

Coordination and Social Learning

Chong Huang*

November 17, 2016

Abstract

This paper studies the interaction between coordination and social learning in a dynamic regime change game. Social learning provides public information, to which players overreact due to the coordination motive. Coordination affects the aggregation of private signals through players' optimal choices. Such endogenous provision of public information results in informational cascades, and thus inefficient herds, with positive probability, even if private signals have an unbounded likelihood ratio property. An extension shows that if players can individually learn, there exists an equilibrium in which inefficient herding disappears, and thus coordination is almost surely successful.

Key Words: Coordination, social learning, inefficient herding, dynamic global game, common belief

JEL Classification: C72, C73, D82, D83

*Chong Huang, Paul Merage School of Business, University of California, Irvine, chong.h@uci.edu. I am indebted to my advisor, George Mailath, for his invaluable guidance and encouragement. I thank Satoru Takahashi (editor), two anonymous referees, Nageeb Ali, David Dillenberger, Hanming Fang, David Hirshleifer, Itay Goldstein, Philipp Kircher, Qingmin Liu, Steven Matthews, Stephen Morris, Mallesh Pai, Alessandro Pavan, Andrew Postlewaite, Alvaro Sandroni, Lones Smith, Jusso Välimäki, Muhamet Yildiz, as well as participants at the micro-lunch seminar at UPenn and attendees at the 2010 EconCon, the 2011 North American summer meeting of the Econometric Society, and the 2011 International Conference on Game Theory, for their helpful comments. All errors are mine.

1 Introduction

This paper analyzes the interaction between coordination and social learning. In many social and economic environments that feature coordination, such as political revolutions, innovations, investments, currency attacks, and bank runs, players must coordinate if they are to achieve a discrete increase in their payoffs. However, players are usually uncertain about the fundamentals that determine the payoffs from their coordination, which causes them to hesitate in choosing a coordination action. Hence, players look at previous players' actions and coordination results, which partially aggregate previous information about the fundamentals. In this way, players "socially learn" the fundamentals. But does social learning help players coordinate? If players' coordination can bring them higher payoffs (although they do not know this *ex ante*), will players eventually coordinate?

In a seminal paper, Smith and Sørensen (2000) show that to the extent that any decision maker's payoff is independent of other players' choices and all players have the same preference, the public history will successfully aggregate private signals if and only if the strength of private signals is unbounded.¹ However, such a well-known result is not obvious when players need to coordinate. In particular, when there are incentives to coordinate, players make inferences from their private signals not only about the unknown fundamentals, but also about their opponents' signals (and hence their opponents' actions). Therefore, even if a player is convinced by an extremely strong signal that coordination can bring good payoffs, she is not confident that her opponent is also receiving such signals. As a result, the player refrains from choosing a coordination action, coordination fails, and the history is prevented from aggregating private information.

In this paper, I study the interaction between coordination and social learning in a dynamic regime change game. There are two possible regimes: the status quo and an alternative. The game continues as long as the status quo is in place. In each period, there are two, new short-lived (one-period) players. They commonly observe previous plays, and they each receive one piece of private information about the status quo. Based on this information, they update their beliefs about the true state of the status quo, which is unknown but fixed. These two new, short-lived players then simultaneously choose to attack or not to attack the status quo. (Attacking the status quo is the coordination action that favors regime change.) A player who chooses to attack receives a positive payoff if the regime changes; she receives a negative payoff otherwise. Not attacking is a safe action, resulting in

¹The literature of observational learning was initiated by Banerjee (1992) and Bikhchandani, Hirshleifer, and Welch (1992). These two papers show that informational cascades and herding behavior emerge when each player receives a binary private signal. Lee (1993) analyzes a social learning model with continuous action space and binary signal space, in which inefficient herding does not arise, because no information goes unused. Bikhchandani, Hirshleifer, and Welch (1998) and Chamley (2004) provide surveys of this literature.

a zero payoff whether the regime changes or not. The strength of the status quo is drawn at the beginning of the game from a set consisting of three elements: weak, medium, and strong. If the status quo is weak, an attack by at least one player changes the regime; if the status quo is medium, then synchronous coordination (i.e., both players choose to attack) is necessary and sufficient to trigger regime change; if the status quo is strong, the status quo can never be beaten.

The main result in this paper is that in any equilibrium, once coordination is commonly known to be necessary for regime change, a herd emerges with positive probability. In such a herd, players refrain from attacking, ignoring their private signals. Unlike in Smith and Sørensen (2000), this result holds regardless of the strength of a player’s signal. It then follows that the medium status quo survives with positive probability in any equilibrium.

That players may join a herd even if they may receive extremely strong private signals arises from the interaction between coordination and social learning. Let’s take as a benchmark the equilibrium outcomes conditional on the weak state. In the weak state, coordination is not necessary and, therefore, the status quo is in place only if there is no attack. But, as the probability of attacking decreases, the “no-attack” outcome provides players with less information. This slow social learning is eventually dominated by players’ extremely informative private signals. As a result, in any equilibrium, if the status quo is weak, an inefficient herd never forms, and the regime changes almost surely.

By contrast, when coordination is commonly known to be necessary (i.e., the weak status quo is ruled out by a previous failed attack that was launched by one player), the current players need to have a sufficiently high prior belief about the medium status quo in order to attack. Hence, when the public history makes a player pessimistic *ex ante*, even if she receives an extremely informative private signal favoring the medium state, she is not confident that her opponent is also receiving an “attack” signal. Hence, in any equilibrium, inefficient herds emerge in the medium state with positive probability, and the regime may not change. Put differently, in the medium state, informational cascades result in the impossibility of coordination, and social learning is a source of coordination failure.²

Coordination incentives lead to such different social learning outcomes because the realizations of players’ signals and the players’ information structure both matter for the possibility of coordination. Specifically, without coordination, because players are uncertain

²“Herd behavior” and “informational cascade” are significantly different terms, as Smith and Sørensen (2000) pointed out. In my model featuring coordination, an informational cascade emerges if the information revealed from the public history makes coordination impossible despite players’ private signals. Obviously, if an informational cascade emerges, players join a “no-attack” herd. However, a no-attack herd may be caused purely by the failure of coordination on the risky action. In particular, if one player chooses not to attack, regardless of her private signals, her opponent will choose not to attack too. If such a coordination failure happens even when coordination is possible, the resulting no-attack herd is not caused by an informational cascade.

only about the fundamentals, their prior beliefs can be overturned by extremely strong realized signals and hence they will never ignore their signals. By contrast, with coordination, players care also about their opponents' signals. Because the precision of the players' signals is finite, any player's estimation of her opponent's signal cannot be arbitrarily accurate. Hence, when the prior belief about the medium status quo is low enough, a player will not trust that her opponent is receiving a signal favoring the medium status quo, and she will not trust that her opponent believes that she is receiving the signal favoring the medium status quo either.

In an extension of the core model, I allow players to individually learn the strength of the status quo. With individual learning, players' estimations about their opponents' signals will become extremely accurate as time goes by. In the limit, if one player is almost convinced by her own private signal that the status quo is medium, she will be confident that her opponent is also convinced by a signal that the status quo is medium, even if the prior belief about the medium status quo is low. It is therefore possible for the two players to coordinate an attack, and individual learning helps them coordinate attacks, because it increases the accuracy of their estimations of their opponents' signals over time. Thus, there are equilibria in which players never ignore their signals and, in such equilibria, informational cascades never arise: if the status quo is medium, the regime changes almost surely.

1.1 Related Literature

The social learning literature (for example, Banerjee, 1992, Bikhchandani, Hirshleifer, and Welch, 1992, and Smith and Sørensen, 2000) has extensively discussed the necessary and sufficient conditions for informational cascades and herding behavior in environments lacking coordination incentives. However, coordination does interact with social learning in various social, economic, and financial environments. This paper's main contribution is the discovery that coordination incentives dramatically change social learning outcomes. With coordination, the realizations of signals and the players' information structure both matter for players' choices. Therefore, the way the public history aggregates information in environments with coordination differs significantly from how it does so in environments without coordination. I show that even if players receive signals of unbounded strength, which is sufficient for the no-informational cascade result when there are no coordination incentives, informational cascades may start. Therefore, in any equilibrium, inefficient herding may exist.

The paper contributes also to the discussion of the social value of public information, which has been studied in a vast literature pioneered by Hirshleifer (1971). Morris and Shin (2002, 2003) analyze the effects of public information in a model with payoff complementarities. They show that increased provision of public information is more likely to lower social

welfare when players have more precise private signals. Angeletos and Pavan (2007) prove that when the degree of coordination in the equilibrium is higher than the socially optimal one, public information can reduce equilibrium welfare. In this literature, public information is exogenous. In my paper, public information evolves endogenously as a result of players' decisions. Therefore, coordination directly affects the provision of public information. Since the public history may fail to aggregate all private signals and may therefore provide biased public information, the players fail to coordinate.

There is a strand of social learning models that discuss herding behavior and asynchronous coordination. In Dasgupta (2000), first movers and late movers can coordinate. Therefore, first movers have incentives to signal their private signals by choosing the coordination action. Under private signal structures with an unbounded likelihood ratio property, there are "weak herd behaviors," in which players do not ignore their private signals. In Choi (1997), coordination is also asynchronous, payoff complementarities arise only from network effects, and learning is complete once an option is taken. Asynchronous coordination imposes a weaker belief requirement than synchronous coordination, which is my paper's key ingredient. In addition, it can be shown that in my model if coordination is asynchronous, inefficient herds never form.

The present paper contributes also to the literature on global games, initiated by Carlsson and Van Damme (1993). Absent of dynamic features, static global games have been applied to currency attacks (Morris and Shin, 1998), bank runs (Goldstein and Pauzner, 2005), debt crises (Morris and Shin, 2004), and political changes (Edmond, 2013). These static regime change games are solvable by iterated elimination of strictly dominated strategies for fixed prior beliefs and arbitrarily informative private signal structures (see Morris and Shin, 2003). In my model, if no attack has yet occurred, players are in a static global game. But, in this static global game, multiple Bayesian Nash equilibria may exist. Such multiplicity follows two characteristics of the model. First, because the state space is discrete, there are some prior beliefs resulting in multiple Bayesian Nash equilibria, no matter how precise private signals are. The necessity of the connectedness of the state space for the equilibrium uniqueness is discussed in Carlsson and Van Damme (1993). Second, the precision of private signals is fixed, while public beliefs evolve endogenously over time.

In addition, this paper contributes to the growing literature on dynamic global games with endogenous information. Angeletos, Hellwig, and Pavan (2007) incorporate both individual learning and social learning into a dynamic regime change game.³ They consider a continuum

³Dynamic regime change games are studied as examples of dynamic global games. Dasgupta (2007) studies a two-period dynamic global game that allows asynchronous coordination. Tarashev (2007) discusses a two-period currency crisis model, in which the interest rate informs investors about other investors' actions (and hence also about their information). Other papers contributing to this growing literature include Giannitsarou and Toxvaerd (2007) and Heidhues and Melissas (2006).

of long-lived agents, each individually learning the true fundamentals. Players also learn from the publicly observable fact that the regime has not changed. Due to the assumption of continuous players, a period in which no attack occurs yields no new information to players. The authors show that some status quos that can be overthrown nevertheless survive.⁴ My model differs from that of Angeletos, Hellwig, and Pavan (2007) mainly in that players, being short-lived, cannot individually learn. By assuming this, I can abstract away players' individual learning and focus on the interaction between social learning and coordination. In an extension of the core model, I allow players individually to learn. Unlike in Angeletos, Hellwig, and Pavan (2007), in my model the absence of an attack in a given period can be informative because there are finitely many players in each period. Moreover, because of the finite state space, the status quo that can be overthrown will be overthrown almost surely.

The present paper is also related to Chen and Suen (2016), who study a different dynamic global game. In their model, a one-period global game of regime change with a continuum of players is played in every period. They assume that regimes in different periods are independent, and so attacks will surely occur in any period. However, the model of the world, which is either tranquil (regime change is difficult) or frantic (overthrowing the regime is easy), is constant (or follows a Markov process) and is unknown to players. Hence, in every period, players will first make inferences about the model of the world from all previous regime change outcomes. Since such inferences are not drawn from previous players' actions, and players never ignore their private information, herding or informational cascades do not emerge in Chen and Suen (2016). In my paper, because there are finitely many players in each period, attacks do not occur even if they can occur with positive probability. Since not only the previous regime change outcomes but also the previous players' decisions are informative about the fixed fundamentals, herding or informational cascades may emerge.

The rest of the paper is organized as follows. In Section 2, I introduce a dynamic regime change game and provide an algorithm to characterize all equilibria. In Section 3, I study the effects of social learning on both the dynamics of attacking and the regime change outcome. Section 4 is devoted to two extensions of the core model. Section 5 concludes. All omitted proofs are presented in the appendix.

⁴ In Dasgupta, Stewart, and Steiner (2012), players can also individually learn, but the success of coordination does not require synchronous actions. In contrast to my results on informational cascades and herding behavior, they show that coordination failure almost never arises in a sufficiently long asynchronous coordination game.

