Penn Institute for Economic Research
Department of Economics
University of Pennsylvania
3718 Locust Walk
Philadelphia, PA 19104-6297
pier@econ.upenn.edu
http://economics.sas.upenn.edu/pier

PIER Working Paper 11-021

"Coordination and Social Learning"

by

Chong Huang

http://ssrn.com/abstract=xxxxxxx

Coordination and Social Learning

Chong Huang*

August 5, 2011

*Contact information: hchong@sas.upenn.edu. I am indebted to my advisor, George Mailath, for his invaluable guidance and encouragement. I thank my committee members Itay Goldstein, Qingmin Liu and Andrew Postlewaite, for their great comments that have improved the paper. I also thank Nageeb Ali, David Dillenberger, Hanming Fang, Philipp Kircher, Steven Matthews, Stephen Morris, Mallesh Pai, Alessandro Pavan, Alvaro Sandroni, Lones Smith, Jusso Välimäki, Muhamet Yildiz, participants at the micro lunch seminar at UPenn, and attendees at 2010 EconCon, the 2011 North American summer meeting of the Econometric Society 2011, and the 2011 International Conference on Game Theory for helpful comments. All errors are mine.

Coordination and Social Learning

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. So coordination affects the aggregation of private signals through players' optimal choices. Such endogenous provision of public information results in inefficient herds with positive probability, even though private signals have an unbounded likelihood ratio property. Therefore, social learning is a source of coordination failure. An extension shows that if players could individually learn, inefficient herding disappears, and thus coordination is successful almost surely. This paper also demonstrates that along the same history, the belief convergence differs in different equilibria. Finally, social learning can lead to higher social welfare when the fundamentals are bad.

JEL Classification: C72, C73, D82, D83

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

1 Introduction

In many economic and social environments featuring coordination motives, such as investments, currency attacks, bank runs and political revolutions, agents are often uncertain about the coordination results. So can coordination outcomes be implemented in these environments if coordination is socially optimal? Learning about the fundamentals is a potential way to resolve this problem: the more information agents possess, no matter whether the incremental information is public or private, the more likely there exists an equilibrium in which agents choose to coordinate. But is this argument true, especially in a dynamic environment?

Imagine a dynamic world in which each individual has a noisy signal about the fundamentals. The public information is the behavior of previous players. If the public history successfully aggregates private signals, the public history conveys arbitrarily accurate information about the fundamentals, and thus coordination outcomes can be reached. Conversely, if a herd forms, that is, if players choose "not coordinate" as most previous players did, ignoring their own private signals (see Banerjee (1992) and Bikhchandani, Hirshleifer, and Welch (1992)), the information aggregation will be unsuccessful. In this case, learning from the public history, or *social learning*, may result in coordination failure. To the extent that any decision maker's payoff is independent of other players' choices, Smith and Sørensen (2000) show that the public history will successfully aggregate private signals, if and only if the strength of private signals is unbounded.[1] However, does this conclusion hold when the economy features a coordination motive?

The above questions can be formulated as the investigation of the interaction between coordination and social learning. In this paper, I study this interaction 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-lived) players. They commonly observe previous plays, and each of them receives one piece of private information about the status quo.[2] Based on this information, they update their beliefs about the true state of the status quo, which is unknown but fixed. Because any individual can observe only one piece of private information, perfect individual learning is impossible. 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,

[1] Lee (1993) analyzes a social learning model with continuous action space and binary signal space. Inefficient herding does not appear, because no information goes unused. Bikhchandani, Hirshleifer, and Welch (1998) and Chamley (2004) provide surveys of this literature.

[2] As in the social learning literature, I assume players are short-lived. First, this assumption makes the model tractable and naturally rules out perfect individual learning. Second, finite long-lived players can coordinate to experiment as their discount factor becomes arbitrarily close to 1, further complicating the analysis.

which favors the regime change.) A player choosing to attack receives a positive payoff if the regime changes; she receives a negative payoff otherwise. Not attacking is a safe action, giving zero payoff whether the regime changes or not. The true state 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 required to trigger the regime change; if the status quo is strong, the status quo can never be beaten.

The main result in this paper is shown by comparing equilibrium outcomes in the medium state with those in the weak state. As a benchmark, in the weak state, coordination is not necessary, so the status quo is in place only if there is no attack. But as the probability of attacking decreases, the "no attack" history tells players less and less. 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, inefficient herds never form, and the regime changes almost surely. Different from the weak state, in the medium state, when coordination is commonly known to be necessary from a previous failing attack by one player, in order to attack, players must form high "common beliefs" about the medium state. Therefore, the information from the public history may dominate any private signal. That is, if the public history makes a player pessimistic ex-ante, then, even if she receives an extremely informative private signal favoring the medium state, she cannot be confident that her opponent receives an "attacking" signal. So, in any equilibrium, in the medium state, inefficient herds emerge with positive probability, and the regime may not change. Therefore, in the medium state, the informational cascades result in the impossibility of coordination, so social learning is a source of coordination failure.

In a herd, public beliefs converge because players stop learning from the public history. So what is the asymptotic public belief along the outcomes in which players never stop learning from the public history? In particular, will social learning be complete if public beliefs keep changing over time? I show that the public belief convergence differs in different equilibria of this model. Take the outcome with an attack by one player in every period as an example. Along this outcome, the weak status quo is ruled out by the failing attack in the first period. Then, in the most aggressive equilibrium, in which players' strategies specify the highest possible probability of attacking in every period, the public belief about the medium state converges to a point strictly between 0 and 1. Therefore, social learning is incomplete in this equilibrium. In another equilibrium, in which players adopt strategies specifying the smallest positive probabilities of attacking (if there are any), the public belief about the medium state converges to 1, so social learning is complete in this case. In this equilibrium, in the limit, the strategy profile purifies the mixed strategy equilibrium of the complete information normal form game at the medium state.

Social learning not only drives the dynamics of attacking and partly determines the eventual fate of the regime, but it also causes social welfare to differ depending on the status quo. Considering the discounted social welfare, under a weak status quo, social learning delays the regime change, resulting in inefficiency. Under the medium status quo, social learning results in the probability of the regime change being less than 1. Consequently, when the discount factor is sufficiently close to 1, social learning leads to lower social welfare, which is inefficient. For a strong status quo, social learning prevents attacking infinitely often with probability 1, which leads to higher social welfare – provided that the discount factor is sufficiently close to 1.

The dynamic regime change game rules out the possibility of perfect individual learning and focuses on a two-player three-state case. I extend the core model in two directions. In the first extension, the economy consists of two types of players – – any player i in period t is of type i ($i = 1, 2$). Suppose in any period t, the player of type i collects all previous private signals of type i players (but not previous private signals of type j players ($i \neq j$)). So, the precision of the private signal is strictly increasing over time, and it goes to ∞ in the limit. That is, I allow perfect individual learning. In this extended model, in all nontrivial equilibria, if the status quo is in the medium state, inefficient herds do not form, and the regime changes almost surely. In the second extension, I analyze the dynamic regime change game with $N + 1$ possible states of the status quo and N new short-lived players in each period. At state n, at least n attacks are needed to trigger the regime change. So the first state is like the weak status quo in the core model, while the $(N+1)$th state is like the strong status quo in the core model. In this second extended model, in any monotone equilibrium of this game, the dynamics of attacking and the eventual outcomes of the regime change are similar to those in the core model.

1.1 Related Literature

The social value of public information has been discussed 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, the public information is exogenous. In my paper, the public information evolves endogenously as a result of agents' decisions. Therefore, coordination directly affects the provision of the public information. Consequently, the public history may fail to aggregate private signals and thus provide biased public information, which results in coordination failure.

There is a strand of social learning models that discuss herding behaviors and asynchronous coordination. In Dasgupta (2000), first movers and late movers can coordinate, so 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 complementarity is 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 the key ingredient of my paper. In addition, it is easy to show that if coordination is asynchronous, inefficient herds never form in my model.

