past attack will also be 0. Hence, $\theta = A_j$ is also the necessary criterion.

and no agent attacks after the regime passes the second viability test. The regime succeeds iff $\theta \geq \alpha x^*$. Therefore, a diffused policy is persuasive if and only if $J > J^*$.

Repeated 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 show that if the stress test is sufficiently tough $k > \hat{k}$, then even in the cutoff equilibrium where agents attack most aggressively, no agent will attack the regime that passes the stress test. In contrast, our paper considers repeated tests with the weakest possible strength, or repeated viability tests. Although a one time viability test may not be "tough" enough, when repeated with a sufficient frequency, it can dissuade the agents from attacking a viable regime.[16]

Any viable regime succeeds when viability tests are repeated sufficiently frequently. In contrast, under one-time stress test with $k > \hat{k}$, a viable regime with strength below k fails. Thus, when the agents move sequentially, from an ex-ante perspective, the principal can do better by adopting a sufficiently frequent viability test policy as compared to a sufficiently tough one-time test policy. 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 succeed. This is not so for viability tests, since if a regime is not viable, the principal cannot do any better by misreporting it.[17]

4 Application : Panic-Based Runs

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 sound fundamentals may fail due to the coordination failure among the creditors in rolling over their debt. Any financial institution that performs liquidity transformation, e.g., commercial

[16] In fact, in the online appendix, we show that as the principal repeats such stress tests more frequently, the tests can be less tough to be persuasive.

[17] Thus, implementing a stress test policy requires commitment power as is standard in the Bayesian persuasion literature (see Kamenica and Gentzkow (2011)), but implementing viability tests do not need 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.

banks, hedge funds and mutual funds, could be exposed to such panic-based runs. Brunnermeier (2009) argues that this coordination risk on account of the maturity mismatch is one of the main causes of the recent financial crisis.

Asynchronous debt structures that prevent creditors from withdrawing simultaneously are fairly common in practice (see He and Xiong (2012)). Choi, Hackbarth and Zechner (2017) document empirical evidence that corporate bond issuers diversify debt rollovers across dates and 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.[18] Can implementing such asynchronous structure help to solve the panic-based runs? We are looking for a theoretical rationale from the persepective of information transmission.

Let us consider a stylized model in which there is a unit mass of creditors and each of them has lent 1 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 has to go through the bankruptcy process, which is public information to all the 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 1.[19]

The asynchronous debt structure separates creditors into different groups based on the dispersed maturity dates. Group j creditors refers to the creditors who decide whether to withdraw ($a_i = 0$), or roll over the debt ($a_i = 1$) at time $(j-1)\alpha$. The return from the illiquid investment R is realized at time 1. Before that time, the borrower has some extra funding source, for example, other liquid assets or credit lines, that can be used to sustain total withdrawals up to θ.

The information environment is assumed to be the same as the one in Section 1. Agents have private information about θ. The borrower chooses the debt maturity structure J prior

[18] 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 when a fund's liquidity position is not favorable.

[19] We do not model the ex-ante lending, but only the roll over problem. However, ceteris paribus, if the borrower is less likely to fail, then the creditors will be more willing to lend in the first place. Thus, such modification will only reinforce our result.

to θ being realized. Thus, the debt structure itself is not informative about θ.[20] All creditors know the maturity date of their debt as well as the overall debt structure. Let w_j stand for 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., $\theta \geq \alpha \sum_{l=1}^{j-1} w_l$, or not.

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 get 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 withdrawal up to time 1; otherwise, he gets 0.

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

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, the withdrawing creditors in group j will split the remaining liquid assets θ_j.[21] Thus, the payoff for group j creditors from withdrawal is

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

This payoff specification is different from the one introduced in Section 1. The payoff from rolling over depends on whether the borrower can sustain all withdrawal to time 1, while that from withdrawing only depends on whether she can withstand the withdrawal from the current group. This asymmetry in payoff may increase the incentive for creditors to withdraw. Moreover, unlike in section 1, the payoff from withdrawing is not a constant when the borrower defaults. This payoff specification is similar to the one in Goldstein

[20] If the borrower only undertakes such 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.

and Pauzner (2005), in which the payoff from withdrawing is negatively dependent on aggregate withdrawal.[22]

In the following proposition we show that despite these differences, a sufficiently asynchronous debt structure, i.e., when maturity dates and hence the viability tests occur sufficiently frequently, prevents panic-based runs.