2 A Dynamic Regime Change Game

2.1 The Model

Time is discrete and indexed by $t \in \{1, 2, \dots\}$. There are two possible regimes: the status quo and an alternative. Denote the state of the regime at the end of period t by $R_t \in \{0, 1\}$: $R_t = 0$ refers to the status quo, and $R_t = 1$ refers to the alternative. Assume that the regime is in the status quo at the beginning of the game; hence, $R_0 = 0$. If $R_t = 1$ for some t , then $R_\tau = 1$ for all $\tau > t$. That is, once the regime changes, it remains in the alternative state forever. The strength of the status quo is described by $\theta \in \Theta \equiv \{w, m, s\}$, where $w, m, s \in \mathbb{R}$ with the order⁵ $w < m < s$. At the beginning of the game, θ is chosen by nature according to a common prior μ_1 , where $\mu_1(\theta) > 0, \forall \theta \in \Theta$. Once picked, θ is fixed forever.

In each period t , there are two new, short-lived players. Each player i chooses $a_{it} \in \{0, 1\}$, where $a_{it} = 1$ denotes “attack,” and $a_{it} = 0$ denotes “no attack.” Player i ’s ex-post payoff depends on both her choice and the state of the regime:

$$u_{it} = (1 - R_{t-1})a_{it}(R_t - c) + R_{t-1}(1 - c).$$

Hence, suppose that the regime changes in period t . If player i attacks, she receives a payoff of $1 - c$ ($c \in (0, 1)$); otherwise, she receives a payoff of 0. Once the regime changes, the game ends.⁶ Conditional on $R_{t-1} = 0$, whether the regime changes or not in period t depends on both the strength of the status quo θ and the number of attacks $a_{1t} + a_{2t}$. For the weak status quo, one player’s attack is sufficient for regime change; for the medium status quo, an attack by one player is not enough, but an attack by two players can trigger regime change; if the status quo is strong, the regime never changes. The following table summarizes the regime change outcomes conditional on $R_{t-1} = 0$:

$a_{1t} + a_{2t}$	$\theta = w$	$\theta = m$	$\theta = s$
0	$R_t = 0$	$R_t = 0$	$R_t = 0$
1	$R_t = 1$	$R_t = 0$	$R_t = 0$
2	$R_t = 1$	$R_t = 1$	$R_t = 0$

Before making a decision, period- t player i observes a private signal $x_{it} = \theta + \xi_{it}$. Here,

⁵Notations w , m , and s refer to “weak,” “medium,” and “strong,” respectively.

⁶This assumption, together with the assumption that the number of attacks in each previous period can be publicly observed, implies that subsequent players may delete some states from the support of their beliefs. This is a key assumption of dynamic regime change games; see, for example, Angeletos, Hellwig, and Pavan (2007). If, however, I assume that the game does not end even when the regime changes, that is, when the previous attack outcomes cannot be observed, then players in each period will play a static global game because they never know that their coordination is necessary.

$\xi_{it} \sim \mathcal{N}(0, 1/\beta)$, where $\beta \in \mathbb{R}_{++}$ denotes the precision of the players' private signals.⁷ The noise term ξ_{it} is independent of θ and independent across i and across t . Thus, all players' private signals are conditionally independent. In addition to private signals, at the beginning of any period $t \geq 2$, players observe the public history about the number of players attacking. Formally, period- t players observe the public history $h^t = (b_1, \dots, b_{t-1})$, where $b_\tau \in \{0, 1, 2\}$ is the number of players attacking in period τ , for all $\tau < t$. Let H^t be the set of all possible public histories at the beginning of period t . I define the strategy of player i in period t as $s_{it} : H^t \times \mathbb{R} \rightarrow \{0, 1\}$. Here, $s_{it}(h^t, x_{it})$ is the action player i chooses, given the public history h^t and the private signal x_{it} . Let $\mu_t|h^t$ be period- t players' common prior belief about θ , conditional on the public history h^t . I call μ_t the public belief in period t .

Definition 1 *An assessment $\{(s_{it})_{t=1, \dots}^{i=1, 2}, (\mu_t)_{t=1, \dots}\}$ is an equilibrium if*

1. *for any t , given μ_t , (s_{1t}, s_{2t}) forms a Bayesian Nash equilibrium in the static game; and*
2. *μ_t is calculated by Bayes' rule on the equilibrium path.*

The first part of the definition is a natural requirement of the assumption that players are short-lived. Because period- t players have no intertemporal incentives to make decisions, their strategies need to form a Bayesian Nash equilibrium in a static game, given their public belief μ_t . The second part of the definition specifies how to calculate public beliefs on the equilibrium path. In fact, as in the definition of a sequential equilibrium, the consistency requirement should be imposed on the off-equilibrium path. However, since players' strategies must form a static Bayesian Nash equilibrium, their equilibrium strategies are not affected by plays on the off-equilibrium path. Therefore, to simplify the analysis, I require only that public beliefs on the equilibrium path be calculated by Bayes' rule.

2.2 Equilibrium Characterization

From the definition, an equilibrium can be characterized in two steps: first, given any μ_t , calculate (s_{1t}, s_{2t}) , which constitutes a Bayesian Nash equilibrium in period t ; second, given μ_1 , h^t , and $(s_{i\tau})_{\tau=1, t-1}^{i=1, 2}$, employ Bayes' rule to calculate μ_t . For simplicity, if the environment

⁷By the Gaussian assumption, it is easy to calculate players' belief updates. In addition, the Gaussian assumption has a very clear description of the precision of players' signals. However, this assumption is not necessary. The distribution of the signal can be fairly general, and the assumptions I have to make are the following: (1) conditional on θ , the players' private signals are independent and identically distributed; (2) the support of x_{it} is an open interval $(\underline{X}, \bar{X}) \subset \mathbb{R}$, and the conditional density $f(x|\theta)$ of the signal is strictly positive for all $x \in (\underline{X}, \bar{X})$ and all $\theta \in \Theta$; (3) unbounded likelihood ratio: $\lim_{x \rightarrow \underline{X}} f(x|\theta)/f(x|\theta') = +\infty$ and $\lim_{x \rightarrow \bar{X}} f(x|\theta)/f(x|\theta') = 0$, whenever $\theta < \theta'$; and (4) monotone likelihood ratio: if $\theta < \theta'$, $f(x|\theta)/f(x|\theta')$ is strictly decreasing in x .

is understood clearly, I use the term “equilibrium” for both the Bayesian Nash equilibrium in the static game and the equilibrium in the dynamic game.

Three facts should be noted before a detailed analysis. First, since $\mu_1(\theta) > 0$ for all $\theta \in \Theta$, $\mu_t(\theta) > 0$ for all $\theta \in \Theta$ after the public history without any attack, because any type of status quo will survive along such a history. Second, if one player attacks in period t and $R_t = 0$, players in the subsequent periods immediately learn that $\theta \neq w$; that is, $\mu_\tau(w) = 0$ for all $\tau > t$. Third, if both players attack in period t and $R_t = 0$, subsequent players immediately learn that $\theta = s$. Therefore, for the analysis of a static game, only three possible public beliefs are relevant: (i) $\mu_t(\theta) > 0$ for all $\theta \in \Theta$; (ii) $\mu_t(w) = 0$, $\mu_t(m) > 0$, and $\mu_t(s) > 0$; (iii) $\mu_t(s) = 1$. Note that in a monotone equilibrium, since $\mu_1(\theta) > 0$ for all $\theta \in \Theta$, $\mu_t(\theta) > 0$ implies that $\mu_t(\theta') > 0$ for $\theta < \theta'$ in any period t . Out of these three cases, the one with $\mu_t(s) = 1$ is trivial. Because no-attack is the dominant action in this case, the unique Bayesian Nash equilibrium is that both players refrain from attacking regardless of their private signals.

Now, suppose that $\mu_t(\theta) > 0$ for all $\theta \in \Theta$. Let $\rho(\cdot|x_{it})$ denote period- t player i 's posterior belief over Θ after receiving signal x_{it} . By Bayes' rule, the posterior belief about θ is

$$\rho(\theta|x_{it}) = \frac{\mu_t(\theta)\phi(\sqrt{\beta}(x_{it} - \theta))}{\sum_{\theta' \in \Theta} \mu_t(\theta')\phi(\sqrt{\beta}(x_{it} - \theta'))},$$

where $\phi(\cdot)$ is the standard normal pdf. Player i 's interim payoff from attacking given her signal x_{it} and player j 's strategy s_{jt} is:

$$\begin{aligned} & E_{x_{jt}} u_{it}(1, x_{it}, s_{jt}) \\ &= \rho(w|x_{it}) + \Pr(s_{jt} = 1, m|x_{it}) - c \end{aligned} \quad (1)$$

$$= \rho(w|x_{it}) + \rho(m|x_{it}) \Pr(s_{jt} = 1|m) - c \quad (2)$$

$$= \frac{\mu_t(w)\phi(\sqrt{\beta}(x_{it} - w))}{\sum_{\theta' \in \Theta} \mu_t(\theta')\phi(\sqrt{\beta}(x_{it} - \theta'))} + \frac{\mu_t(m)\phi(\sqrt{\beta}(x_{it} - m))}{\sum_{\theta' \in \Theta} \mu_t(\theta')\phi(\sqrt{\beta}(x_{it} - \theta'))} \Pr(s_{jt} = 1|m) - c. \quad (3)$$

Here, (1) implies (2) because players' private signals are independent conditional on θ . Note that $\rho(w|x_{it}) \rightarrow 1$ as $x_{it} \rightarrow -\infty$; hence, from the regime change rule, attacking is the dominant action for player i , when x_{it} is extremely negative. By continuity of the interim payoff function, there exists an $\underline{x}_t \in \mathbb{R}$ such that $E_{x_{jt}} u_{it}(1, x_{it}, s_{jt}) > E_{x_{jt}} u_{it}(0, x_{it}, s_{jt}), \forall x_{it} \leq \underline{x}_t$, and $\forall s_{jt}$. I call the set $(-\infty, \underline{x}_t]$ the *dominant region of attacking*. Similarly, there is an $\bar{x}_t \in \mathbb{R}$ such that $E_{x_{jt}} u_{it}(1, x_{it}, s_{jt}) < E_{x_{jt}} u_{it}(0, x_{it}, s_{jt}), \forall x_{it} \geq \bar{x}_t$, and $\forall s_{jt}$. I call the set $[\bar{x}_t, +\infty)$ the *dominant region of not attacking*. Therefore, in the case $\mu_t(\theta) > 0$ for all $\theta \in \Theta$, players in period t play a static global game. Proposition 1 below not only proves the existence of a Bayesian Nash equilibrium, but also provides the equation to characterize the equilibrium in this case.

Proposition 1 *In a static game with $\mu_t(\theta) > 0$ for all $\theta \in \Theta$, a Bayesian Nash equilibrium exists. In any Bayesian Nash equilibrium, players follow a symmetric cutoff strategy with threshold point⁸ $x_t^* \in \mathbb{R}$:*

$$s_t^* = \begin{cases} 1, & \text{if } x \leq x_t^*, \\ 0, & \text{if } x > x_t^*. \end{cases}$$

In addition, x_t^ solves the equation*

$$G(x, \mu_t) = \frac{\mu_t(w)\phi(\sqrt{\beta}(x-w))}{\sum_{\theta' \in \Theta} \mu_t(\theta')\phi(\sqrt{\beta}(x-\theta'))} + \frac{\mu_t(m)\phi(\sqrt{\beta}(x-m))}{\sum_{\theta' \in \Theta} \mu_t(\theta')\phi(\sqrt{\beta}(x-\theta'))} \Phi(\sqrt{\beta}(x-m)) - c = 0, \quad (4)$$

where $\Phi(\cdot)$ is the standard normal cdf.

Now consider the boundary public belief case: $\mu_t(w) = 0$, $\mu_t(m) > 0$ and $\mu_t(s) > 0$. Because $\mu_t(w) = 0$, there is no *dominant region of attacking*. Since $\mu_t(s) > 0$, the *dominant region of not attacking* still exists. Hence, $\mu_t(m) > 0$ implies that the regime may change if and only if both players choose to attack. Hence, players in this case are playing a coordination game. Obviously, the strategy profile with coordination failure (i.e., both players refrain from attacking) is an equilibrium.

Let's analyze whether there exists an equilibrium in which period- t players attack with positive probability. I call such an equilibrium an *equilibrium with attacks*. Because the only state in which the regime can change is $\theta = m$, players choose to attack only if both their beliefs about $\theta = m$ and their beliefs about their opponents choosing to attack are sufficiently high. Therefore, if players' public beliefs about $\theta = m$ are high (players are optimistic), cooperation is possible; conversely, when players' public beliefs about $\theta = m$ are low (players are pessimistic), then even if one player observes an extremely negative signal and is convinced that $\theta = m$, she will not attack. This is because she believes that the probability of her opponent observing a signal favoring $\theta = m$ is very low. Proposition 2 below formally shows conditions under which an equilibrium with attacks exists.