This paper contributes to the global game literature, initiated by Carlsson and Van Damme (1993). Absent the dynamic aspect, static regime change 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, 2008). 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 happened yet, players are in a static global game. But in this static global game, multiple Bayesian Nash equilibria may exist. Such multiplicity follows two characterizations of the model. First, because the state space is discrete, there are some prior beliefs resulting in multiple Bayesian Nash equilibria, no matter how large the precision of private signals is. The necessity of the connectedness of the state space for the equilibrium uniqueness is discussed in Carlsson and Van Damme (1993), and my model provides a counter example. Second, the precision of private signals is fixed, while public beliefs evolve endogenously over time.

In a recent paper, Angeletos, Hellwig, and Pavan (2007) incorporate both individual learning and social learning into a dynamic regime change game.[3] They consider a continuum of long-lived agents, each learning the true state eventually by collecting one piece of private information in every period. Players also learn from the publicly observable fact that the regime has not changed. Furthermore, in an extension, all players observe public and private signals about the previous attacking sizes. My model differs from Angeletos, Hellwig, and Pavan (2007) mainly in that individual learning is impossible. As shown in my first extended model, individual learning overturns the social learning effect, so inefficient herding does not emerge, even when coordination is commonly known to be necessary. Hence, to analyze the

[3]Dynamic regime change games are studied as examples of dynamic global games. Dasgupta (2007) studies a two-period model dynamic global game, which allows asynchronous coordination. Other papers contributing to this growing literature include Giannitsarou and Toxvaerd (2007) and Heidhues and Melissas (2006).

interaction between coordination and social learning, one needs to go one step back to a model without individual learning. Related to the interaction between individual learning and coordination is the paper by Cripps, Ely, Mailath, and Samuelson (2008). They show that with a finite state space and conditional independent private signals, individual learning implies common learning. Therefore, since common learning is exogenous, and social learning may be suspended endogenously, high common beliefs about the medium state are formed infinitely often. As a result, inefficient herds disappear, and thus, the regime changes almost surely if the status quo is medium.

The rest of this 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 the dynamics of attacking, the eventual fate of the status quo, and the social welfare. Section 4 is devoted to two extensions of the core model. Section 5 concludes. All omitted proofs are presented in the Appendix.

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 $R_t \in \{0, 1\}$: $R_t = 0$ means the status quo, and $R_t = 1$ means the alternative. Assume the regime is in the status quo at the beginning of the game, so $R_0 = 0$; if $R_t = 1$ for some t, then $R_\tau = 1$ for all $\tau > t$; that is, once the regime changes, it stays in the alternative 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 order $w < m < s$.[4] At the beginning of the game, θ is chosen by nature according to a commonly known distribution 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 $a_{it} = 1$ means "attack," and $a_{it} = 0$ means "not 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 the regime is in the status quo at the beginning of period t and the regime changes in that period. If player i chooses to attack, she receives the payoff $1 - c$ ($c \in (1/2, 1)$); if in such a period, player i chooses not to attack, she receives the payoff 0. If the regime is in the alternative state at the beginning of period t, then no matter what player i chooses, she receives payoff $1 - c$.[5] Thus, once the regime changes, the game essentially ends. Conditional on $R_{t-1} = 0$, whether the regime changes or not in period t

[4] Notations w,m and s refer to "weak," "medium" and "strong," respectively.

[5] Assuming players receive payoff $1-c$ after the regime changes to the alternative is equivalent to assuming that the game ends once the regime changes, in terms of the strategic analysis. But this assumption makes the welfare analysis in subsection 3.3 much easier.

depends on both the strength of the status quo θ and the number of attacks $a_{1t} + a_{2t}$. For the weak status quo, the attack by one player is sufficient for the regime change; for the medium status quo, the attack by one player is not enough, but the attack by two players can trigger the 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 the decision, period t player i observes a private signal $x_{it} = \theta + \xi_{it}$. $\xi_{it} \sim \mathcal{N}(0, 1/\beta)$, where $\beta \in \mathbb{R}_{++}$ is the common precision of players' private signals.[6] ξ_{it} is independent of θ and independent across i and across t. Thus, all players' private signals are conditionally independent. Besides private signals, at the beginning of any period $t \geq 2$, players are aware of the public history about the number of players choosing to attack ($\tau < t$). Denote a typical public history by $h^t = (b_1, \ldots, b_{t-1})$, 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 a period t player i's strategy by $s_{it} : H^t \times \mathbb{R} \to \{0, 1\}$. So $s_{it}(h^t, x_{it})$ is the action player i chooses, given the public history h^t and 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. Call μ_t the public belief in period t.

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

1. *For any t, given μ_t, (s_{1t}, s_{2t}) forms a Bayesian Nash equilibrium in the static game;*
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 incentive when making 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 only specifies how to calculate public

[6] By the Gaussian assumption, it is easy to calculate players' belief updates. In addition, the Gaussian assumption has a very clear description about 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: $\lim_{x \to \underline{X}} f(x|\theta)/f(x|\theta') = +\infty$ and $\lim_{x \to \bar{X}} f(x|\theta)/f(x|\theta') = 0$, whenever $\theta < \theta'$; and (4) monotone likelihood ratio: if $\theta < \theta'$, $f(x|\theta)/f(x|\theta')$ is strictly decreasing in x.

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 game Bayesian Nash equilibrium, their equilibrium strategies are not affected by plays on the off-equilibrium path. Therefore, to simplify the analysis, I only require 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. To simplify, if the environment 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 detailed analysis. First, since $\mu_1(\theta) > 0$ for all $\theta \in \Theta$, and the regime does not change with probability 1 if there is no attack, $\mu_t(\theta) > 0$ for all $\theta \in \Theta$ after the public history without any attack. Second, if one player attacks in period t and $R_t = 0$, players in the subsequent periods learn immediately that $\theta \neq w$, that is, $\mu_\tau(w) = 0$ for all $\tau > t$. Third, if both players attack in period t and $R_t = 0$, players learn immediately that $\theta = s$. Therefore, as 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, since $\mu_1(\theta) > 0$ for all $\theta \in \Theta$, in any period t, $\mu_t(\theta) > 0$ implies $\mu_t(\theta') > 0$ for $\theta < \theta'$. Among these three cases, the one with $\mu_t(s) = 1$ is trivial. Because "not attack" is the dominant action in this case, the unique Bayesian Nash equilibrium is that both players choose not to attack for all their private signals.

Now, suppose $\mu_t(\theta) > 0$ 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}. Then from Bayes' rule, the posterior belief about θ is:

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

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

$$\begin{aligned}
& E_{x_{jt}} u_{it}(1, x_{it}, s_{jt}) \\
= \ & \rho(w|x_{it}) + \Pr(s_{jt} = 1, m|x_{it}) - c && (1) \\
= \ & \rho(w|x_{it}) + \rho(m|x_{it}) \Pr(s_{jt} = 1|m) - c && (2) \\
= \ & \frac{\mu_t(w)\phi(\sqrt{\beta}(x_{it} - w))}{\sum\limits_{\theta' \in \Theta} \mu_t(\theta')\phi(\sqrt{\beta}(x_{it} - \theta'))} + \frac{\mu_t(m)\phi(\sqrt{\beta}(x_{it} - m))}{\sum\limits_{\theta' \in \Theta} \mu_t(\theta')\phi(\sqrt{\beta}(x_{it} - \theta'))} \Pr(s_{jt} = 1|m) - c. && (3)
\end{aligned}$$

The fact that (1) implies (2) is because players' private signals are independent conditional on θ. Note $\rho(w|x_{it}) \to 1$ as $x_{it} \to -\infty$; 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 $\forall s_{jt}$, so the set $[\bar{x}_t, +\infty)$ is called 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 point $x_t^* \in \mathbb{R}$:*[7]

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

In addition, x_t^* is the solution to the equation

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

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; $\mu_t(m) > 0$ implies that the regime may change if both players choose to attack. Hence, players in this case are playing a coordination game. Obviously, the strategy profile with coordination failure, that is, both players choose not to attack, is an equilibrium. So is no attack the unique Bayesian Nash equilibrium in period t? That is, are there any equilibria with positive probability of attacking?