Proposition 2 *There exists $\alpha^* > 0$ such that any asynchronous debt structure with $J > \frac{1}{\alpha^*}$ eliminates the risk of panic-based debt runs.*

Our result is robust to the case when payoff from attacking (or withdrawal) does not depend on 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, as 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 application of our policy. Although the payoffs upon success or failure may depend on the fundamental and withdrawals, as long as they are bounded, 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 $\overline{p}$. Thus, a sufficiently asynchronous debt structure with $J > 1/\alpha^*(\overline{p}, \tau)$ eliminates the chance of panic-based runs.

5 Discussion

This debt run application of our theory 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 specified in Section 1). In this section, we consider some variations of our set up to understand what could make it harder or easier to dissuade the agents from attacking. Also, to understand what are 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 the information to the agents.

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

Full Disclosure

Suppose that the principal knows the underlying fundamental and discloses the true θ at the beginning. In addition, very frequently she discloses the information regarding past attacks. Consider $\theta < 1$ and suppose that the agents follow the strategy - attack regardless of the history. It is easy to see that this is an equilibrium, and in this equilibrium, the regime will fail after more than θ fraction of agents make their decision.

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, the currency attack game as in Morris and Shin (1998) - attacking is costly but would be profitable if agents can coordinate on doing that ($b_0 > 0, c_1 < 0$), while not attacking is not profitable ($b_1 = c_0 = 0$). 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 and one agent is equivalent to half mass, and the principal discloses all the information about θ and the past attack - whether the first agent has attacked or not. Suppose that $\theta \in (1/2, 1)$.[23] Then the second agent will attack if the first agent attack ($b_0 > b_1$) and does 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 will attack. Thus, it follows from backward induction that both agents will take the payoff dominant action - attack. This simple example shows that full disclosure may fail to dissuade the 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 agent play the payoff dominant action even generalizes to repeated games. In contrast with the folk theorem results, Dutta (2012) (also see 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 regardless of the current state, the agents will soon move to playing the payoff dominant action.[24]

In contrast, our selective partial disclosure policy is persuasive regardless of whether

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

[24] Zhou and Chen (2015) show that in an investment game when agents move sequentially and perfectly observe past actions, aggregate investment increases. They investigate the optimal sequencing strategy when the agents have heterogeneous payoff externalities.

attacking is the payoff dominant action or not. Does attacking being the payoff dominant action makes it harder to dissuade the agents from attacking? Let us now move to the comparative statics question - what makes it harder or easier to persuade?

Comparative Statics

In order to persuade the agents not to attack, the first hurdle that the principal needs to overcome is the agents' intrinsic reluctance, captured by the parameter p. One possible reason for why agents are more reluctant to follow the principal's recommendation is that agents' incentive is not aligned with the principal, since p is higher when attacking is the payoff dominant action ($b_1 < b_0$).

The following Proposition shows that if agents are more reluctant to not attack, a policy needs to be more diffused to be persuasive.[25] The argument is simple. By definition, α^* is such that, when $\alpha = \alpha^*$, given all other agents are playing the cutoff strategy $\hat{s}$ (for some $\hat{s}$), the marginal agent with signal $\hat{s}$ believes that the chance of success is p. Therefore, if p is higher, the belief needs to be more optimistic, and the belief can be made more optimistic when the group size α is smaller.

The second hurdle in persuading the agents to not attack is their private information. If τ is higher, then the marginal agent in group j with signal $\hat{s}_j$ is less optimistic about the success of the regime (see Equation 4). Thus, persuading the agents to ignore their private information and follow the principal's recommendation requires higher diffusion. The following proposition summarizes the comparative statics - how $J^*(p,\tau)$ changes with p and τ.

Proposition 3 *Given p, τ, if $J^*(p,\tau)$ is the critical diffusion needed to eliminate the coordination risk, then*

1. *For $p' > p$, $J^*(p',\tau) > J^*(p,\tau)$.*
2. *For $\tau' > \tau$, $J^*(p,\tau') > J^*(p,\tau)$.*[26]

The result follows from the argument preceding this proposition. Now, let us revisit the basic set up. We consider a regime change game where agents have noisy private infor-

[25] Note that our policy can work with any $p < 1$ (Theorem 1). If $p \geq 1$, then attacking becomes the dominant action and there is no room for persuasion.