Proposition 2 *In the static game with $\mu_t(w) = 0$, $\mu_t(m) > 0$, and $\mu_t(s) > 0$, there is a unique $\tilde{\mu} \in (0, 1)$ such that*

1. *if $\mu_t(m) < \tilde{\mu}$, there is no equilibrium with attacks;*
2. *if $\mu_t(m) > \tilde{\mu}$, there are at least two equilibria with attacks that are symmetric and in cutoff strategies with the threshold bounded from below by $m + \Phi^{-1}(c)/\sqrt{\beta}$; and*

⁸To simplify the notation, I denote a strategy by x_t^* when there is no confusion. In particular, $x_t^* = -\infty$ represents the strategy of not attacking for all signals, and $x_t^* = +\infty$ represents the strategy of attacking for all signals.

3. if $\mu_t(m) = \tilde{\mu}$, there exists a unique $\tilde{x} \in \mathbb{R}$ such that there is a unique equilibrium with attacks in which both players employ cutoff strategies with threshold point \tilde{x} .

The key step in proving Proposition 2 is to analyze the solution to equation (4), with the parameter $\mu_t(w) = 0$, $\mu_t(m) > 0$, and $\mu_t(s) > 0$: if $G(x, \mu_t) = 0$ does not have a solution, then the no-attack strategy profile is the unique equilibrium; if $G(x, \mu_t) = 0$ has a solution x_t^* , then there exists an equilibrium in which players employ a symmetric cutoff strategy with threshold point x_t^* . The parameter $\tilde{\mu}$ is determined by the following equation:

$$\max_{x \in \mathbb{R}} G(x, \tilde{\mu}) = 0, \quad (5)$$

and \tilde{x} is the solution to the equation $G(x, \tilde{\mu}) = 0$.

With Proposition 1 and Proposition 2, the following algorithm characterizes all equilibria of the dynamic model:

1. in any period t , given μ_t , compute all solutions to $G(x; \mu_t) = 0$ and pick any solution x_t^* to be the threshold point in period t . If $G(x; \mu_t) = 0$ does not have a solution, then let $x_t^* = -\infty$;
2. on the equilibrium path, conditional on $R_t = 0$, and given μ_t , x_t^* , and b_t , employ Bayes' rule to calculate μ_{t+1} .

3 Dynamics of Attacking, Regime Change, and Social Learning

In this section, I describe the equilibrium dynamics of attacking, and, based on these dynamics, I analyze the regime change outcomes conditional on the strength of the status quo. Because social learning is the driving force behind the dynamics of attacking, I investigate how social learning plays a role in the dynamics of attacking and in determining regime change outcomes.

Let $\hat{h}^t \equiv (0, \dots, 0)$ denote the history without any attack. When $\theta = w$, \hat{h}^t is the only history leading to $R_{t-1} = 0$. Thus, the dynamics of attacking along \hat{h}^t determine the eventual outcome of the regime change, conditional on $\theta = w$.

Proposition 3 *Fix any equilibrium. Along \hat{h}^t , $\mu_t(w) \rightarrow 0$ and $\Pr(b_t > 0 | \hat{h}^t) \rightarrow 0$. But if $\theta = w$, the regime changes almost surely.*

The conclusions in Proposition 3 show that if coordination is not publicly known to be necessary for regime change, private signals of unbounded strength will completely eliminate informational cascades; that is, along \hat{h}^t , players will never ignore their private signals when

making decisions. Hence, the existence of a dominant region of attacking implies that there are always private signals that lead to attacks and trigger regime change at state $\theta = w$. These results are essentially the same as those in Smith and Sørensen (2000), and they provide a good benchmark for analyzing the interaction between coordination and social learning when coordination is commonly known to be necessary for regime change. For the sake of completeness, the proof of Proposition 3 is included in the appendix.

Now let's turn to the analysis of the medium status quo. In the first period, because $\mu_1(\theta) > 0$ for all $\theta \in \Theta$, Proposition 1 implies positive probability of attacking in the first period. It directly follows that the medium status quo changes with positive probability. Hence, I focus on the question of whether the medium status quo changes almost surely.

For any prior beliefs, each player receives a private signal with the unbounded likelihood ratio property, and so the player will be convinced by her private signal that $\theta = m$ with positive probability. Therefore, it seems plausible to conjecture that in a nontrivial equilibrium (in which players choose to attack with positive probability whenever possible), with probability one, there is a period in which both players receive private signals favoring $\theta = m$, and, as a result, the medium status quo should change almost surely.

Surprisingly, I show that the medium status quo survives with positive probability in any equilibrium, including equilibria in which players attack with positive probability whenever possible. Consider the outcome $\bar{h}^\infty \equiv (1, 0, \dots)$, in which there is an attack by one player in the first period, but no subsequent attacks. If this outcome is reached, the medium status quo does not change. This outcome may be consistent with many equilibria: for any given subset $\mathcal{Q} \subset \mathbb{N} \setminus \{1\}$, only period- $t \in \mathcal{Q}$ players attack the status quo with positive probability, and players in all other periods adopt the no-attack strategy. In the strategy profile with finite \mathcal{Q} , the outcome \bar{h}^∞ is reached with strictly positive probability. Hence, in any equilibrium, unless \mathcal{Q} is unbounded, the regime survives with positive probability if the status quo is medium.

Suppose there is an equilibrium in which \mathcal{Q} is unbounded. Due to the failed attack in the first period, players rule out the weak state. By Proposition 2, in any $t \in \mathcal{Q}$, players employ the cutoff strategy with the threshold x_t^* bounded from below. Consequently, the probability of attacking is bounded away from 0. Hence, the no-attack outcome in period $t \in \mathcal{Q}$ makes subsequent players more pessimistic, because it implies that both period- t players receive signals above x_t^* . This then leads to a decrease in $\mu_{t+1}(m)$, and so along the outcome $\bar{h}^t \equiv (1, 0, \dots)$, players will become more and more pessimistic ex ante. Recall that a necessary condition for players to attack with positive probability in a period $t \in \mathcal{Q}$ is $\mu_t(m) \geq \tilde{\mu}$, where $\tilde{\mu}$ is a constant for any given β ; then, after a period $T \in \mathcal{Q}$, $\mu_T(m)$ will be smaller than $\tilde{\mu}$, implying that subsequent players cannot coordinate attacks. Hence, \mathcal{Q} cannot be unbounded, and so the median status quo survives with positive probability.

Formally:

Proposition 4 *There exists $Q \in \mathbb{N}$, such that in any equilibrium, along the outcome \bar{h}^∞ , there are at most Q periods in which players' strategies specify positive probabilities of attack. Hence, in any equilibrium, $\mathbb{P}_m(\bar{h}^\infty) > 0$, and the medium status quo survives with positive probability.*

The outcome \bar{h}^∞ is just an example of all outcomes that result in the survival of the medium status quo and are realized with positive probability in any equilibrium. All these outcomes share three features: (1) there is no period in which both players attack; (2) there are finitely many periods (at least one) in which one player attacks; and, (3) after a certain period in which one player attacks, no subsequent attack occurs.

Proposition 4 shows that herding is an outcome in any equilibrium: players in later periods join the no-attack herd, ignoring their own private signals no matter how strong such signals are. Importantly, this result is not due to pure coordination failure, but arises from the informational cascades formed by players' social learning. In particular, players learn from the failed attack in the first period that the status quo is not weak, and therefore they need to coordinate. Hence, they learn from the no-attack history after the first period that the status quo is very likely to be strong.

The effect of the interaction between social learning and coordination can be seen by comparing the equilibrium outcomes in the weak state with those in the medium state. Similar to that in Smith and Sørensen (2000), the no-attack herd does not emerge in the weak state. This is because players have a dominant region of attacking, such that their own signals may determine their actions: once the realized signal lands in the dominant region of attacking, a player will choose to attack, independent of her opponent's choice. In consequence, players never ignore their signals and, therefore, neither informational cascades nor no-attack herds will emerge.

If, however, the status quo is medium, players will join a no-attack herd with positive probability in any equilibrium. The difference stems from the coordination requirement in the medium state and the signal structure of global games: when coordination is commonly known to be necessary for regime change, not only the realizations of signals but also the precision of the players' private information matters for the players' decisions. In particular, in the medium state, once the weak status quo is ruled out, it is common knowledge that coordination is necessary for successful attacks. As a result, one player will make her decision based on her estimations about both the status quo's strength and her opponent's signal. On the one hand, along the outcome \bar{h}^∞ , players become increasingly pessimistic about the medium status quo and, eventually, the public belief about $\theta = m$ will decrease below $\tilde{\mu}$. On the other hand, because period- t players' signals are $x_{kt} = \theta + \xi_{kt}$ ($k = 1, 2$), we

have $x_{jt} = x_{it} - \xi_{it} + \xi_{jt}$, and thus the variance of player i 's estimation of player j 's signal (conditional on player i 's signal x_{it}) is $2/\beta$ in period t . Since β is fixed, the accuracy of such an estimation is also fixed. Then, even if a player is convinced by an extreme signal realization that $\theta = m$, she believes that it is very unlikely that the other player is also receiving an “attack” signal and thus attacking. (Because players believe that $\theta = m$ is very unlikely ex ante, Bayes’ rule requires that the precision of one player’s private signal be large enough for her to receive a signal sufficiently favoring $\theta = m$.) Such a player refrains from attacking. Consequently, informational cascades and the no-attack herd form.

One interesting observation from both Proposition 3 and Proposition 4 is that the no-attack history delivers information to subsequent players. That is, players’ beliefs update along history \hat{h}^t in any equilibrium and along history \bar{h}^∞ in any nontrivial equilibrium. Such an observation differs significantly from that in Angeletos, Hellwig, and Pavan (2007), where players do not update their beliefs from the no-attack history. The difference stems from the different numbers of players: in this paper, there are two players in each period, while in Angeletos, Hellwig, and Pavan (2007), there is a continuum of players. With a continuum of players, if the threshold of attacking is finite, the law of large numbers implies that a positive measure of players attack; equivalently, if no player attacks, it must be that players adopt the no-attack strategy in that period, regardless of their signals. Hence, in Angeletos, Hellwig, and Pavan (2007), the no-attack outcome in any period is not informative about players’ signals, and hence cannot deliver any information to subsequent players.

4 Extensions

4.1 Individual Learning

In the model described in Section 2, no player can individually learn the true status quo, because the precision of the players’ signals is fixed at $\beta < +\infty$. In this section, I analyze the effect of players’ individual learning. In particular, I focus on whether the no-attack herd in Proposition 4 will emerge, and whether the medium status quo can survive with positive probability once individual learning is allowed.

A straightforward way to model individual learning is to extend the model in Section 2 by assuming that period- t players receive private signals of precision $t\beta$. Then, as time goes by, the precision of players’ signals monotonically increases and diverges to ∞ . In this way, players’ signals are increasingly precise, even though they are still short-lived, such that the model remains tractable.⁹

⁹In Angeletos, Hellwig, and Pavan (2007), players are long-lived, and so they can individually learn by collecting signals in each period. However, in their model, there is a continuum of players, so long-lived players behave myopically.

The most interesting case is $\theta = m$, especially when $\theta = w$ is ruled out by one player's failed attack. In such a case, players lose the dominant region of attacking, and it is common knowledge that they need to attack simultaneously to trigger regime change.

However, unlike Proposition 4, there are equilibria in which informational cascades do not form, and the medium status quo will change almost surely. The key reason for this result is that as the precision of players' signals increases over time and diverges to ∞ , $\tilde{\mu}_t$, the critical public belief above which period- t players can possibly attack, continues to decrease and converges to zero.

Lemma 1 *Suppose $\theta = w$ has been ruled out. There is $T \geq 1$, such that for all $t \geq T$, $\tilde{\mu}_t$ is decreasing in t . As $t \rightarrow +\infty$, $\tilde{\mu}_t$ converges to 0.*

By Proposition 2, there are equilibria in which attacks can occur with positive probability in period t if and only if $\mu_t(m) \geq \tilde{\mu}_t$. In addition, when $\mu_t(m) \geq \tilde{\mu}_t$, period- t players attack with a probability bounded away from zero. Then, since $\tilde{\mu}_t$ goes to 0 by Lemma 1 above, unless the status quo is abandoned, attacks can occur in infinitely many periods. Therefore, the medium status quo changes almost surely in such equilibria.

Proposition 5 *There exist equilibria in which the regime changes almost surely when $\theta = m$.*

Individual learning leads to regime change when $\theta = m$ in some equilibria, because an informational cascade never starts in such equilibria. That is, with individual learning, some players will not ignore their signals when making decisions. Yet, because even period-1 players observe normal signals of non-uniformly bounded strength, the reason for the disappearance of informational cascades differs from that given in Smith and Sørensen (2000). By contrast, players do not ignore their signals in this case precisely because, as their signals become more accurate, they can estimate their opponents' signals more accurately. Therefore, based on her signal, one player can better predict the other player's behavior. This is an application of "common learning," introduced by Cripps, Ely, Mailath, and Samuelson (2008): because of individual learning, players can form a common q -belief about the true θ infinitely often, and the medium status quo is attacked infinitely often (with the attacking probability bounded away from zero in each attacking period).