The only state for which the regime can change is $\theta = m$, so 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 belief about $\theta = m$ is high (players are optimistic), cooperation is possible; conversely, when players' public belief about $\theta = m$ is low (players are pessimistic), even if one player observes an extremely negative signal and is convinced that $\theta = m$, she won't attack. This is because she believes that the probability

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

of her opponent observing a signal favoring $\theta = m$ is very low. With the sufficiently informative signals assumption given below (I maintain this assumption throughout), Proposition 2 formally shows the above intuition.

Assumption 1 *Private signals are sufficiently informative:* $\Phi(\frac{\sqrt{\beta}}{2}(s-m)) > c$.

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

1. *If* $\mu_t(m) < \tilde{\mu}(m)$, *there is no equilibrium with attack;*

2. *If* $\mu_t(m) > \tilde{\mu}(m)$, *there are two equilibria with attack, which are symmetric and in cutoff strategies. The threshold point for any equilibrium is in* $(m, +\infty)$*;*

3. *If* $\mu_t(m) = \tilde{\mu}(m)$, *there exists a unique* $\tilde{x} \in (m, +\infty)$ *such that* $(\tilde{x}, \tilde{x})$ *is the unique equilibrium with attack.*

The key step to prove 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 $\tilde{\mu}(m)$ is determined by the following equation

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

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

Since when $\mu(w) = 0$, it is hard to analyze solutions to $G(x, \mu) = 0$, I define

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

Then when $\mu(w) = 0$, $G(x, \mu) = \rho(m|x) g(x, \mu)$. Because $\rho(m|x) > 0$ for all $x \in \mathbb{R}$, when $\mu(w) = 0$, $G(x, \mu) = 0$ if and only if $g(x, \mu) = 0$. And since $g(x, \mu)$ has nice properties, whenever $\mu(w) = 0$, instead of $G(x, \mu) = 0$, I analyze solutions to the equation $g(x, \mu) = 0$.

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

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 $R_t = 0$, given μ_t, x_t^* and b_t, employ Bayes' rule to calculate μ_{t+1}.

2.3 Multiple Equilibria

From the above equilibrium characterization algorithm, multiple equilibria exist. The multiplicity is not due to plays on the off-equilibrium path but stems from the structure of the dynamic regime change game.

From Proposition 2, when $\theta = w$ is ruled out (after observing a failing attack by one player in some period), if $\mu_t(m) \geq \tilde{\mu}(m)$, $G(x, \mu_t) = 0$ has at least one solution; so there are multiple Bayesian Nash equilibria in period t. Thus, there must be multiple equilibria in the dynamic regime change game. This kind of multiplicity is driven by the coordination property when $\mu(w) = 0$, as in Angeletos, Hellwig, and Pavan (2007).[8]

However, in the dynamic regime change game, there is another source for multiplicity. After a history h^t without any attack, period t players are still in a global game. In the static global game literature, the state space is usually assumed to be connected. Then, for a given prior belief, as the precision of private signals becomes arbitrarily large, a unique equilibrium of the static regime change game can be established by interim iterated elimination of strictly dominated strategies (Morris and Shin (2003)). But in my model, two features lead to multiple Bayesian Nash equilibria in the stage game after the no attack history. First, as discussed by Carlsson and Van Damme (1993), with a finite state space, multiple Bayesian Nash equilibria exist when players' common prior beliefs put a sufficiently high weight on the medium state (no matter how large β is). Second, the precision of private signals is fixed, while players' public beliefs evolve endogenously. Since there exist public beliefs in period t with $\mu_t(\theta) > 0$ for all $\theta \in \Theta$, such that multiple Bayes Nash equilibria exist in period t (see the discussion in the proof of Proposition 1 in the Appendix for details), the dynamic regime change game has multiple equilibria.

The equilibrium definition requires that all players commonly know the strategies adopted by previous players, that is, x_t^* is common knowledge to all players arriving after period t. The emergence of multiple equilibria makes this assumption even stronger. For example, when state w is ruled out and $\mu_t(m) > \tilde{\mu}_t(m)$, no attack in period t is consistent with both the strategy profile $(-\infty, -\infty)$ (the pure no attack strategy profile) and the strategy profile (x_t^*, x_t^*) with $x_t^* > m$ (when both players' private signals land above x_t^*). However, the belief consistency is part of an equilibrium in the dynamic regime change game, so this common knowledge assumption is natural when an equilibrium is fixed. Therefore, in all the analysis in section 3, the strategy profile is fixed. In addition, I will show that attacking dynamics in different equilibria are driven by different social learning processes and, thus, have dramatically different properties.

Two classes of equilibria have nice properties and are easy to analyze, which I define

[8] In their model, after the first period, players lose the dominant region of attacking, so players are in a coordination game from the second period on.

as follows:

Definition 2 $E^{max} = \{(x_t^*, x_t^*)_{t=1,\ldots}, (\mu_t^*)_{t=1,\ldots}\}$ *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 3 $E^{min} = \{(x_t^*, x_t^*)_{t=1,\ldots}, (\mu_t^*)_{t=1,\ldots}\}$ *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 exists),* $x_t^* \in \mathbb{R}$ *and* $x_t^* \leq x_t'$*; if there is no such* x_t'*,* $x_t^* = -\infty$*.*

So 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 $\mu(m) > \tilde{\mu}(m)$.

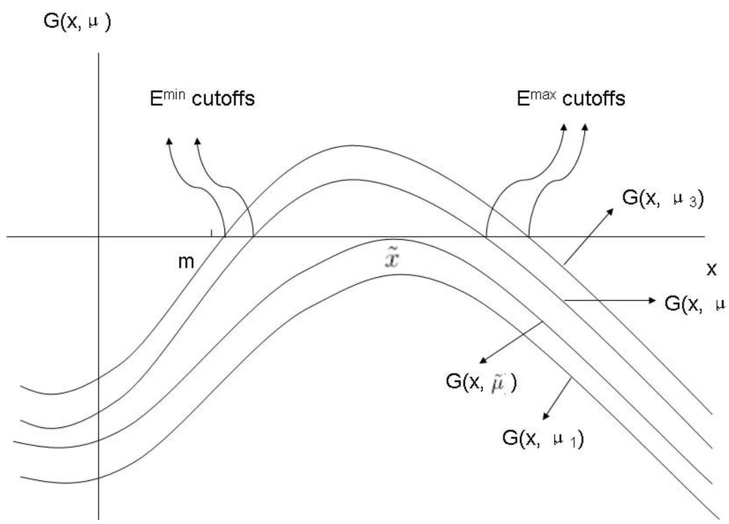

Figure 1: Function $G(x;\mu)$ with $\mu_3(m) > \mu_2(m) > \tilde{\mu}(m) > \mu_1(m)$.

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 eventual outcome of the regime change 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. Additionally, as shown in subsection 3.3, in different states, social learning has different effects on social welfare.

3.1 Dynamics of Attacking and Regime Change

Let $\hat{h}^t \equiv (0, \ldots, 0)$ denote the history without any attack. When $\theta = w$, conditional on $R_{t-1} = 0$, $\hat{h}^t$ is the only history period t players can observe. So the dynamics of attacking along $\hat{h}^t$ determine the eventual outcome of the regime change, conditional on $\theta = w$. Hence, I first study the dynamics of attacking along $\hat{h}^t$.