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

mation as is standard in the global game literature. Is this private information environment essential? What if - the noisy information was public and accordingly agents shared a homogeneous belief over θ?

Homogeneous Belief

Consider the same set up as in Section 1 with the following modification. The agents receive a noisy public signal $s = \theta + \sigma\epsilon$ instead of private signals. The principal also observes this signal and announces the policy after learning s. Conditional on the realized signal s, the agents have a homogeneous belief over θ. Under homogeneous belief, the only way to dissuade the agents from attacking is to convince them "even if others attack, the regime is very likely to succeed". Otherwise, it is always a possible equilibrium where all the agents attack. Can sufficient diffusion dissuade the agents from attacking?

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 sufficiently frequently (or $J > J^*(s)$), it can make the probability that "θ is in the upper dominance region" high enough to dissuade the agents from attacking a viable regime. It is easy to see that the lower the s, the lower the $\mathbb{P}(\theta \geq 1|s)$ and consequently the harder it is 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 believe that if all the agents attack, then the regime will certainly fail. Then, regardless of however diffused the policy, all attack is a possible equilibrium.

In contrast, under heterogeneous beliefs, there is a uniform J^* such that regardless of the private signals, no agent attacks. Even the agent who receives a private signal $-\sigma/2$ and privately knows that $\theta = 0$, can be persuaded not to attack. The crucial difference with the heterogeneous belief case is as follows. Under noisy private information, when θ is low, all the agents may privately learn that $\theta < 1$, but it is not a common 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 succeed." However, she can still convince him that "others are not very likely to attack," and thus the regime is very likely to succeed.

Proposition 4 *Under the information environment with 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.[27]

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 common belief that $\theta < 1$ and thus no diffused policy can dissuade the agents from attacking. It is always a possible equilibrium where all the 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.

6 Conclusion

This paper proposes a simple policy, called repeated viability tests, to eliminate the strategic uncertainty in a global game of regime change. The repeated viability tests diffuses 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 to not 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 small enough group of agents will follow the recommendation as well (argument N_j).

[27] Under unbounded noise (for example, the Guassian 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.

This contributes to the 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 sufficiently asynchronous debt structure eliminates the possibility of panic.

Readers might wonder that since the agents move sequentially, it is possible that over time some exogenous shock hits the fundamental or new information arrives. Will sufficient diffusion be still persuasive? We relegate such extensions to the online appendix. We show that the result is robust to small perturbations to 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 think these are promising directions for future research. We briefly discuss these issues in the online appendix.

Appendix

Proof of Claim 1 $G(x,\tau,\alpha)$ is decreasing in α and $\lim_{\alpha\to 0} G(x,\tau,\alpha) = 1$ for all $x \in (0,1]$. When $x \to 0$, using L'Hospital rule, we get

$$\begin{aligned}\lim_{x\to 0} G(x,\tau,\alpha) &= \lim_{x\to 0} \frac{1}{f(F^{-1}(x)+\alpha\sqrt{\tau}x)} \cdot \frac{1}{\frac{1}{f(F^{-1}(x))}+\alpha\sqrt{\tau}} \\ &= \frac{1}{1+\alpha\sqrt{\tau}\times f(-\frac{1}{2})}.\end{aligned}$$

Since $f(-\frac{1}{2}) < \infty$, $\lim_{\alpha\to 0}\left(\lim_{x\to 0} G(x,\alpha)\right) = 1$. Let us define

$$x(\tau,\alpha) := \arg\min_{x\in[0,1]} G(x,\tau,\alpha) \text{ and } y(\tau,\alpha) := G(x(\tau,\alpha),\tau,\alpha). \tag{6}$$

Following theorem of maximum, we know $y(\tau,\alpha)$ is well defined and is continuous in α. Now for all $\alpha_1 < \alpha_2$, we have $y(\tau,\alpha_1) = G(x(\tau,\alpha_1),\tau,\alpha_1) > G(x(\tau,\alpha_1),\tau,\alpha_2) \geq G(x(\tau,\alpha_2),\tau,\alpha_2) = y(\tau,\alpha_2)$. The first inequality comes from the fact that $G(x,\tau,\alpha)$ is decreasing in α, and the second follows from the definition of $x(\tau,\alpha)$. Hence, $y(\tau,\alpha)$ is decreasing in α and $\lim_{\alpha\to 0} y(\tau,\alpha) = 1 > p$. By definition, α^* is such that $y(\tau,\alpha^*) = p$ and if $y(\tau,1) > p$, then $\alpha^* = 1$. Therefore, for any $\alpha < \alpha^*$, we have $G(x,\tau,\alpha) > p$ for all $x \in [0,1]$. $\square$

Proof of Proposition 2 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.