Proposition 5 differs from Angeletos, Hellwig, and Pavan (2007), which features coordination and individual learning. In Angeletos, Hellwig, and Pavan (2007), the strength of the status quo is distributed over the whole real line. Therefore, if all players attack, the status quo of the strength less than 1 cannot survive. However, the authors show that, given the prior belief, if the status quo's strength is in an open interval $(z, 1)$, where z is strictly less than 1, the regime will not change in any equilibrium. This result derives from the assumption that the strength of the status quo is drawn from a continuous state space, which

causes higher frictions for players' coordination. In this paper, I show that with a discrete state space, individual learning allows players to form common beliefs infinitely often and thus coordinate attacks infinitely often. Therefore, if the status quo can be overthrown, it will be overthrown almost surely in some equilibria.

4.2 Persistent Status Quo

In the model described in Section 2, the strength of the status quo is drawn at the beginning of the game and is forever fixed. One interesting extension is that the strength of the status quo persists, but is not constant over time. In this section, I analyze the regime change result when the strength of the status quo is persistent. A more detailed analysis is included in an online appendix.

To model the persistent status quo, I consider that θ , the strength of the status quo, varies over time. Without loss of generality, I consider two cases. First, both the medium status quo and the strong status quo may become weak. Formally, I assume that in this case,

$$\Pr(\theta_{t+1} = w | \theta_t = k) = \alpha_k \in (0, 1), \forall k \in \{w, m, s\}. \quad (6)$$

Without loss of generality, I assume that

$$1 > \alpha_w > \alpha_m > \alpha_s > 0. \quad (7)$$

In such a case, in every period, the weak status quo is in the support of players' beliefs. In particular, even if the posterior belief about $\theta_{t-1} = w$ in period $t - 1$ is zero (for example, $b_{t-1} = 1$ but the regime does not change), then in period t , θ_t can become weak with a probability of at least α_s . Therefore, in any period t , players have their dominant region of attacking, and the probability that any player will attack is bounded from below by the attacking probability when the public belief is $\mu_t(w) = \alpha_s$, $\mu_t(m) = 0$, and $\mu_t(s) = 1 - \alpha_s$. Consequently, in any period, the regime changes with a probability bounded away from zero, which directly implies that the regime will change almost surely.

Second, I consider the case where the weak status quo is constant, but the strength of the status quo may vary between the medium state and the strong state. By Proposition 3, the weak status quo changes almost surely. Therefore, without loss of generality, I only need to consider the case where $\theta_t = w$ is ruled out. Thus, I assume that

$$\Pr(\theta_{t+1} = w | \theta_t = m) = \Pr(\theta_{t+1} = w | \theta_t = s) = 0, \quad (8)$$

and¹⁰

$$\Pr(\theta_{t+1} = m | \theta_t = m) = \Pr(\theta_{t+1} = s | \theta_t = s) = \gamma \in \left(\frac{1}{2}, 1\right). \quad (9)$$

¹⁰The symmetry assumption here is just for simplicity. The result is robust when the transition rule is asymmetric.

Suppose that in an equilibrium $\theta_T = w$ has been ruled out by the beginning of period T . Therefore, without any other belief updating, there is a stationary objective distribution of θ over $\{m, s\}$, in which $\Pr(\theta = m) = 1/2$. Furthermore, without any other endogenous belief updating, if $\mu_T(m) > 1/2$, $\mu_t(m)$ is strictly decreasing for all $t \geq T$ and converges to $1/2$, and if $\mu_T(m) < 1/2$, $\mu_t(m)$ is strictly increasing for all $t \geq T$ and converges to $1/2$. The regime change outcome in this case depends on whether $\tilde{\mu} \geq 1/2$, because the coordination of an attack by period- t players depends, by Proposition 2, on whether $\mu_t(m) \geq \tilde{\mu}$.

When $\tilde{\mu} \geq 1/2$, the status quo survives with positive probability in any equilibrium, because an informational cascade will start in an equilibrium with attacks. In particular, suppose that there is an equilibrium in which period- t players attack with positive probability if $t \in \mathcal{Q}$. In that case, a long history without any attack will make players sufficiently pessimistic, such that $\mu_{T'} < \tilde{\mu}$ for some $T' \in \mathcal{Q}$. Since $\tilde{\mu} \geq 1/2$, $\mu_t(m)$ is either strictly decreasing to $1/2$ (when $\tilde{\mu} > \mu_{T'} > 1/2$) or strictly increasing to $1/2$ (when $\tilde{\mu} > 1/2 > \mu_{T'}$), it follows that $\mu_t(m) < \tilde{\mu}$ for all $t \geq T'$. Therefore, no attack can occur after period T' . Since T' is finite, there are at most Q attacks, where Q is also finite. Because the status quo can survive such Q attacks with positive probability, the regime change probability is strictly less than one.

By contrast, when $\tilde{\mu} < 1/2$, there are equilibria in which the regime changes almost surely, because the belief updating due to the exogenous state varying can make players sufficiently optimistic. The idea of constructing such an equilibrium is as follows. In any period $t > T$, if $\mu_t(m) \geq \tilde{\mu}$, period- t players attack if and only if their signals land below a threshold that is bounded from below according to Proposition 2. If $\mu_t(m) < \tilde{\mu}$, players refrain from attacking. Such a no-attack outcome will not change subsequent players' beliefs, because period- t players' choices are not based on their signals. Since $\mu_t(m) < \tilde{\mu}$, $\mu_t(m) < 1/2$. Therefore, while players refrain from attacking, the public belief keeps increasing. Importantly, suppose that in period τ , $\mu_\tau(m) = 1 - \gamma$ (because of the transition rule in equation (9), $\mu_\tau(m) \geq 1 - \gamma$); then, there is T' , such that $\mu_{\tau+T'}(m) > \tilde{\mu}$, because $\mu_t(m)$ converges to $1/2$. Hence, the public belief about $\theta = m$ is higher than $\tilde{\mu}$ at least every T' periods, implying that players can attack (with a probability bounded away from zero) every T' periods. Since the regime changes with a probability bounded away from zero in infinitely many periods, the regime will change almost surely.

The key reason for the different regime change outcomes in these two cases is that the exogenous change in the status quo makes the public beliefs adjust toward $1/2$. Once the public belief drops below $\tilde{\mu}$, players stop attacking, ignoring their private information. It follows that subsequent players will stop learning from the public history (as in the model with constant status quos), and so the only reason for the change in their public beliefs is that the status quo may change exogenously. Since the exogenous change in the status quo

makes the public beliefs adjust toward $1/2$, in the case where $\tilde{\mu} \geq 1/2$, once the public belief drops below $\tilde{\mu}$, it stays there forever; but in the case where $\tilde{\mu} < 1/2$, the public belief will go above $\tilde{\mu}$ after finitely many periods, making attacks possible again.

5 Conclusion

In this paper, I analyze the interaction between coordination and social learning in a dynamic regime change game. I show that contrary to the existing conclusion in the social learning literature, informational cascades will form with coordination, even if players' signals have the unbounded likelihood ratio property. Hence, in any equilibrium, inefficient herding emerges, and the status quo that can be overthrown by players' coordination survives with positive probability.

The key intuition for this result is that players' decisions depend not only on the realizations of their signals but also on their information structure. When players cannot individually learn, the accuracy of one player's estimation about the other player's private signal is bounded. Therefore, as the public history makes the players pessimistic ex ante, they refrain from attacking even if their private signals favor the attacking choice, because they cannot be confident that their opponents are also receiving "attacking" signals. If players are allowed to individually learn, one player's estimation about her opponent's signal is increasingly accurate over time and becomes perfect. Thus, even if players are ex-ante pessimistic, when one player receives an extremely strong signal favoring the attack action, she will believe that her opponent is also receiving a signal favoring $\theta = m$, because her estimation about the other player's signal is accurate. Hence, with individual learning, there are some equilibria in which the status quo is attacked infinitely many times, and thus the status quo that can be overthrown will change almost surely.

In addition to its theoretical contribution, this paper also has an applied contribution. There are many social, economic, and financial environments in which social learning and coordination interact with each other. For example, in foreign direct investments, new investors are uncertain about the host country's investment environments. They look at past investors' decisions and investment outcomes to socially learn. Given the coordination feature in the foreign direct investment problem, the host country may want to design an information disclosure mechanism that can increase the correlation between the foreign investors' information, so that it can encourage coordination and boost foreign direct investments.

Appendix A Omitted Proofs

This section includes proofs of propositions and lemmas that are stated in the text but not proved.

Proof of Proposition 1:

Since this proposition is about the static game, I do not use any time index in order to simplify the notation. I first show that if a Bayesian Nash equilibrium exists, it is in the cutoff strategies. Because signals are conditionally independent, in equation (3) fix any s_j , $\Pr(s_j = 1|m)$ is a constant number less than or equal to 1. Therefore, for any fixed s_j , player i 's interim payoff $E_{x_j} u_i(1, x_i, s_j)$ is strictly decreasing in x_i . Note also that since $\lim_{x_i \rightarrow -\infty} E_{x_j} u_i(1, x_i, s_j) = 1 - c > 0$ (dominant region of attacking) and $\lim_{x_i \rightarrow +\infty} E_{x_j} u_i(1, x_i, s_j) = -c < 0$ (dominant region of not attacking), the best response to any s_j is a cutoff strategy with threshold point $\hat{x}_i \in \mathbb{R}$. Therefore, if a Bayesian Nash equilibrium exists, it is in the cutoff strategies. Hence, I represent an equilibrium profile by (\hat{x}_1, \hat{x}_2) .

Second, I show that if a Bayesian Nash equilibrium exists, it is symmetric; that is, $\hat{x}_1 = \hat{x}_2$. Suppose, then, that there is an equilibrium (\hat{x}_1, \hat{x}_2) with $\hat{x}_1 > \hat{x}_2$. Because players are ex-ante homogeneous, there exists another equilibrium $(\hat{\hat{x}}_1, \hat{\hat{x}}_2) = (\hat{x}_2, \hat{x}_1)$. Because $E_{x_j} u_i(1, x_i, s_j)$ is strictly supermodular and $E_{x_j} u_i(1, \hat{x}_i, \hat{x}_j) = 0$, \hat{x}_i is strictly increasing in \hat{x}_j . Thus $\hat{\hat{x}}_2 = \hat{x}_1 > \hat{x}_2$ implies that $\hat{x}_2 = \hat{\hat{x}}_1 > \hat{x}_1$, a contradiction.

Now consider any symmetric cutoff strategy profile (x, x) . Let's fix any public belief μ ($\mu(\theta) > 0$ for all $\theta \in \Theta$). Then, the interim payoff from attacking given the signal x and the opponent's cutoff strategy with threshold point x can be written as:

$$G(x, \mu) = \underbrace{\frac{\mu(w)\phi(\sqrt{\beta}(x-w))}{\sum_{\theta' \in \Theta} \mu(\theta')\phi(\sqrt{\beta}(x-\theta'))}}_{\text{posterior belief about } \theta=w} + \underbrace{\frac{\mu(m)\phi(\sqrt{\beta}(x-m))}{\sum_{\theta' \in \Theta} \mu(\theta')\phi(\sqrt{\beta}(x-\theta'))}}_{\text{posterior belief about } \theta=m} \underbrace{\Phi(\sqrt{\beta}(x-m))}_{\text{probability } j \text{ attacks}} - c.$$

Because $G(x, \mu)$ is continuous in x , the dominant region of attacking and the dominant region of not attacking imply that there exists $x^* \in \mathbb{R}$ such that (x^*, x^*) is an equilibrium.

Finally, I claim that for any fixed β , there exists μ with $\mu(\theta) > 0$ for all $\theta \in \Theta$, such that multiple equilibria exist in this static regime change game. To show this, I only need to show that there exists μ such that there is more than one solution to $G(x, \mu) = 0$. Because $c \in (0, 1)$, there is $x' < x''$ such that $\Phi(\sqrt{\beta}(x' - m)) < c < \Phi(\sqrt{\beta}(x'' - m))$. Then $\lim_{\mu(m) \rightarrow 1} G(x', \mu) = \Phi(\sqrt{\beta}(x' - m)) - c < 0$ and $\lim_{\mu(m) \rightarrow 1} G(x'', \mu) = \Phi(\sqrt{\beta}(x'' - m)) - c > 0$. Therefore, when $\mu(m)$ is sufficiently large, the dominant region of attacking, the dominant region of not attacking, and the continuity of $G(x, \mu)$ in x imply that there are three solutions to $G(x, \mu) = 0$: one in $(-\infty, x')$, one in (x', x'') , and one in $(x'', +\infty)$.