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

$$\mu_t(\theta) = \frac{\mu_1(\theta) \prod_{\tau=1}^{t-1} [\Phi(\sqrt{\beta}(\theta - x_\tau^*))]^2}{\sum_{\theta' \in \Theta} \mu_1(\theta') \prod_{\tau=1}^{t-1} [\Phi(\sqrt{\beta}(\theta' - x_\tau^*))]^2}, \quad \forall \theta \in \Theta, \tag{6}$$

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 $\theta \in \Theta$, $x_1^* \in \mathbb{R}$ from Proposition 1, which in turn implies that $\mu_2(\theta) > 0$ for all $\theta \in \Theta$. Then, by induction, along the 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 happen with the highest probability in the weak state. Thus, the no attack outcome in period t lowers period $t+1$ players' belief about $\theta = w$. As a result, along $\hat{h}^t$, $\{\mu_t(w)\}_t$ is a bounded and strictly decreasing sequence, which converges to $\mu_\infty(w) \geq 0$. If $\mu_\infty(w) > 0$, the dominant region of attacking has a positive measure in the limit; so the probability of attacking is bounded away from 0, then $\mu_t(w) \to 0$ from the belief updating equation (6). Hence, $\mu_t(w)$ must converge to 0. For the probability of attacking along $\hat{h}^t$, suppose 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, that the probability of attacking is bounded away from 0 implies that the corresponding cutoff points are bounded below. This contradiction implies that the probability of attacking is converging to 0. These arguments are formally stated in the following proposition:

Proposition 3 *Fix any equilibrium. Along $\hat{h}^t$, $\mu_t(w) \to 0$ and* $\Pr(b_t > 0 | \hat{h}^t) \to 0$.

When the status quo is weak, the attack by one player can trigger the regime change. But because of the dominant region of not attacking, in any equilibrium, players choose not to attack with positive probability in any period. Also, because the probability of attacking converges to 0, the asymptotic probability of the regime change when $\theta = w$ depends on the speed at which the probability of attacking converges to 0. If it is too fast, the status quo may not fall. But while players continue to learn, they learn slowly after some period T, because the probability of attacking is arbitrarily small for all $t > T$. This slow social learning process implies that the probability of attacking cannot converge to 0 too fast. As a result, the regime changes eventually if the status quo is weak.

Proposition 4 *If $\theta = w$, in any equilibrium, the regime changes almost surely.*

Proof. Fix an 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 $\mathbb{P}_w$ and $\hat{\mathbb{P}}$ are probability measures induced on $\Theta \times \{0,1,2\}^{\infty}$ by $\mathbb{P}$, conditioning 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 $\mathbb{P}$. So $\{\mu_t(w)\}_t$ 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 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, 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 bigger than $\frac{1-\mu_0(w)}{\mu_0(w)}$ for very big t since $\mu_\infty(w)(a) = 0, \quad \forall a \in A$, 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 $\hat{h}^\infty$ is the only outcome in which the weak status quo never falls).

■

The conclusion that the weak status quo falls almost surely is due to the assumptions about private signals. In my model, private signals are normally distributed, so they have continuous support and an unbounded likelihood ratio property. In particular, the existence of a dominant region of attacking implies that there are always private signals that can overturn the public beliefs.

Now let me 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 that the probability of attacking in the first period is positive. Hence, the medium status quo falls with positive probability. So does the medium status quo fall almost surely? Because each player's private signal has the unbounded likelihood ratio property, for any prior beliefs, there are signals making her posterior belief about $\theta < s$ arbitrarily close to 1. Therefore, it seems that in a nontrivial equilibrium (in which players choose to attack with positive probability whenever possible), with probability 1, there is one period in which both players in that period receive signals informing them that $\theta < s$. As a result, the medium status quo should fall almost surely. But is this argument right?

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 attacks afterward. If this outcome is reached, the medium status quo does not fall. 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 pure no attack strategy. In the strategy profile with finite $\mathcal{Q}$, the outcome $\bar{h}^\infty$ is reached with strictly positive probability. Hence, unless there is an equilibrium in which $\mathcal{Q}$ is infinite, in any equilibrium, the regime dose not change with positive probability if the status quo is medium.

From the failing attack in the first period, players rule out the weak state. By Proposition 2, in any $t \in \mathcal{Q}$, the probability of attacking is bounded away from 0. Then the observation of "no attack" in period $t \in \mathcal{Q}$ makes subsequent players increasingly pessimistic. Also, the necessary condition for period $t \in \mathcal{Q}$ players to attack with positive probability is $\mu_t(m) \geq \tilde{\mu}(m)$. So if $\mathcal{Q}$ is infinite, after a period $T \in \mathcal{Q}$, players' public beliefs about $\theta = m$ drop below $\tilde{\mu}(m)$, which contradicts the assumption that $\mathcal{Q}$ is infinite.

Proposition 5 *There exists $Q \in \mathbb{N}$, such that in any equilibrium, along the outcome $\bar{h}^\infty$, there are at most Q periods in which players' strategies specify positive probabilities of attacking.*

Corollary 1 *In any equilibrium, $\mathbb{P}_m(\bar{h}^\infty) > 0$, so the regime does not change with positive probability when the status quo is medium.*

The outcome $\bar{h}^\infty$ is just an example of all outcomes, which result in the survival of the status quo and are realized with positive probabilities in any equilibrium, conditional on the medium state. All these outcomes share three features: (1) there is no period in which both players choose to attack; (2) there are finite (at least one) periods in which one player chooses to attack; and (3) after the last period in which one player chooses to attack, no attack happens ever again. These are herding outcomes; that is, players in later periods join the "not attack" herds, ignoring their own private signals no matter how informative such signals are. These herding outcomes are inefficient, because conditional on the medium status quo, the $(attacking, attacking)$ strategy profile Pareto dominates the no attacking strategy profile. I will analyze the social welfare due to these herding outcomes in more detail in subsection 3.3.

Comparing the equilibrium outcomes in the weak state with those in the medium state, we can see the interaction between coordination and social learning. "Not attack" herds do not appear in the weak state, as in Smith and Sørensen (2000). But in any equilibrium, players join the "not attack" herds with positive probability when the status quo is medium. The difference stems from the coordination requirement in the medium state and the signal structure of global games. Consider an equilibrium in which whenever strategies inducing positive probability of attacking can form a Bayesian Nash equilibrium in any period t, period t players adopt such strategies. When the belief about the weak status quo is positive, extreme negative signals lead players to attack, no matter what their opponents choose. This is the reason why "not attack" herds never start in the weak state. In the medium state, once the weak status quo is ruled out, it is common knowledge that coordination is necessary for their attacks to succeed. Therefore, in order to coordinate, players in the same period must form high common beliefs about $\theta = m$. That is, fix any $q \in (0, 1)$, players must believe $\theta = m$ with probability at least q, must believe with probability at least q that their opponents believe $\theta = m$ with probability at least q, and so on (see, *e.g.*, Monderer and Samet (1989) and Morris and Shin (2007)). When the public belief about $\theta = m$ is sufficiently low, to form high common belief about $\theta = m$, the precision of private signals must be big enough. However, in the model, the precision of private signals is fixed, and along the outcome $\bar{h}^\infty$ players' public beliefs about $\theta = m$ decrease. Therefore, players cannot form high common beliefs about $\theta = m$ eventually, and thus the "not attack" herds start.

The herding outcome is due to the informational externality, which is different from the pure coordination failure. In some equilibria, after the weak status quo is ruled out, players simply choose not to attack. This "no attack" outcome is purely because of coordination failure. Proposition 5 shows, however, that the informational externality causes the impossibility of coordination in all equilibria.

3.2 Belief Convergence in Different Equilibria

After the weak status quo is ruled out, Proposition 5 implies that, in all equilibria, along the outcome $\bar{h}^\infty$, public beliefs about the medium status quo converge to $\mu_\infty(m) \in (0, 1)$. This convergence is just 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. Then, 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.

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

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 $\tilde{x} < \frac{m+s}{2}$. (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 $G(x, \mu') = 0$. Consider the outcome $\tilde{h}^\infty \equiv (1, 1, \dots)$ (exactly one player chooses to attack in every period).