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

$$u(0,\theta,\{w_l\}_{l=1}^J) = \begin{cases} c_1(\theta,\{w_l\}_{l=1}^J) & \text{if } \theta \geq \alpha \sum_{l=1}^{j} w_l \\ b_0(\theta,\{w_l\}_{l=1}^J) & \text{if } \theta < \alpha \sum_{l=1}^{j} w_l. \end{cases}$$

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 $\theta \gtreqless \alpha \sum_{l=1}^{J} w_l$ is equivalent to $\theta \gtreqless \alpha \sum_{l=1}^{j} w_l$. 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

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

It follows from Assumption 1 that, regardless of $\hat{s}_j$ and A_j, the RHS of the above inequality is lower than $\overline{p} \equiv \frac{1}{1+\frac{\underline{n}}{\overline{m}}}$. Therefore, $\alpha < \alpha^*(\overline{p}, \tau)$ 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 3 By definition of $\alpha^*(p, \tau)$ and $y(\tau, \alpha)$, $\alpha < \alpha^*(p, \tau)$ is the necessary and sufficient condition for $y(\tau, \alpha) > p$. Also, $y(\tau, \alpha)$ is continuously decreasing in τ.

1. Given $p' > p$, if $\alpha < \alpha^*(p', \tau)$, then $y(\tau, \alpha) > p' > p$. Therefore, $\alpha^*(p', \tau) < \alpha^*(p, \tau)$. The inequality is strict since by continuity of $y(\tau, \alpha)$, there is some $\alpha \in (\alpha^*(p', \tau), \alpha^*(p, \tau))$ for which $y(\tau, \alpha) \in (p, p')$.

2. Given $\tau' > \tau$, we know $y(\tau, \alpha) > y(\tau', \alpha) > p$ for all $\alpha < \alpha^*(p, \tau')$. Therefore, $\alpha^*(p, \tau') < \alpha^*(p, \tau)$. $\square$

Proof of Proposition 4 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 succeed if it can sustain the attacks from the current group. Hence, if $\theta \geq \tilde{w} + \alpha$, it succeeds 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

$$\frac{F(\sqrt{\tau}(s-1))}{F(\sqrt{\tau}(s-1+\alpha))} > p, \tag{7}$$

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 (7) 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 (7) and thus the inductive argument N_j hold true. Moreover, since the LHS of Inequality (7) 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 (7) 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. $\square$

References

Acharya, Avidit, and Kristopher W Ramsay. 2013. “The calculus of the security dilemma.” *Quarterly Journal of Political Science*, 8(2): 183–203.

Alonso, Ricardo, and Odilon Câmara. 2016*a*. “Bayesian Persuasion with Heterogeneous Priors.” *Journal of Economic Theory*, 165: 672–706.

Alonso, Ricardo, and Odilon Câmara. 2016*b*. “Persuading Voters.” *The American Economic Review*, 106(11): 3590–3605.

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.

Bardhi, Arjada, and Yingni Guo. 2016. “Modes of Persuasion Toward Unanimous Consent.” http://yingniguo.com/wp-content/uploads/2017/04/BardhiGuo0407.pdf.

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. 2016. "Information Design, Bayesian Persuasion, and Bayes Correlated Equilibrium." *The American Economic Review*, 106(5): 586–591.

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.

Bueno de Mesquita, Ethan. 2014. "Regime Change and Equilibrium Multiplicity." *Unpublished manuscript, Harris School, University of Chicago*.

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.

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.

Eső, Péter, and Balazs Szentes. 2007. "Optimal information disclosure in auctions and the handicap auction." *The Review of Economic Studies*, 74(3): 705–731.

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."

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.

Mathevet, Laurent, Jacopo Perego, and Ina Taneva. 2017. "Information Design: The Epistemic Approach." http://www.laurentmathevet.com/wp-content/uploads/2017/02/MPT.pdf.

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.

Taneva, Ina A. 2016. "Information Design." https://sites.google.com/site/itaneva13/research.

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.

Zhou, Junjie, and Ying-Ju Chen. 2015. "Key leaders in social networks." *Journal of Economic Theory*, 157: 212–235.