Proof of Proposition 2:

This proposition is also about the static game and, as before, I exclude the time index. In this proof, the public belief μ satisfies $\mu(w) = 0$, $\mu(m) > 0$, and $\mu(s) > 0$. First, when $\mu(w) = 0$, for a fixed s_j , $E_{x_j} u_i(1, x_i, s_j)$ is strictly decreasing in x_i , and the regime change game is supermodular. Similar to the proof of Proposition 1, if an equilibrium with attacks exists, it is symmetric and in the cutoff strategies. Denote a symmetric cutoff-strategy profile by (x, x) ; then (x^*, x^*) is a nontrivial equilibrium of the regime change game if and only if $G(x^*, \mu) = 0$. Therefore, the conditions for the existence of a nontrivial equilibrium are equivalent to those for the existence of a solution to $G(x, \mu) = 0$. Note that $G(x, \mu)$ can be equivalently written as $G(x, \mu) = \rho(m|x)g(x, \mu)$, where

$$g(x, \mu) = [\Phi(\sqrt{\beta}(x - m)) - c] - c\left(\frac{1}{\mu(m)} - 1\right) \exp\left[\frac{\beta}{2}(s - m)(2x - s - m)\right].$$

For any $x \in \mathbb{R}$, $\rho(m|x) > 0$; therefore x^* is a solution to $G(x, \mu) = 0$ if and only if it is a solution to $g(x, \mu) = 0$. The rest of this proof relies on the following sequence of lemmas.

Lemma 2 *There exist $\bar{\mu}(m), \underline{\mu}(m) \in (0, 1)$ with $\bar{\mu}(m) > \underline{\mu}(m)$, such that for all $\mu(m) \in (0, \underline{\mu}(m)]$, there is no solution to $g(x, \mu) = 0$, and for all $\mu(m) \in [\bar{\mu}(m), 1)$, there is $x^* \in \mathbb{R}$ such that $g(x^*, \mu) = 0$.*

Proof: First consider the case in which $\mu(m)$ is close to 1. Since $c < 1$, there is an $x' \in \mathbb{R}$ such that $\Phi(\sqrt{\beta}(x' - m)) > c$. Because there exists a $\mu(m)$ close to 1 such that $c\left(\frac{1}{\mu(m)} - 1\right) \exp\left[\frac{\beta}{2}(s - m)(2x' - s - m)\right]$ is very close to zero, $g(x', \mu) > 0$ when $\mu(m)$ is large enough. Note that for all $\mu(m) \in (0, 1)$, $g\left(m + \frac{\Phi^{-1}(c)}{\sqrt{\beta}}, \mu\right) < 0$ and $\lim_{x \rightarrow +\infty} g(x, \mu) < 0$, and so by the continuity of $g(x, \mu)$ in x , there exist $\hat{x} \in \left(m + \frac{\Phi^{-1}(c)}{\sqrt{\beta}}, x'\right)$ and $\hat{\hat{x}} \in (x', +\infty)$ such that $g(\hat{x}, \mu) = 0$ and $g(\hat{\hat{x}}, \mu) = 0$. Therefore, there exists $\bar{\mu}(m) \in (0, 1)$ such that solutions to $g(x, \mu) = 0$ exist for all $\mu(m) \in [\bar{\mu}(m), 1)$.

Now consider $\mu(m)$ close to 0. The last term of $g(x, \mu)$ is very negative for any x larger than $m + \frac{\Phi^{-1}(c)}{\sqrt{\beta}}$, so $g(x, \mu) < 0$ for all $x > m + \frac{\Phi^{-1}(c)}{\sqrt{\beta}}$. Combining this with the fact that $g(x, \mu) < 0$ for all $x \leq m + \frac{\Phi^{-1}(c)}{\sqrt{\beta}}$, we get $\underline{\mu}(m) \in (0, 1)$ such that for all $\mu(m) \in (0, \underline{\mu}(m)]$, $g(x, \mu) < 0, \forall x \in \mathbb{R}$.

Finally, because $\bar{\mu}(m)$ can be picked as a number very close to 1 and $\underline{\mu}(m)$ can be picked as a number very close to 0, $\bar{\mu}(m) > \underline{\mu}(m)$.

Lemma 3 *There exists $\tilde{\mu} \in (\bar{\mu}(m), \underline{\mu}(m))$, such that for all $\mu(m) \in (0, \tilde{\mu})$, there is no solution to $g(x, \mu) = 0$, and for all $\mu(m) \in (\tilde{\mu}, 1)$, there are two solutions to $g(x, \mu) = 0$. Therefore, claims (1) and (2) in Proposition 2 are true.*

Proof: Suppose $1 > \mu'(m) > \mu''(m) > 0$ and $\exists x'' \in (m + \Phi^{-1}(c)/\sqrt{\beta}, +\infty)$ such that $g(x'', \mu'') = 0$ (because all $x \leq m + \Phi^{-1}(c)/\sqrt{\beta}$ cannot be a solution to $g(x, \mu'') = 0$). Since $g(x, \mu)$ is strictly increasing in $\mu(m)$ for any fixed $x \in \mathbb{R}$, $g(x'', \mu') > g(x'', \mu'') = 0$. Then by the continuity of $g(x, \mu')$ and $\lim_{x \rightarrow +\infty} g(x, \mu') < 0$, there exists $x' \in (x'', +\infty)$ such that $g(x', \mu') = 0$. Similarly, if $1 > \mu'(m) > \mu''(m) > 0$ and $g(x; \mu') < 0$ for all $x \in \mathbb{R}$, then $g(x, \mu'') < 0$ for all $x \in \mathbb{R}$. Define $\tilde{\mu} = \inf\{\mu(m) \in (0, 1) : \exists x \in \mathbb{R} \text{ such that } g(x, \mu) = 0\} = \sup\{\mu(m) \in (0, 1) : g(x, \mu) < 0 \ \forall x \in \mathbb{R}\}$ (since for a given μ , $g(x, \mu)$ either has a solution or does not have a solution). Obviously, $\tilde{\mu} \in (\bar{\mu}(m), \underline{\mu}(m))$.

For all $\mu(m) \in (\tilde{\mu}, 1)$, note that $\frac{\partial^2 g}{\partial x^2} < 0$ for all $x \geq m$ and $g(x, \mu)$ has a single peak in $(m, +\infty)$. Therefore, when $\mu(m) \in (\tilde{\mu}, 1)$, there are two solutions to $g(x, \mu) = 0$.

Lemma 4 *There exists a unique $\tilde{x} \in (m, +\infty)$ such that $g(\tilde{x}, \tilde{\mu}) = 0$. Therefore, claim (3) in Proposition 2 is true.*

Proof: Suppose $\forall x \in \mathbb{R}$, $g(x, \tilde{\mu}) < 0$. Recall that because $\mu(m) < 1$, for any $x \in (\bar{x}(\tilde{\mu}), +\infty)$, $g(x, \tilde{\mu}) < 0$ (dominant region of not attacking), where $\bar{x}(\mu(m)) = \inf\{x \in \mathbb{R} : E_{x_j} u_i(1, x_i, s_j) < 0 \text{ for all } s_j\}$. Since $E_{x_j} u_i(1, x_i, s_j)$ is increasing in $\mu(m)$, $\bar{x}(\mu(m))$ is an increasing function in $\mu(m)$. As a result, $g(x, \tilde{\mu}) < 0$ for all $x > \bar{x}(\bar{\mu}(m))$ because $\tilde{\mu} < \bar{\mu}(m)$. Since $c \in (0, 1)$, there is an x' such that for any μ , $g(x, \mu) < 0$ for all $x < x'$. Now consider the compact set $[m, \bar{x}(\bar{\mu}(m))]$. By the continuity of $g(x, \mu)$ in x , $\exists \hat{x} \in [m, \bar{x}(\bar{\mu}(m))]$ such that $g(x, \tilde{\mu}) \leq g(\hat{x}; \tilde{\mu}) < 0$. Pick a sequence $\{\mu^k(m)\}$ such that $\mu^k(m) \in (\tilde{\mu}, \bar{\mu}(m))$, $\mu^k(m) > \mu^{k+1}(m)$ and $\mu^k(m) \rightarrow \tilde{\mu}$. Since $g(x, \mu)$ is continuous in $\mu(m)$ for any $x \in [m, \bar{x}(\bar{\mu}(m))]$, $\lim_{k \rightarrow +\infty} g(x, \mu^k) = g(x, \tilde{\mu})$. Defining $M^k = \sup_{x \in [m, \bar{x}(\bar{\mu}(m))]} |g(x, \mu^k) - g(x, \tilde{\mu})|$, it can be calculated that

$$\begin{aligned} M^k &= \sup_{x \in [m, \bar{x}(\bar{\mu}(m))]} \left| \frac{1}{\mu^k(m)} - \frac{1}{\tilde{\mu}} \right| c \exp\left[\frac{\beta}{2}(s-m)(2x-s-m)\right] \\ &= \left| \frac{1}{\mu^k(m)} - \frac{1}{\tilde{\mu}} \right| c \exp\left[\frac{\beta}{2}(s-m)(2\bar{x}(\bar{\mu}(m))-s-m)\right]. \end{aligned}$$

Therefore, $\forall \epsilon > 0$, $\exists K$ such that for all $k > K$, $\left| \frac{1}{\mu^k(m)} - \frac{1}{\tilde{\mu}} \right| < \frac{\epsilon}{c \exp[\frac{\beta}{2}(s-m)(2\bar{x}(\bar{\mu}(m))-s-m)]}$, which implies that $M^k < \epsilon$. Therefore, $g(x, \mu^k)$ converges to $g(x, \tilde{\mu})$ uniformly, and there exists K' such that for all $k > K'$, $g(x, \mu^k) - g(x, \tilde{\mu}) < \frac{|g(\hat{x}, \tilde{\mu})|}{2}$, thus $g(x, \mu^k) < -\frac{|g(\hat{x}, \tilde{\mu})|}{2} < 0$ for all $x \in [m, \bar{x}(\bar{\mu}(m))]$. Note that for any $x < m$ and $x > \bar{x}(\bar{\mu}(m))$, $g(x, \mu^k) < 0$, and so, for all $x \in \mathbb{R}$, $g(x, \mu^k) < 0$. But by the definition of $\tilde{\mu}$, there must be some $x' \in \mathbb{R}$ such that $g(x', \mu^k) = 0$. Therefore, when $\mu(m) = \tilde{\mu}$, there exists \tilde{x} such that $g(\tilde{x}, \tilde{\mu}) = 0$.

Now suppose that $x' \neq \tilde{x}$ and $g(x', \tilde{\mu}) = 0$. Because $\frac{\partial^2 h}{\partial x^2} < 0$ for all $x \geq m$, there must be x'' between x' and \tilde{x} such that $g(x'', \tilde{\mu}) > 0$. Since $g(x'', \mu)$ is continuous in $\mu(m)$, for any fixed $\epsilon \in (0, \frac{g(x'', \tilde{\mu})}{2})$, there exists $\gamma > 0$ such that for all $\mu'(m) \in (\tilde{\mu} - \gamma, \tilde{\mu})$, $g(x'', \mu) >$

$g(x'', \tilde{\mu}) - \epsilon > 0$. Therefore, there exists $x''' \in (m, x'')$ such that $g(x'''; \mu) = 0$. This contradicts the definition of $\tilde{\mu}$. Therefore, there exists a unique $\tilde{x} \in \mathbb{R}$, such that $g(\tilde{x}, \tilde{\mu}) = 0$.

Q.E.D.

Proof of Proposition 3:

Given μ_1 and a fixed strategy profile, conditional on $R_{t-1} = 0$, period- t players' public beliefs μ_t after \hat{h}^t can be calculated as

$$\mu_t(\theta) = \frac{\mu_1(\theta) \prod_{\tau=1}^{t-1} [\Phi(\sqrt{\beta}(\theta - x_\tau^*))]^2}{\sum_{\theta' \in \Theta} \mu_1(\theta') \prod_{\tau=1}^{t-1} [\Phi(\sqrt{\beta}(\theta' - x_\tau^*))]^2}, \quad \forall \theta \in \Theta, \quad (10)$$

where $[\Phi(\sqrt{\beta}(\theta - x_\tau^*))]^2$ is the conditional (on θ) probability that both players in period τ observe signals landing above the cutoff point x_τ^* . Because $\mu_1(\theta) > 0$ for all $\theta \in \Theta$, $x_1^* \in \mathbb{R}$ by Proposition 1, which in turn implies that $\mu_2(\theta) > 0$ for all $\theta \in \Theta$. Then, by induction, along history \hat{h}^t , $\mu_t(\theta) > 0$ for all $\theta \in \Theta$, and $x_t^* \in \mathbb{R}$. Because of the monotone likelihood ratio property of private signals, in any particular period t attacks occur with highest probability in the weak state. Thus, the no-attack outcome in period t lowers period- $t+1$ players' beliefs about $\theta = w$. As a result, along \hat{h}^t , $\{\mu_t(w)\}_t$ is a bounded and strictly decreasing sequence that converges to $\mu_\infty(w) \geq 0$. If $\mu_\infty(w) > 0$, the dominant region of attacking has a positive measure in the limit. Therefore, the probability of attacking is bounded away from 0, then $\mu_t(w) \rightarrow 0$ by the belief updating equation (10). Hence, $\mu_t(w)$ must converge to 0. For the probability of attacking along \hat{h}^t , suppose that there is an infinite subsequence of attacking probabilities that are bounded away from 0; that is, there exists $\epsilon > 0$ such that all terms in this subsequence are greater than ϵ . Therefore, there exists T such that after the history \hat{h}^T , players' public beliefs about $\theta < s$ are so low that the cutoff points after period T should be arbitrarily negative. But, in the subsequence, the fact that the probability of attacking is bounded away from 0 implies that the corresponding cutoff points are bounded from below. This contradiction implies that the probability of attacking converges to 0. Therefore, in any fixed equilibrium, along \hat{h}^t , $\mu_t(w) \rightarrow 0$ and $\Pr(b_t > 0 | \hat{h}^t) \rightarrow 0$.