Proposition 6 *Suppose $\tilde{x} < \frac{m+s}{2}$. Along the outcome $\tilde{h}^\infty$.*

1. *In E^{max}, $\mu_t(m)$ monotonically converges to $\mu'(m) \in (0, 1)$. So 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.*

The detailed proof is in the Appendix, but the intuition 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 $\tilde{h}^\infty$ in E^{max} and E^{min} respectively. In Figure 2, $\mu'(m)$ is the unique fixed point in $(\tilde{\mu}(m), 1)$. Therefore, in E^{max}, along $\tilde{h}^\infty$, $\mu_t(m) \to \mu'(m)$, 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 $\tilde{\mu}(m)$ is 1, and it is stable. This means along $\tilde{h}^\infty$, $\mu_t(m) \to 1$. Hence, the probability of attacking converges to c. So 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 6 can be seen in Figure 1: if $\mu_{t+1}(m) > \mu_t(m)$, $g(x, \mu_{t+1}) > g(x, \mu_t), \forall x > m$, so $x_{t+1}^* > x_t^*$ in E^{max} and $x_{t+1}^* < x_t^*$ in E^{max}.

Proposition 6 shows that along the outcome $\tilde{h}^\infty$, the convergence of $\mu_t(m)$ differs in E^{max} and E^{min}. In E^{max}, the failing 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). Then $b_{t+1} = 1$ 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

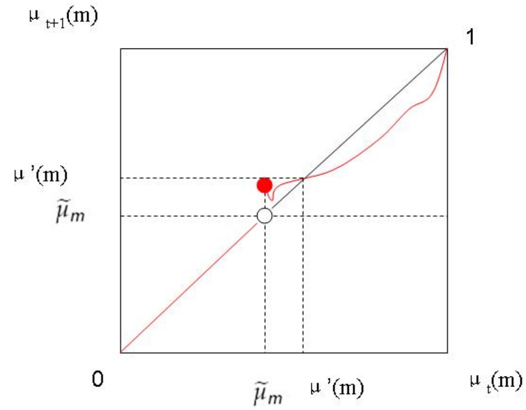

Figure 2: Evolution of Public Beliefs in E^{max}.

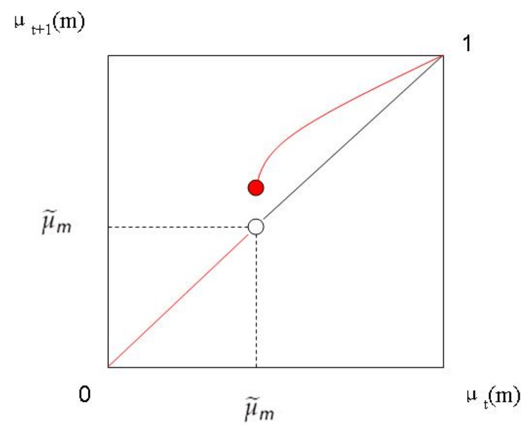

Figure 3: Evolution of Public Beliefs in E^{min}.

informative over time, because the threshold points of players' strategies are decreasing over time (from Figure 1, all threshold points in E^{min} are smaller than $\frac{s+m}{2}$). As a result, in E^{min}, along $\tilde{h}^\infty$, $\mu_t(m)$ converges to 1. Put differently, along the same outcome $\tilde{h}^\infty$, in E^{max}, the public history provides decreasingly informative evidence over time, so the social learning is incomplete. But in E^{min}, the public history provides increasingly informative evidence over time, so the social learning is complete.

3.3 Social Welfare

In this section, I analyze the effect of social learning on social welfare. Imagine the scenario that all players believe they are in a static regime change game, so that there is no social learning. From Proposition 1, in any equilibrium, players attack the status quo with positive probability (bounded away from 0) in each period. Consider the ex-post social welfare $\mathcal{W} = (1-\delta)\sum_{t=1}^{\infty}\delta^{t-1}(u_{1t}+u_{2t})$, where $\delta \in (0,1)$ is the discount factor. By comparing the social welfare functions of the regime change game with and without social learning, the effect of social learning can be seen from the following proposition:

Proposition 7 *Compare with the scenario without social learning,*

1. *if* $\theta = w$*, in both* E^{max} *and* E^{min}*, for any* $\delta \in (0,1)$*, social learning leads to inefficiency; as* δ *goes to* 1*, this inefficiency disappears;*

2. *if* $\theta = m$*, in any equilibrium, there exists* $\delta_m \in (0,1)$*, such that social learning is inefficient for any* $\delta \in (\delta_m, 1)$*;*

3. *if* $\theta = s$*, in any equilibrium, there exists* $\delta_s \in (0,1)$*, such that social learning leads to a higher social welfare value for all* $\delta \in (\delta_s, 1)$*.*

The intuition behind the first part of this proposition comes from the delayed regime change due to social learning. Because the delay is only finitely long, when the discount factor is sufficiently large, such inefficiency disappears. In the second part, when $\theta = m$, social learning leads players to stop attacking with positive probability, so the regime does not change with positive probability, which is inefficient because positive utilities after the regime change cannot be collected. In the third case, the smaller the number of attacks, the higher the social welfare, because the regime cannot be beaten. Since social learning can prevent infinitely many attacks, it is more efficient.

Consider the ex-ante social welfare. Part 1 of Proposition 7 implies that social welfare with and without social learning are almost the same when δ is sufficiently close to 1. As a result, whether social learning leads to inefficiency depends only on the comparison between $\mu_1(m)$ and $\mu_1(s)$. In particular, as $\delta \to 1$, the higher $\mu_1(m)/(\mu_1(m) + \mu_1(s))$, the lower the ex-ante social welfare with social learning.

4 Extensions

In this section, I consider two extensions of the dynamic regime change game. In the first extension, I assume period t player i is of type i ($i = 1, 2$). Then suppose period t player i collects all previous private signals of type i players (but not previous private signals of type j players ($i \neq j$)). Hence, the precision of private signals increases to $+\infty$. This extension captures the idea of inter-generational information transmissions and demonstrates the effect of individual learning in the dynamic regime change game. The second extension models an N-player $(N+1)$-state dynamic regime change game. Restricted to the monotone equilibrium, I get results very similar to those in the core model.

4.1 Individual Learning

In the dynamic regime change game, because players are short-lived, they can observe only one piece of private information. Hence, no one can individually learn the true state, given that the precision of private signals, β, is a constant. Consequently, when $\theta = w$ is ruled out, the critical belief $\tilde{\mu}(m)$ is constant over time. (Recall that, if the weak status quo has been ruled out and players' public beliefs about the medium status quo are below $\tilde{\mu}(m)$, they choose not to attack, ignoring their private signals.) As a result, in any equilibrium, along the history $\bar{h}^t \equiv (1, 0, \ldots, 0)$, there is a finite subset $\mathcal{Q} \subset \mathbb{N} \setminus \{1\}$ such that after the first period attack, only in period $t \in \mathcal{Q}$, do players adopt strategies inducing a positive probability of attacking. Since after period $\max\{\mathcal{Q}\}$, players choose not to attack for all their own private signals, no subsequent player updates her public belief about $\theta = m$ from the no attack outcome after period $\max\{\mathcal{Q}\}$. So $\mu_t(m) = \mu_{\max\{\mathcal{Q}\}+1}(m) < \tilde{\mu}(m)$ for all $t \geq \max\{\mathcal{Q}\} + 1$, and period t players will choose not to attack for any private signals. This analysis leads to Proposition 5.

So what happens if individual learning is allowed? In particular, does the "not attack" herd in Proposition 5 start in a dynamic regime change game with individual learning? Furthermore, can players with private learning change the regime when the status quo is medium? To answer these questions, I incorporate individual learning into the dynamic regime change game in this extension. A straightforward way to model individual learning is to assume that period t player i is of type i, and any player i collects private signals of previous type i players. But period t player i could not observe previous private signals of type j players. Because of the normal distribution, define recursively $z_{it} = \frac{t-1}{t} z_{it-1} + \frac{1}{t} x_{it}$, where $z_{i1} = x_{i1}$. Then in terms of belief updating, period t player i's private signals could be summarized by the sufficient statistic $z_{it} \sim \mathcal{N}(\theta, 1/(t\beta))$. Hence, the precision of private signals increases over time and goes to ∞.

The most interesting case is when $\theta = m$. Consider the history $\bar{h}^t$. A failing attack by one player in the first period rules out $\theta = w$, so players lose the dominant region of attacking. While the public belief about $\theta = m$ keeps decreasing over time, the critical belief $\tilde{\mu}_t(m)$ changes over time. From Proposition 2, in any equilibrium, attacks can occur with positive probability in period t, if and only if $\mu_t(m) \geq \tilde{\mu}_t(m)$. Since $\tilde{\mu}_t(m)$ goes to 0 (by Lemma 5 in the Appendix), unless the status quo is abandoned, attacks occur in infinitely many periods. Therefore, there are equilibria in which, when $\theta = m$, the regime changes almost surely.

There are some equilibria in which the herding result in Proposition 5 disappears because the precision of private signals increases without bound. As the precision of players' private signals goes to $+\infty$ (and this is common knowledge), the correlation of players' private signals (in the same period) is arbitrarily close to 1. Therefore, when one player gets the private signal indicating that $\theta = m$, she assigns an arbitrarily high probability that her opponent's private signal also indicates $\theta = m$. As a result, for a fixed public belief (if, in an equilibrium, period t players' strategies are not to attack for all signals, then period $t+1$ players' prior beliefs are the same as period t players'), there is a sufficiently large precision of private signals such that attacking with positive probability is consistent with an equilibrium.

This result is an example in the common learning literature. Modify the common learning definition of Cripps, Ely, Mailath and Samuelson (2008) as follows:

Definition 4 $\theta \in \Theta$ *is commonly learned, if conditional on* θ, *for any fixed* $q \in (0,1)$, *a common* q*-belief about* θ *is formed infinitely often.*

Then for any β, conditional on θ, players commonly learn θ in any equilibrium, because the precision of private signals increases to $+\infty$. Since conditional on $\theta = m$, a common q-belief about $\theta = m$ is formed infinitely often, there exist equilibria in which the status quo is attacked infinitely many times. As a result, the medium status quo falls almost surely.

This extended model with individual learning is a modified model of Angeletos, Hellwig, and Pavan (2007). It shows that an economy can move back and forth between "tranquility" phases (when no attack is the unique rationalizable action) and "distress" phases (when attack is also rationalizable for a positive measure of private signals). Proposition 5 shows that this transition cannot happen in the dynamic regime change game without individual learning.

4.2 N-Player $(N+1)$-State Dynamic Regime Change Game

The core model has three possible levels of the 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 strength of the status quo.

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

Different from Proposition 1 and Proposition 2, given some strategy profile of other players s_{-i}, the best response of player i may 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 (so attack for low private signals and not attack for high private signals). Lemma 6 in the Appendix shows that if all other players are following these kinds of strategies, player i's best response is a cutoff strategy with attack for low signals and not attack for high signals. Because of the strategic complementarity, if a monotone equilibrium exists, it is symmetric. Following in a similar way the proof of Proposition 1, a monotone equilibrium can be shown to exist for any interior prior belief. Also, there exists a $\tilde{\mu}(m_{N+1}) \in (0, 1)$ such that when m_k is ruled out (so that all states $m_{k'}, k' < k$ are ruled out) and the public belief about $\theta = m_{N+1}$ is greater than $\tilde{\mu}(m_{N+1})$, not attack is the unique rationalizable action for any private signals. 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 falls almost surely. When $\theta = m_{N+1}$, by assumption, the status quo never falls. For any state m_k $(1 < k < N+1)$, inefficient herds form, and the informational cascades result in coordination failure. This result could be shown by analyzing the dynamic of attacking along the outcome $\bar{h}^\infty = (1, 0, 0, \ldots)$. The attack in the first period by one player cannot be successful conditional on m_k $(1 < k < N+1)$, so players lose the dominant region of attacking from the second period on. Then the fact that no attack happens in a large number of periods promote players' public beliefs about $\theta = m_{N+1}$ above $\tilde{\mu}(m_{N+1})$. Consequently, inefficient herds emerge, and coordination is impossible.

5 Conclusion

In this paper, I analyze the interaction between coordination and social learning in a dynamic regime change game without individual learning. I show that the inefficient herding phenomenon and the coordination failure reinforce each other. When a player believes she can change the regime by choosing to attack herself (so coordination is not necessary), she does not ignore her private signals. As a result, inefficient herds do not form. If players commonly know that coordination is necessary to trigger the regime change, they choose not to attack as the unique rationalizable action for all private signals, when they are sufficiently pessimistic about the coordination outcome. Consequently, private signals are not aggregated, and inefficient herds form. Such informational cascades result in the impossibility of coordination, so social learning is a source of coordination failure. However, once the individual learning is allowed in the dynamic regime change game, players commonly learn the true state. Therefore, the inefficient herding phenomenon disappears, and coordination is successful in some equilibria.

The interaction between social learning and coordination determines the dynamics of attacking, predicts the eventual regime change outcome, and affects social welfare in the dynamic regime change game without individual learning. If the status quo is weak, the regime changes with probability 1, though the probability of attacking converges to 0 along the no attack outcome. Conditional on the medium state, because "not attack" herds form with a positive probability after the weak state is ruled out, the regime does not change with a positive probability. I also show that along the outcome in which exactly one player chooses to attack in every period, the public belief convergence exhibits very different properties in different equilibria. In the most aggressive equilibrium, the public belief about the medium state converges to a point strictly between 0 and 1, so social learning is incomplete. But in the least coordination equilibrium, the public belief about the medium state converges to 1, so social learning is complete. As a result, in the least coordination equilibrium, the strategy profile in the limit purifies the mixed strategy Nash equilibrium in the complete information normal form game when the status quo is medium. Finally, I show that social welfare in the dynamic regime change game depends on the true state. In the weak status quo, social learning delays the regime change and thus leads to inefficiency. But as the discount factor is close to 1, such inefficiency disappears. In the medium status quo, the interaction between social learning and coordination results in inefficient herds; so when the discount factor is sufficiently large, social welfare is strictly lower. But in the strong status quo, because social learning prevents infinitely many attacks, it leads to higher social welfare.

The dynamic regime change game without individual learning has been applied to numerous economic and social issues. For instance, the investment problem. A lot of projects require synchronous coordination, and no investor can individually learn the true economic fundamentals of the project. The analysis in this paper suggests that if the project is not invested for a long time, it may never be invested, even though it is profitable and investors' research data could be arbitrarily informative. The model can also be used to explain herding phenomena in currency attacks, bank runs, debt crises, and political revolutions, in which synchronous coordination plays a critical role and players' private signals are assumed to be rich and powerful.

Appendix A Omitted Proofs

This section includes proofs of Propositions and Lemmas, which 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 notation. I first show that if a Bayesian Nash equilibrium exists, it is in 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 $\lim\limits_{x_i \to -\infty} E_{x_j} u_i(1, x_i, s_j) = 1 - c > 0$ (dominant region of attacking) and $\lim\limits_{x_i \to +\infty} E_{x_j} u_i(1, x_i, s_j) = -c < 0$ (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 cutoff strategies. So 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 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 $\hat{x}_2 = \hat{\hat{x}}_1 > \hat{x}_1$, a contradiction.