Now, let us consider the regime change probability at the state $\theta = w$. In a fixed equilibrium, the strategy profile and the prior belief induce a probability measure \mathbb{P} on the outcome space $\Theta \times \{0, 1, 2\}^\infty$. Suppose that \mathbb{P}_w and $\hat{\mathbb{P}}$ are probability measures induced on $\Theta \times \{0, 1, 2\}^\infty$ by \mathbb{P} , conditioned on the state w and the set $\{m, s\}$, respectively. Hence, $\mathbb{P} = \mu_0(w)\mathbb{P}_w + (1 - \mu_0(w))\hat{\mathbb{P}}$.

The sequence $\{\mu_t(w)\}_t$ is a bounded martingale adapted to the filtration \mathcal{F}^t , which is generated by the history H^t under the measure \mathbb{P} . Therefore, $\{\mu_t(w)\}_t$ converges \mathbb{P} – almost surely to $\mu_\infty(w)$. Since \mathbb{P}_w is absolutely continuous with respect to \mathbb{P} , $\mu_t(w) \rightarrow \mu_\infty(w)$, \mathbb{P}_w – almost surely.

Now, suppose that there is a set $A \subset \{0, 1, 2\}^\infty$ such that $\mu_\infty(w)[a] = 0, \forall a \in A$, and $\mathbb{P}_w(A) > 0$. Bayes' rule implies that the odds ratio $\{(1 - \mu_t(w))/\mu_t(w)\}_t$ is a \mathbb{P}_w – martingale, and so $\mathbb{E}[\frac{1 - \mu_t(w)}{\mu_t(w)}] = \frac{1 - \mu_0(w)}{\mu_0(w)}$ for all t . However, $\mathbb{E}[\frac{1 - \mu_t(w)}{\mu_t(w)}] = \mathbb{E}[\frac{1 - \mu_t(w)}{\mu_t(w)} \chi(A)] + \mathbb{E}[\frac{1 - \mu_t(w)}{\mu_t(w)} (1 - \chi(A))]$, where χ is the indicator function. Obviously, the second term is nonnegative, while the first term is greater than $\frac{1 - \mu_0(w)}{\mu_0(w)}$ for very large t , since $\mu_\infty(w)(a) = 0, \forall a \in A$, which leads to a contradiction. Therefore, $\mu_\infty(w) > 0, \mathbb{P}_w$ – almost surely.

Since, along \hat{h}^t , $\mu_t(w) \rightarrow \mu_\infty(w) = 0$, \hat{h}^t is a 0 measure event under \mathbb{P}_w . As a result, the regime changes almost surely when the status quo is weak (because \hat{h}^∞ is the only outcome in which the weak status quo never changes).

Q.E.D.

Proof of Lemma 1:

As t increases, the precision of players' signals increases. Because players are still short-lived, how $\tilde{\mu}_t$ changes over time is equivalent to how $\tilde{\mu}$ changes as β increases in the static game.

Recall that $\tilde{\mu}$ is the belief about $\theta = m$, where there is a unique $\tilde{x} \in \mathbb{R}$ such that $G(\tilde{x}, \tilde{\mu}) = 0$ (where $\tilde{\mu}(w) = 0$). Since $G(x, \tilde{\mu}) < 0$ for all $x \neq \tilde{x}$, $G'(\tilde{x}, \tilde{\mu}) = 0$. As in Proposition 2, instead of studying $G(x, \tilde{\mu})$ directly, it is easier to study the function $g(x, \tilde{\mu}) = [\Phi(\sqrt{\beta}(x - m)) - c] - c(\frac{1}{\tilde{\mu}} - 1) \exp(\frac{\beta}{2}(s - m)(2x - s - m))$. Since \tilde{x} is also the unique solution to $g(x, \tilde{\mu}) = 0$, $g'(\tilde{x}, \tilde{\mu}) = 0$. That is,

$$\begin{aligned} [\Phi(\sqrt{\beta}(\tilde{x} - m)) - c] - c(\frac{1}{\tilde{\mu}} - 1) \exp(\frac{\beta}{2}(s - m)(2\tilde{x} - s - m)) &= 0 \\ \phi(\sqrt{\beta}(\tilde{x} - m)) - \sqrt{\beta}(s - m)c(\frac{1}{\tilde{\mu}} - 1) \exp(\frac{\beta}{2}(s - m)(2\tilde{x} - s - m)) &= 0. \end{aligned}$$

A comparative static analysis shows that, for large β , $\tilde{\mu}$ is decreasing in β .

A necessary condition for the above system of equations is $\Phi(\sqrt{\beta}(\tilde{x} - m)) - c = \frac{\phi(\sqrt{\beta}(\tilde{x} - m))}{\sqrt{\beta}(s - m)}$. The right-hand side obviously goes to 0, as β goes to $+\infty$. Therefore, as β goes to $+\infty$, $\Phi(\sqrt{\beta}(\tilde{x} - m))$ goes to c , which implies that $\sqrt{\beta}(\tilde{x} - m)$ goes to $\Phi^{-1}(c)$. Hence, as $\beta \rightarrow +\infty$, $\exp(\frac{\beta}{2}(s - m)(2\tilde{x} - s - m))$ goes to $\exp(-\frac{(s - m)^2}{2}\beta + \Phi^{-1}(c)(s - m)\sqrt{\beta})$. Suppose that $\tilde{\mu}$ is bounded away from 0 as β goes to $+\infty$; then $(\frac{1}{\tilde{\mu}} - 1) \exp(-\frac{(s - m)^2}{2}\beta + \Phi^{-1}(c)(s - m)\sqrt{\beta})$ and $\sqrt{\beta}(\frac{1}{\tilde{\mu}} - 1) \exp(-\frac{(s - m)^2}{2}\beta + \Phi^{-1}(c)(s - m)\sqrt{\beta})$ both go to 0. Hence, $g'(\tilde{x}, \tilde{\mu}) > 0$, which leads to a contradiction. Hence, as $\beta \rightarrow +\infty$, $\tilde{\mu} \rightarrow 0$.

Q.E.D.

Proof of Proposition 5:

From part 2 of Proposition 2, after $\theta = w$ is ruled out, when $\mu_t(m) > \tilde{\mu}_t$, period- t players can employ a cutoff strategy that they attack if and only if their signals are below $x_t^* \geq m + \Phi^{-1}(c)/\sqrt{t\beta}$. Then, as t goes to $+\infty$, x_t^* goes to m . That is, for any $\epsilon > 0$, there is a $T \in \mathbb{N}$ such that for all $t \geq T$, $x_t^* \geq m - \epsilon$. Therefore, x_t^* is bounded from below.

Now we can construct an equilibrium in which the medium status quo changes almost surely. When $\theta = w$ has not been ruled out, period- t players employ the strategy specified in Proposition 1. If θ_w is ruled out, period- t players refrain from attacking if $\mu_t(m) < \tilde{\mu}_t$, and they employ the strategy with a finite threshold as in part 2 and part 3 of Proposition 2 if $\mu_t(m) \geq \tilde{\mu}_t$. Finally, if both $\theta = w$ and $\theta = m$ are ruled out, no subsequent player attacks.

I am now in the position to show that the medium status quo changes almost surely in such an equilibrium. First, as in Proposition 3, the outcome of no attack in any period cannot occur with positive probability. Therefore, $\theta = w$ will be ruled out almost surely. Second, after $\theta = w$ is ruled out, if $\mu_t(m) < \tilde{\mu}_t$ in some period t , players do not attack, no matter what their signals are. Thus, the no-attack outcome will not convey new information to subsequent players, implying that the public belief about $\theta = m$ stays at $\mu_t(m)$ until some period $\tau > t$ with $\mu_\tau(m) \geq \tilde{\mu}_\tau(m)$. By Proposition 1, such a τ always exists for any t . Therefore, period- τ players will attack if and only if their signals are above x_τ^* , which is bounded from below. This implies that in period τ , the medium status quo changes with a probability that is bounded away from zero. Since there are infinitely many periods in which the medium status quo changes with a probability bounded away from zero, the medium status quo changes almost surely.

Q.E.D.

References

- [1] Angeletos, G., Hellwig C., Pavan A.: Dynamic Global Games of Regime Change: Learning, Multiplicity, and the Timing of Attacks. *Econometrica* **75(3)**, 711–756 (2007).
- [2] Angeletos, G., Pavan A.: Efficient Use of Information and Social Value of Information. *Econometrica* **75(4)**, 1103–1142 (2007).
- [3] Banerjee, A.V.: A simple Model of Herd Behavior. *Q. J. Econ.* **107(3)**, 797–817 (1992).
- [4] Bikhchandani, S. Hirshleifer D., Welch I.: A Theory of Fads, Fashion, Custom, and Culture Change as Informational Cascades. *J. Polit. Econ.* **100(5)** 992–1026 (1992).
- [5] Bikhchandani, S. Hirshleifer D., Welch I.: Learning from the Behavior of Others: Conformity, Fads, and Informational Cascades. *J. Econ. Perspect.* **12(3)**, 151–170 (1998).
- [6] Carlsson, H., Van Damme, E.: Global Games and Equilibrium Selection. *Econometrica* **61(5)**, 989–1018 (1993).
- [7] Chamley, C.: *Rational Herds*. Cambridge University Press, Cambridge (2004).
- [8] Chen, H., Suen W.: Falling Dominoes: A Theory of Rare Events and Crisis Contagion. *AEJ: Micro* **8(1)**, 228–255 (2016).
- [9] Choi, J.P.: Herd Behavior, the “Penguin Effect,” and the Suppression of Informational Diffusion: An Analysis of Informational Externalities and Payoff Interdependency. *Rand. J. Econ.* **28(3)**, 407–425 (1997).
- [10] Cripps, M., Ely J., Mailath G., Samuelson L.: Common Learning. *Econometrica* **76(4)**, 909–933 (2008).
- [11] Dasgupta, A.: Social Learning and Payoff Complementarities. <http://personal.lse.ac.uk/DASGUPT2/research.html> (2000).
- [12] Dasgupta, A.: Coordination and Delay in Global Games. *J. Econ. Theory.* **134(1)**, 195–225 (2007).
- [13] Dasgupta, A., Steiner J., Stewart C.: Dynamic Coordination with Individual Learning. *Game. Econ. Behav.* **74(1)**, 83–101 (2012).
- [14] Edmond, C.: Information Manipulation, Coordination, and Regime Change. *Rev. Econ. Stud.* **80(4)**, 1422–1458 (2013).

- [15] Giannitsarou, C., Toxvaerd, F.: Recursive Global Games. CEPR Discussion Paper No. DP6470 (2007).
- [16] Goldstein, I., Pauzner, A.: Demand Deposit Contracts and the Probability of Bank Runs. *J. Financ.*, **60(3)**, 1293–1328 (2005).
- [17] Heidhues, P., Melissas, N.: Equilibria in a Dynamic Global Game: The Role of Cohort Effects. *Econ. Theory* **28**, 531–57 (2006).
- [18] Hirshleifer, J.: The Private and Social Value of Information and the Reward to Inventive Activity. *Am. Econ. Rev.* **61(4)**, 561–574 (1971).
- [19] Lee, I.H.: On the Convergence of Informational Cascades. *J. Econ. Theory.* **61(2)**, 396–411 (1993).
- [20] Morris, S., Shin, H.S.: Unique Equilibrium in a Model of Self-Fulfilling Currency Attacks. *Am. Econ. Rev.* **88(3)**, 587–597 (1998).
- [21] Morris, S., Shin, H.S.: Social Value of Public Information. *Am. Econ. Rev.* **92(5)**, 1521–1534 (2002).
- [22] Morris, S., Shin, H.S.: Global Games: Theory and Applications. In: Dewatripont, M., Hansen, L., Turnovsky S. (eds.) *Advances in Economics and Econometrics: Theory and Applications*, Eighth World Congress, Vol. 1, pp. 57–114 (2003).
- [23] Morris, S., Shin, H.S.: Coordination Risk and the Price of Debt. *Eur. Econ. Rev.* **48(1)**, 133–153 (2004).
- [24] Smith, L., Sørensen P.: Pathological Outcomes of Observational Learning. *Econometrica* **68(2)**, 371–398 (2000).
- [25] Tarashev, N.: Speculative Attacks and the Information Role of the Interest Rate. *J. Eur. Econ. Assoc.* **5(1)**, 1–36 (2007).