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

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

Because $G(x, \mu)$ is continuous in x, the dominant region of attacking and 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 a μ with $\mu(\theta) > 0$ for all $\theta \in \Theta$, such that multiple equilibria exist in this static regime change game. To show this claim, I just need to show that there exists a μ such that there are more than one solution to $G(x, \mu) = 0$. Note that $\lim\limits_{\mu(m) \to 1} G(\frac{w+m}{2}, \mu) = \Phi(\frac{\sqrt{\beta}}{2}(w - m)) - c < 0$ (by $c > \frac{1}{2}$) and $\lim\limits_{\mu(m) \to 1} G(\frac{m+s}{2}; \mu) = \Phi(\frac{\sqrt{\beta}}{2}(s - m)) - c > 0$ (by Assumption 1). Therefore, 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 $(-\infty, \frac{w+m}{2})$, one in $(\frac{w+m}{2}, \frac{m+s}{2})$, and one in $(\frac{m+s}{2},$

$+\infty$). Furthermore, if we fix μ such that $\frac{\mu(w)}{\mu(w)+\mu(m)} < c < \frac{\mu(m)}{\mu(m)+\mu(s)}$, then for any β, there are three Bayesian Nash equilibria in the static game.

Q.E.D.

Proof of Proposition 2:

This Proposition is also about the static game, so 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. So similar to the proof of Proposition 1, if an equilibrium with attacks exists, it is symmetric and in cutoff strategies. Denote a symmetric cutoff strategy profile by (x, x), then (x^*, x^*) is a non-trivial equilibrium of the regime change game if and only if $G(x^*, \mu) = 0$. Therefore, conditions for the existence of a nontrivial equilibrium are equivalent to those for the existence of a solution to $G(x, \mu) = 0$. Note $G(x, \mu)$ can be equivalently written as $G(x, \mu) = \rho(m|x) g(x, \mu)$, where

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

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 where $\mu(m)$ is close to 1. Since $\Phi(\frac{\sqrt{\beta}}{2}(s-m)) > c$, $g(\frac{s+m}{2}; \mu) > 0$. Note that for all $\mu(m) \in (0, 1)$, $g(m, \mu) < 0$ and $\lim_{x \to +\infty} g(x, \mu) < 0$, so by continuity of $g(x; \mu)$ in x, there exist $\hat{x} \in (m, \frac{s+m}{2})$ and $\hat{\hat{x}} \in (\frac{s+m}{2}, +\infty)$ 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 $\mu(m)$ is close to 0. The last term of $g(x; \mu)$ is very negative for any x larger than m, so $g(x, \mu) < 0$ for all $x > m$. Combined with the fact that $g(x, \mu) < 0$ for all $x \le m$, there exists $\underline{\mu}(m) \in (0, 1)$ 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 $\tilde{\mu}(m) \in (\bar{\mu}(m), \underline{\mu}(m))$, such that for all $\mu(m) \in (0, \tilde{\mu}(m))$, there is no solution to $g(x, \mu) = 0$; and for all $\mu(m) \in (\tilde{\mu}(m), 1)$, 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 $\exists x'' \in (m, +\infty)$ such that $g(x'', \mu'') = 0$ (because all $x \leq m$ 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 $\lim_{x \to +\infty} g(x, \mu') < 0$, 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 $\tilde{\mu}(m) = \inf\{\mu(m) \in (0,1) : \exists x \in \mathbb{R}$ such that $g(x, \mu) = 0\} = \sup\{\mu(m) \in (0,1) : g(x, \mu) < 0 \quad \forall x \in \mathbb{R}\}$ (since for a given μ, $g(x, \mu)$ either has a solution or does not have a solution). Obviously, $\tilde{\mu}(m) \in (\bar{\mu}(m), \underline{\mu}(m))$.

For all $\mu(m) \in (\tilde{\mu}(m), 1)$, 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 $\mu(m) \in (\tilde{\mu}(m), 1)$, 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 $x \in (\bar{x}(\tilde{\mu}(m)), +\infty)$, $g(x, \tilde{\mu}(m)) < 0$ (dominant region of not attacking), where $\bar{x}(\mu(m)) = \inf\{x \in \mathbb{R} : E_{x_j} u_i(1, x_i, s_j) < 0$ 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, $g(x, \tilde{\mu}(m)) < 0$ for all $x > \bar{x}(\bar{\mu}(m))$, because $\tilde{\mu}(m) < \bar{\mu}(m)$. Since $c > \frac{1}{2}$, for any μ, $g(x, \mu) < 0$ for all $x < m$. Now consider the compact set $[m, \bar{x}(\bar{\mu}(m))]$. From 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 $\mu^k(m) \in (\tilde{\mu}(m), \bar{\mu}(m))$, $\mu^k(m) > \mu^{k+1}(m)$ and $\mu^k(m) \to \tilde{\mu}(m)$. Since $g(x, \mu)$ is continuous in $\mu(m)$ for any $x \in [m, \bar{x}(\bar{\mu}(m))]$, $\lim_{k \to +\infty} g(x, \mu^k) = g(x, \tilde{\mu})$. 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

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

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

Now suppose $x' \neq \tilde{x}$ 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 $g(x'', \tilde{\mu}) > 0$. Then because $g(x'', \mu)$ is continuous in $\mu(m)$, fix any $\epsilon \in (0, \frac{g(x'',\tilde{\mu})}{2})$, there exists $\gamma > 0$ such that for all $\mu'(m) \in (\tilde{\mu}(m) - \gamma, \tilde{\mu}(m))$, $g(x'', \mu) > g(x'', \tilde{\mu}) - \epsilon > 0$. So there exists $x''' \in (m, x'')$ such that $g(x'''; \mu) = 0$. This contradicts the definition of $\tilde{\mu}(m)$. Therefore, there exists a unique $\tilde{x} \in \mathbb{R}$, such that $g(\tilde{x}, \tilde{\mu}(m)) = 0$. ■

Q.E.D.

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

Since 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 $y = \frac{1}{2}$,

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

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 $x \in (m, \frac{m+s}{2})$,

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

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

Q.E.D.

Proof of Proposition 6:

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 $(\tilde{\mu}(m), 1)$. Therefore, I only need to show that this fixed point is stable, which is equivalent to show that the slope $0 < d\mu_{t+1}(m)/d\mu_t(m) < 1$. Since

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

$\left.\frac{\partial \mu_{t+1}(m)}{\partial \mu_t(m)}\right|_{\tilde{\mu},\tilde{x}} = 1$, and $-1 < \left.\frac{\partial \mu_{t+1}(m)}{\partial x_t^*}\right|_{\tilde{\mu},\tilde{x}} \left.\frac{\partial x_t^*}{\partial \mu_t(m)}\right|_{\tilde{\mu},\tilde{x}} < 0$ 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. Together with the fact that $x_t^* > m$, $\mu_t(m) \to 1$. Since $g(x_t^*, \mu_t) = 0$ for all t, $\Phi[\sqrt{\beta}(x_t^* - m)] \to c$.

Q.E.D.