Appendix B Online Appendix (Not for Publication)

Appendix B.1 Belief Convergence in Different Equilibria

After the weak status quo is ruled out, Proposition 4 implies that, in all equilibria, along the outcome \bar{h}^∞ , public beliefs about the medium status quo converge to $\mu_\infty(m) \in (0, 1)$. This convergence exists because players stop learning from the public history after some period. Now, let's consider an outcome along which players never stop learning from the public history after the weak status quo is ruled out; that is, public beliefs about the medium state change over time. Will $\mu_t(m)$ necessarily converge to 1, if the status quo is medium? This subsection shows that the answer to this belief convergence question differs in different equilibria.

Let's first define two classes of equilibria that have nice properties and are easily analyzed. Specifically,

Definition 1 $E^{max} = \{(x_t^*, x_t^*)_{t=1, \dots}, (\mu_t^*)_{t=1, \dots}\}$ is the equilibrium in which, given μ_t^* , for any (x'_t, x'_t) forming a static equilibrium in period t with prior belief μ_t^* , $x'_t \leq x_t^*$.

Definition 2 $E^{min} = \{(x_t^*, x_t^*)_{t=1, \dots}, (\mu_t^*)_{t=1, \dots}\}$ is the equilibrium in which, given μ_t^* , for any (x'_t, x'_t) forming an equilibrium in period t with prior belief μ_t^* and $x'_t \in \mathbb{R}$ (if it exists), $x_t^* \in \mathbb{R}$ and $x_t^* \leq x'_t$; if there is no such x'_t , then $x_t^* = -\infty$.

Therefore, E^{max} is the equilibrium in which players choose the most aggressive strategy in their own period, and E^{min} is the equilibrium in which players choose the lowest possible cooperation strategy in their own period. Figure 1 shows how the cutoff points in E^{max} and E^{min} are determined, when $\mu(w) = 0$ and $\mu(m) > \tilde{\mu}$.




Figure 1: Function $G(x; \mu)$ with $\mu_3(m) > \mu_2(m) > \tilde{\mu}(m) > \mu_1(m)$.

Fix an equilibrium. Suppose the weak status quo has been ruled out by period t , and there is an attack by one player in period t . Then, given μ_t , how period- $t + 1$ players form their public beliefs about the medium status quo depends on the threshold point of period- t players' strategies.

Lemma 1 *Given μ_t , suppose the attack by one player in period t fails. Then $|\mu_{t+1}(m) - \mu_t(m)|$ is strictly increasing in $|x_t^* - \frac{m+s}{2}|$. Moreover, $\mu_{t+1}(m) \geq \mu_t(m)$ if and only if $x_t^* \leq \frac{m+s}{2}$.*

Proof 1 (Proof of Lemma 1) *Suppose the weak status quo has been ruled out by period t . Given μ_t , $b_t = 1$, and x_t^* , Bayes' rule implies*

$$\begin{aligned} \mu_{t+1}(m) &= \frac{\mu_t(m)\Phi[\sqrt{\beta}(x_t^* - m)]\Phi[\sqrt{\beta}(m - x_t^*)]}{\mu_t(m)\Phi[\sqrt{\beta}(x_t^* - m)]\Phi[\sqrt{\beta}(m - x_t^*)] + (1 - \mu_t(m))\Phi[\sqrt{\beta}(x_t^* - s)]\Phi[\sqrt{\beta}(s - x_t^*)]} \\ &= \frac{\mu_t(m)}{\mu_t(m) + (1 - \mu_t(m))\frac{\Phi[\sqrt{\beta}(x_t^* - s)]\Phi[\sqrt{\beta}(s - x_t^*)]}{\Phi[\sqrt{\beta}(x_t^* - m)]\Phi[\sqrt{\beta}(m - x_t^*)]}}. \end{aligned}$$

Obviously, if $x_t^* = \frac{m+s}{2}$, $\mu_{t+1}(m) = \mu_t(m)$.

In the equilibrium, when $\mu_t(w) = 0$, $x_t^* > m$. Now, consider the case x' and x , such that $m < x' < x < \frac{m+s}{2}$. (The case $\frac{m+s}{2} > x > x'$ is similar.) Then

$$\Phi[\sqrt{\beta}(x' - s)] < \Phi[\sqrt{\beta}(x - s)] < \frac{1}{2} < \Phi[\sqrt{\beta}(x' - m)] < \Phi[\sqrt{\beta}(x - m)].$$

Because the function $f(y) = y(1-y)$ is strictly concave and has the maximum value at $y = \frac{1}{2}$, we get

$$\frac{\Phi[\sqrt{\beta}(x - s)]\Phi[\sqrt{\beta}(s - x)]}{\Phi[\sqrt{\beta}(x - m)]\Phi[\sqrt{\beta}(m - x)]} > \frac{\Phi[\sqrt{\beta}(x' - s)]\Phi[\sqrt{\beta}(s - x')]}{\Phi[\sqrt{\beta}(x' - m)]\Phi[\sqrt{\beta}(m - x')]},$$

which implies that given μ_t and $b_t = 1$, $\mu_{t+1}(m)$ is strictly decreasing in x_t^* . Since $\Phi[\sqrt{\beta}(x - s)] < \frac{1}{2} < \Phi[\sqrt{\beta}(x - m)]$ for $x \in (m, \frac{m+s}{2})$, we get

$$\frac{\Phi[\sqrt{\beta}(x - s)]\Phi[\sqrt{\beta}(s - x)]}{\Phi[\sqrt{\beta}(x - m)]\Phi[\sqrt{\beta}(m - x)]} < 1.$$

Therefore, $\mu_{t+1}(m) > \mu_t(m)$ if $x_t^* < \frac{m+s}{2}$.

Because only the medium state and the strong state are in the support of period- $t + 1$ players' public beliefs, if $x_t^* = \frac{m+s}{2}$, then $b_t = 1$ is neutral in terms of belief updating. Hence, the further x_t^* is away from $\frac{m+s}{2}$, the more informative the attack by one player in period t is.

To simplify the analysis, I assume that $\tilde{x} < \frac{m+s}{2}$. (Recall that \tilde{x} is the solution to the equation $G(x, \tilde{\mu}) = 0$, where $\tilde{\mu}(w) = 0$ and $\max_x G(x, \tilde{\mu}) = 0$.) Also, denote by μ' the public belief that $\mu'(w) = 0$ and $\frac{m+s}{2}$ is the largest solution to the equation $G(x, \mu') = 0$. I focus on the evolutions of the public belief along the history $\tilde{h}^\infty \equiv (1, 1, \dots)$, in which exactly one

player chooses to attack in every period. Note that though \tilde{h}^∞ occurs with 0 probability, any t -period long history $\tilde{h}^t = (1, 1, \dots, 1)$ occurs with positive probability.

Proposition 1 *When $\theta = w$ is ruled out, the public belief about $\theta = m$ evolves differently in different equilibria along the same history. In particular, suppose that $\tilde{x} < \frac{m+s}{2}$; then, along the outcome \tilde{h}^∞ , we have*

1. *in E^{max} , $\mu_t(m)$ monotonically converges to $\mu'(m) \in (0, 1)$. Thus, conditional on $\theta = m$, the probability of attacking converges to $\Phi[\frac{\sqrt{\beta}}{2}(s - m)]$;*
2. *in E^{min} , $\mu_t(m)$ converges to 1. The probability of attacking is strictly decreasing over time; and conditional on $\theta = m$, the probability of attacking converges to c , the cost of attacking.*

Proof 2 (Proof of Proposition 1) Part 1: *When the weak status quo is ruled out, $\mu_{t+1}(m)$ is a function of $\mu_t(m)$. Then Lemma 1 implies that this function has a unique fixed point in $(\tilde{\mu}, 1)$. Therefore, I only need to show that this fixed point is stable, which is equivalent to showing that the slope $0 < d\mu_{t+1}(m)/d\mu_t(m) < 1$. Since*

$$\mu_{t+1}(m) = \frac{\mu_t(m)\Phi[\sqrt{\beta}(x_t^* - m)]\Phi[\sqrt{\beta}(m - x_t^*)]}{\mu_t(m)\Phi[\sqrt{\beta}(x_t^* - m)]\Phi[\sqrt{\beta}(m - x_t^*)] + (1 - \mu_t(m))\Phi[\sqrt{\beta}(x_t^* - s)]\Phi[\sqrt{\beta}(s - x_t^*)]},$$

it follows that $\frac{\partial \mu_{t+1}(m)}{\partial \mu_t(m)} \Big|_{\tilde{\mu}, \tilde{x}} = 1$, and $-1 < \frac{\partial \mu_{t+1}(m)}{\partial x_t^} \Big|_{\tilde{\mu}, \tilde{x}} \frac{\partial x_t^*}{\partial \mu_t(m)} \Big|_{\tilde{\mu}, \tilde{x}} < 0$ when β is large. Therefore, $0 < d\mu_{t+1}(m)/d\mu_t(m) < 1$.*

Part 2: *Along the history \tilde{h}^t ,*

$$\mu_{t+1}(m) = \frac{\mu_1(m)}{\mu_1(m) + (1 - \mu_1(m)) \prod_{\tau=1}^t \frac{\Phi[\sqrt{\beta}(x_\tau^* - s)]\Phi[\sqrt{\beta}(s - x_\tau^*)]}{\Phi[\sqrt{\beta}(x_\tau^* - m)]\Phi[\sqrt{\beta}(m - x_\tau^*)]}}.$$

Because $\tilde{x} < \frac{m+s}{2}$, the smallest solution to the equation $g(x, \mu_t) = 0$ is strictly less than $\frac{m+s}{2}$ for all t . Combining this with the fact that $x_t^ > m$, we get $\mu_t(m) \rightarrow 1$. Since $g(x_t^*, \mu_t) = 0$ for all t , $\Phi[\sqrt{\beta}(x_t^* - m)] \rightarrow c$.*

The intuition of Proposition 1 is illustrated in the following figures. Figure 2 and Figure 3 present $\mu_{t+1}(m)$ as a function of $\mu_t(m)$ along the outcome \tilde{h}^∞ in E^{max} and E^{min} , respectively. In Figure 2, $\mu'(m)$ is the unique fixed point in $(\tilde{\mu}, 1)$. Therefore, in E^{max} , along \tilde{h}^∞ , $\mu_t(m) \rightarrow \mu'(m)$, which implies that the cutoff point converges to $\frac{m+s}{2}$ and that the probability of attacking converges to $\Phi[\frac{\sqrt{\beta}}{2}(s - m)]$. In Figure 3, the unique fixed point greater than $\tilde{\mu}$ is 1, and it is stable. This means that along \tilde{h}^∞ , $\mu_t(m) \rightarrow 1$. Hence, the probability of attacking converges to c . Thus, the equilibrium strategy profile in the limit purifies the mixed strategy Nash equilibrium of the complete-information normal-form game when $\theta = m$. (In fact, in

the complete-information normal-form game, when $\theta = m$ the mixed strategy equilibrium is the lowest possible coordination equilibrium.) The monotonicity stated in Proposition 1 can be seen in Figure 1: if $\mu_{t+1}(m) > \mu_t(m)$, $g(x, \mu_{t+1}) > g(x, \mu_t), \forall x > m$, and so $x_{t+1}^* > x_t^*$ in E^{max} and $x_{t+1}^* < x_t^*$ in E^{max} .




Figure 2: Evolution of Public Beliefs in E^{max} .




Figure 3: Evolution of Public Beliefs in E^{min} .

Proposition 1 shows that along the outcome \tilde{h}^∞ , the convergence of $\mu_t(m)$ differs in E^{max} and E^{min} . In E^{max} , the failed attack by one player in period t results in $\mu_{t+1}(m)$ strictly between $\mu_t(m)$ and $\mu'(m)$, which in turn implies that x_{t+1}^* is strictly between x_t^* and $\frac{m+s}{2}$ (see Figure 1). Hence, $b_{t+1} = 1$ is less informative than $b_t = 1$ by Lemma 1. Because the attack by one player is less informative over time, the public belief about the medium status quo converges to $\mu'(m) \in (0, 1)$. Conversely, in E^{min} the attack by one player is increasingly informative over time because the threshold points of players' strategies are decreasing over time (in Figure 1, all threshold points in E^{min} are less than $\frac{s+m}{2}$). As a result, in E^{min} , along \tilde{h}^∞ , $\mu_t(m)$ converges to 1. Put differently, along the same outcome \tilde{h}^∞ , in E^{max} , the public history provides decreasingly informative evidence over time, and so the social learning is

incomplete. But in E^{min} , the public history provides increasingly informative evidence over time, and so the social learning is complete.

Appendix B.2 Persistent Status Quo

In this section, I study the extension with persistent, but not constant, status quos (described in more detail in Section 4.2 of the paper).

I assume that the strength of the status quo, θ , changes over time according to a first-order Markov process. To make the symbols simple, I consider two cases separately. In the first case,

$$\Pr(\theta_{t+1} = w | \theta_t = k) = \alpha_k, k \in \{w, m, s\}. \quad (1)$$

Without loss of generality, I assume that $\alpha_w > \alpha_m > \alpha_s > 0$. In the second case, I assume that

$$\Pr(\theta_{t+1} = w | \theta_t = m) = \Pr(\theta_{t+1} = w | \theta_t = s) = 0, \quad (2)$$

and¹

$$\Pr(\theta_{t+1} = m | \theta_t = m) = \Pr(\theta_{t+1} = s | \theta_t = s) = \gamma \in (\frac{1}{2}, 1). \quad (3)$$

That is, once the status quo becomes medium or strong, it will no longer be weak. Yet, a medium status quo can become strong, and a strong status quo can become medium.

Let us first analyze the first case. Denote by ζ_{t-1} the posterior belief about θ at the end of period $t-1$, given the period $t-1$ public belief (μ_{t-1}), period- $t-1$ players' actions (b_{t-1}), and the no-regime-change outcome in period- $t-1$ ($R_t = 0$). Then, ζ_{t-1} , together with the transition rule in equation (1), leads to period- t players' public beliefs about $\theta_t = w$:

$$\mu_t(w) = \zeta_{t-1}(w)\alpha_w + \zeta_{t-1}(m)\alpha_m + \zeta_{t-1}(s)\alpha_s \geq \alpha_s. \quad (4)$$

Therefore, in each period t , the players' public beliefs about $\theta = s$ are at least α_s . Now, consider the most pessimistic public belief in period t :

$$\mu_t(w) = \alpha_s, \mu_t(m) = 0, \text{ and } \mu_t(s) = 1 - \alpha_s.$$

Equation (3) in the paper thus implies that each period- t player attacks if and only if her signal is below x_t^* , which is determined by

$$\frac{\alpha_s \phi(\sqrt{\beta}(x_t^* - w))}{\alpha_s \phi(\sqrt{\beta}(x_t^* - w)) + (1 - \alpha_s) \phi(\sqrt{\beta}(x_t^* - s))} = c. \quad (5)$$

Due to the monotone likelihood ratio property and the unbounded likelihood ratio property of the standard normal distribution, there is a unique solution, x^* , to equation (5), which is determined by α_s , β , and c .

¹This is robust if the transition rule is asymmetric.

In any period t , because $\mu_t(m) = \alpha_s > 0$, players have the dominant region of attacking. Thus, in every period, players will employ a cutoff strategy in which they attack if and only if their signals land below some threshold. In particular, the threshold in any period is bounded from below by x^* . This leads to the prediction of the status quo.

Proposition 2 *When the status quo is persistent, but not constant, over time, and the transition rule satisfies the property in Equation (1), the regime changes almost surely.*

Proof 3 (Proof of Proposition 2) *In each period t , the regime changes with probability*

$$\begin{aligned}
& \Pr(b_t \geq 1 \text{ and } \theta_t = w) \\
&= \Pr(b_t \geq 1 | \theta_t = w) \Pr(\theta_t = w) \\
&> \Phi(\sqrt{\beta}(x_t^* - w)) \mu_t(w) \\
&\geq \Phi(\sqrt{\beta}(x^* - w)) \alpha_s \\
&> 0.
\end{aligned}$$

Since the regime change probability is bounded away from zero in every period, the regime changes almost surely.

Now, let us analyze the second case, in which once the status quo becomes medium or strong, the status quo can no longer be weak. In this case, I allow the strength of the status quo to vary between $\theta = m$ and $\theta = s$, and the transition probability to follow equation (5).

If the weak status quo is constant over time, Proposition 3 shows that the weak status quo is overthrown almost surely. If the weak status quo is only persistent, but not constant, over time, then, because the state will not change from medium or strong to weak, the status quo will become medium or strong almost surely. Therefore, without loss of generality, I analyze the case where $\theta = w$ is ruled out. Proposition 3 below summarizes the regime change results in this case.

Proposition 3 *Suppose the strength of the status quo is not constant but persistent. In addition, let the transition rule of the strength of the status quo satisfy equation (2) and equation (3); then*

1. *if $\tilde{\mu} \geq 1/2$, in any equilibrium the status quo survives with positive probability; and*
2. *if $\tilde{\mu} < 1/2$, there are equilibria in which the status quo changes almost surely.*

Proof 4 (Proof of Proposition 3) *Without loss of generality, I assume that $\theta = w$ is ruled out in period T_0 due to a failed attack by one player. Because of equation (2), all subsequent players know that θ must be either m or s in subsequent periods. It follows from*

the transition rule described by equation (3) that there exists a unique stable distribution over $\{m, s\}$ in which the probability of $\theta = m$ is $1/2$. I can now show that the regime change result depends upon whether $\tilde{\mu} \geq 1/2$.

Let us look at part (i) first. There are two cases. First, if $\mu_{T_0+1}(m) < \tilde{\mu}$, no player attacks in period $T_0 + 1$, and thus the no-attack history will not convey any new information. Yet, subsequent players' public beliefs about $\theta = m$ will all be below $\tilde{\mu}$: if $\mu_{T_0+1}(m) < 1/2$, the public beliefs are increasing over time, but it is bounded above by $1/2$, and so, for any $t > T_0$, $\mu_t(m) < \tilde{\mu}$; if $1/2 < \mu_{T_0+1}(m) < \tilde{\mu}$, the public beliefs are decreasing, and so, for any $t > T_0$, $\mu_t(m) < \tilde{\mu}$. Therefore, in any equilibrium, if $\mu_{T_0+1}(m) < \tilde{\mu}$ (recall that period T_0 is the time when $\theta = w$ is ruled out), the regime will not change.

In the second case, $\mu_{T_0+1}(m) \geq \tilde{\mu}$, and so, at least in period $T_0 + 1$, players can employ a strategy that induces a positive attacking probability. However, if there is no attack in a period when players attack with positive probability, the subsequent players' public beliefs about $\theta = m$, $\mu_t(m)$, will decrease. Then, after Q attacking periods, $\mu_t(m)$ will be below $\tilde{\mu}$, which leads to no subsequent attack. Since Q is finite, such a history will be reached with positive probability. Hence, in any period, with positive probability the regime will not change.

Next, I show that part (ii) is valid. Let us consider the lowest public belief about $\theta = m$. Because even if $\theta_t = s$, $\Pr(\theta_{t+1} = m) = 1 - \gamma$, the lowest public belief about $\theta = m$ is $1 - \gamma$. I construct an equilibrium as follows. When $\mu_t(m) < \tilde{\mu}$, no player attacks; if the attack in period T fails, all players refrain from attacking from period $T + 1$ to period $T + Q$, and period- $T + Q + 1$ players attack. Here, Q is sufficiently large, so that if in period t , $\mu_t(m) = 1 - \gamma$, we still have $\mu_{t+Q}(m) > \tilde{\mu}$. Since $\tilde{\mu} < 1/2$, when $\mu_t(m) < \tilde{\mu}$, the public belief keeps increasing, and there will be a finite Q such that the above conditions hold. Therefore, in such an equilibrium, there are a potentially infinite number of attacking periods.

In each attacking period, we must have $\mu_t(m) > 1 - \gamma$. Also, given this, we also see that the players' threshold in an attacking period is higher than \tilde{x} , implying that if period t is an attacking period, the probability that both period t players attack is bounded from below. Now, since there are infinitely many periods in which the regime change probability is bounded from below, the probability that there is one period in which $\theta = m$ and both players attack is one. That is, the regime changes almost surely.

Appendix B.3 N -Player $(N + 1)$ -State Dynamic Regime Change Game

The core model has three possible levels of strength of the status quo ($|\Theta| = 3$) and two new, short-lived players in each period. I now extend the core model to a dynamic regime change game with N new, short-lived players in each period and $N + 1$ possible levels of the status quo strength.

Suppose that $\Theta = \{m_1, \dots, m_{N+1}\}$ with $N > 2$ and that $m_1 < m_2 < \dots < m_{N+1}$ is the set of states. When $\theta = m_k$, the regime changes if and only if at least k players choose to simultaneously attack. Hence, when $\theta = m_1$, attacking is the dominant action for players because one player can trigger regime change by attacking. When $\theta = m_{N+1}$, not attacking is the dominant action for players, because at least $N + 1$ attacks are required to trigger regime change, but the maximum number of possible attacks is N . In all other states, players may cooperate at different levels.

Unlike in Proposition 1 and Proposition 2, given some strategy profile of other players s_{-i} , the best response of player i need not be a cutoff strategy. For example, when all other players choose to attack the status quo if and only if their private signals are in a small neighborhood of m_N , player i 's best response is to attack the status quo when x_i convinces player i that $\theta = m_N$. Therefore, in this extension, I focus only on monotone equilibria in which players' strategies are decreasing in their own private signals (i.e., the players attack for low private signals, and do not attack for high private signals). Lemma 2 below shows that if all other players follow these kinds of strategies, player i 's best response is a cutoff strategy to attack for low signals, and not to attack for high signals.

Lemma 2 *In an N -player $N + 1$ -state static regime change game, if S_j is a cutoff strategy such that $S_j = 1$ if $x_j \leq \bar{x}_j$ and $S_j = 0$ if $x_j > \bar{x}_j$ for all players $j \neq i$, then player i 's best response is a cutoff strategy such that $S_i = 1$ if $x_i \leq \bar{x}_i$ and $S_i = 0$ if $x_i > \bar{x}_i$.*

Proof 5 (Proof of Lemma 2) *I prove this lemma using the general private signal structure mentioned in footnote 4 of the paper. Let $L_k(x) = \frac{f(x|m_k)}{f(x|m_1)}$ be the likelihood ratio, then $L_k(x)$ is increasing in x for all k and $L_k(x)/L_{k'}(x)$ is increasing in x for any $k > k'$. Given S_j such that $S_j = 1$ if $x_j \leq \bar{x}_j$ and $S_j = 0$ if $x_j > \bar{x}_j$ for all players $j \neq i$, if player i chooses to attack, then conditional on $\theta = m_k$, the probability of regime change is $Z_k = \Pr(\text{there are at least } k - 1 \text{ players choosing to attack besides player } i | m_k)$. Note that Z_k is independent of x_i , and $Z_{k+1} \leq Z_k \leq 1$ for all $k = 1, 2, \dots, N$. Hence, the interim payoff of player i when she observes private signal x_i and chooses to attack is:*

$$u_i(x_i, S_{-i}) = \frac{\sum_{k=1}^N \mu_k L_k(x_i) Z_k}{\sum_{k=1}^{N+1} \mu_k L_k(x_i)}.$$

Now, consider two private signals of player i : x and x' with $x < x'$. Denote $L_k(x) = L_k$ and

$L_k(x') = L'_k$. Then,

$$\begin{aligned}
& u_i(x, S_{-i}) - u_i(x', S_{-i}) \\
&= \frac{1}{Q} \left[\left(\sum_{k=1}^N \mu_k L_k Z_k \right) \left(\sum_{k=1}^{N+1} \mu_k L'_k \right) - \left(\sum_{k=1}^N \mu_k L'_k Z_k \right) \left(\sum_{k=1}^{N+1} \mu_k L_k \right) \right] \\
&= \frac{1}{Q} \left\{ \sum_{k=1}^N \sum_{q \leq k} \mu_k \mu_q (L_k L'_q - L'_k L_q) (Z_k - Z_q) + \mu_{N+1} \sum_{k=1}^N (L'_{N+1} L_k - L_{N+1} L'_k) Z_k \right\}.
\end{aligned}$$

Each term in the first part is positive because $L_k L'_q - L'_k L_q < 0$ and $Z_k - Z_q < 0$ for all $q \leq k$. Every term in the second part is also positive because $L'_{N+1} L_k - L_{N+1} L'_k > 0$ for all $k \leq N$. Therefore, $u_i(x, S_{-i})$ is decreasing in x . Together with the dominant region of attacking and the dominant region of not attacking, this monotonicity implies that player i 's best response is also a cutoff strategy such that $S_i = 1$ if $x_i \leq \bar{x}_i$ and $S_i = 0$ if $x_i > \bar{x}_i$.

Because of the strategic complementarity, if a monotone equilibrium exists, it is symmetric. Similarly to the proof of Proposition 1, a monotone equilibrium can be shown to exist for any interior prior belief. Also, there exists $\tilde{\mu}(m_{N+1}) \in (0, 1)$ such that when m_k is ruled out (and hence all states $m_{k'}, k' < k$ are ruled out) and the public belief about $\theta = m_{N+1}$ is greater than $\tilde{\mu}(m_{N+1})$, the unique rationalizable action for any private signals is not to attack. Hence, restricting attention to monotone equilibria, the static game analysis is similar to that of the core model.

For any fixed monotone equilibrium, the eventual outcome of regime change in this extended model is similar to that of the core model. When $\theta = m_1$, an $N + 1$ state version of Proposition 3 implies that the status quo changes almost surely. When $\theta = m_{N+1}$, by assumption, the status quo never changes. For any state m_k ($1 < k < N + 1$), inefficient herds form, and the informational cascades result in coordination failure. This result can be shown by analyzing the dynamic of attacking along the outcome $\bar{h}^\infty = (1, 0, 0, \dots)$. The attack in the first period by one player cannot be successful conditional on m_k ($1 < k < N + 1$), and so players lose the dominant region of attacking from the second period on. Then the fact that no attack occurs in a large number of periods promotes players' public beliefs about $\theta = m_{N+1}$ above $\tilde{\mu}(m_{N+1})$. Consequently, inefficient herds emerge, and coordination is impossible.