Proof of Proposition 7:

Let $\mathcal{W}^L$ be the social welfare with social learning and $\mathcal{W}^N$ be the social welfare without social learning. Part 1 is due to the decreasing probability of attacking in both E^{max} and E^{min}. When $\theta = w$, let κ be the regime change time. With social learning, $\mathbb{P}_L(\kappa \geq t|\theta = w) = 1 - \sum_{\tau=1}^{t-1} \mathbb{P}(\kappa = \tau|\theta = w)$. Define p_t^L to be the probability that an attack happens in period t conditional on no attack before with social learning when $\theta = w$, then $\mathbb{P}_L(\kappa = 1|\theta = w) = p_1^L$. (So the probability that an attack happens in period t conditional on no attack before without social learning is $p_t^N = p_t^L$ for all t.) So, by induction, $\mathbb{P}_L(\kappa \geq t|\theta = w) = \prod_{\tau=0}^{t-1}(1 - p_\tau)$, where $p_0^L = p_0^N \equiv 0$. By the same way, $\mathbb{P}_N(\kappa \geq t|\theta = w) = (1 - p_1^L)^{t-1}$. Because $\{p_t^L\}_t$ is a decreasing sequence, $\mathbb{P}_L(\kappa \geq t|\theta = w) \leq \mathbb{P}_N(\kappa \geq t|\theta = w)$ for all $t = 1, 2, \ldots$. Therefore, the cumulative distribution function of κ without social learning first order stochastic dominates that with social learning, which implies that the expected regime change time is longer with social learning. Let $\mathcal{V}_t$ be the discounted value conditional that the regime changes in period t, then $\mathcal{V}_t > \mathcal{V}_\tau$ if $t < \tau$, given any $\delta \in (0, 1)$, so social learning leads to inefficiency when $\theta = w$ for any $\delta \in (0, 1)$. However, for any $\epsilon > 0$, there is a T such that $\left|\sum_{t=T+1}^{\infty} \mathbb{P}_L(\kappa = t|\theta = w)\mathcal{V}_t - \sum_{t=T+1}^{\infty} \mathbb{P}_N(\kappa = t|\theta = w)\mathcal{V}_t\right| < \frac{\epsilon}{2}$ for all $\delta \in (0, 1)$. Fix this T, as $\delta \to 1$, $\left|\sum_{t=1}^{T} \mathbb{P}_L(\kappa = t|\theta = w)\mathcal{V}_t - \sum_{t=1}^{T} \mathbb{P}_N(\kappa = t|\theta = w)\mathcal{V}_t\right| < \frac{\epsilon}{2}$. Therefore, for any $\epsilon > 0$, there is a $\delta_w \in (0,1)$ such that for all $\delta \in (\delta_w, 1)$, $\left|\sum\limits_{t=1}^{\infty} \mathbb{P}_L(\kappa = t|\theta = w)\mathcal{V}_t - \sum\limits_{t=1}^{\infty} \mathbb{P}_N(\kappa = t|\theta = w)\mathcal{V}_t\right| < \epsilon$. That is, as the discount factor goes to 1, the inefficiency due to the delay of the regime change caused by social learning disappears.

Part 2 is a consequence of Corollary 1. On one hand, because $\mathbb{P}_N(\text{regime changes}|\theta = m) = 1$, as δ goes to 1, $\mathcal{W}^N$ converges to $2(1-c)$. On the other hand, Corollary 1 implies that $\mathbb{P}_L(\text{regime changes}|\theta = m) < 1$, which in turns implies that $\mathcal{W}^L$ is strictly less than $2(1-c)$. Note that infinitely many attacks are prevented with or without social learning and that the discounted value of the cost from finitely many attacks goes to 0 as δ goes to 1. Therefore, there is a $\delta_m \in (0,1)$ such that for all $\delta \in (\delta_m, 1)$, social learning leads to a lower social welfare value.

For Part 3, while with social learning, $\mathbb{P}_L(\text{attack, i.o.}|\theta = s) = 0$, since either the probability of attacking is constant at 0 from some finite period onward or the probability of two attacks in every period is bounded away from 0, without social learning, $\mathbb{P}_N(\text{attack, i.o.}|\theta = s) > 0$. Because the strong status quo won't fall, the fewer attacks, the higher the social welfare. In particular,

$$\begin{aligned}\mathcal{W}^L &= (1-\delta)\sum_{t=1}^{\infty} \mathbb{P}_L(\text{no attack after period } t|\theta = s)V_t^L \\ \mathcal{W}^N &= (1-\delta)\sum_{t=1}^{\infty} \mathbb{P}_N(\text{no attack after period } t|\theta = s)V_t^N\end{aligned}$$

where V_t^L and V_t^N are the expected discounted social welfare (conditional on the event that no attack occurs after period t) with and without social learning respectively. Note for any t, both V_t^L and V_t^N are finite. Because $\mathbb{P}_L(\text{attack, i.o.}|\theta = s) = 0$, for any $\epsilon > 0$, there is a T such that $\left|\sum\limits_{t=T}^{\infty} \mathbb{P}_L(\text{no attack after period } t|\theta = s)V_t^L\right| < \epsilon$. Because $\mathbb{P}_N(\text{attack, i.o.}|\theta = s) > 0$, for any T', there is a $\epsilon > 0$ such that $\left|\sum\limits_{t=T}^{\infty} \mathbb{P}_N(\text{no attack after period } t|\theta = s)V_t^N\right| > \epsilon'$. Fix such ϵ' and T',

$$\mathcal{W}^L > (1-\delta)\sum_{t=1}^{T'} \mathbb{P}_L(\text{no attack after period } t|\theta = s)V_t^L - \epsilon',$$

while

$$\mathcal{W}^N < (1-\delta)\sum_{t=1}^{T'} \mathbb{P}_N(\text{no attack after period } t|\theta = s)V_t^N - \epsilon'.$$

Therefore, there is a δ_s such that for all $\delta \in (\delta_s, 1)$, $\mathcal{W}^L > \mathcal{W}^N$.

Q.E.D.

Appendix B Useful Lemmas

In this section, I state and prove two lemmas, which are used in section 4.

Lemma 5 *Suppose $\theta = w$ has been ruled out. Given large β, $\tilde{\mu}(m)$ is decreasing in β. As $\beta \to +\infty$, $\tilde{\mu}(m)$ converges to* 0.

Proof.

Recall that $\tilde{\mu}(m)$ is the belief about $\theta = m$, at which 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 $g(x, \tilde{\mu}) = [\Phi(\sqrt{\beta}(x - m)) - c] - c(\frac{1}{\tilde{\mu}(m)} - 1)\exp(\frac{\beta}{2}(s - m)(2x - s - m))$. Since $\tilde{x}$ is also the unique solution to $g(x, \tilde{\mu}) = 0$, $g'(\tilde{x}, \tilde{\mu}) = 0$. That is,

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

Comparative static analysis shows that, for large β, $\tilde{\mu}(m)$ 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 $\beta \to +\infty$, $\exp(\frac{\beta}{2}(s-m)(2\tilde{x}-s-m))$ goes to $\exp(-\frac{(s-m)^2}{2}\beta + \Phi^{-1}(c)(s-m)\sqrt{\beta})$. Suppose $\tilde{\mu}(m)$ is bounded away from 0 as β goes to $+\infty$, then $(\frac{1}{\tilde{\mu}(m)} - 1)\exp(-\frac{(s-m)^2}{2}\beta + \Phi^{-1}(c)(s-m)\sqrt{\beta})$ and $\sqrt{\beta}(\frac{1}{\tilde{\mu}(m)} - 1)\exp(-\frac{(s-m)^2}{2}\beta + \Phi^{-1}(c)(s - m)\sqrt{\beta})$ both go to 0. So $g'(\tilde{x}, \tilde{\mu}) > 0$, which leads to the contradiction. As a result, as $\beta \to +\infty$, $\tilde{\mu}(m) \to 0$. ■

Lemma 6 *In an N-players $N+1$-states 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.

I show this lemma with the general private signal structure mentioned in footnote 4. Let $L_k(x) = \frac{f(x|m_k)}{f(x|m_1)}$ to 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 the regime change is $Z_k = \Pr(\text{there are at least } k-1 \text{ players choosing to attack besides player } i|m_k)$. Note Z_k is independent of x_i, and $Z_{k+1} \leq Z_k \leq 1$ for all $k = 1, 2, \ldots, N$.

Then the interim payoff of player i when she observes private signal x_i and chooses to attack is:

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

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,

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

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 $x_i > \bar{x}_i$.

■

References

[1] Angeletos, George-Marios, Christian Hellwig and Alessandro Pavan (2007): “Dynamic Global Games of Regime Change: Learning, Multiplicity, and the Timing of Attacks,” *Econometrica*, 75(3): 711-56.

[2] Angeletos, George-Marios, and Alessandro Pavan (2007): “Efficient Use of Information And Social Value of Information,” *Econometrica*, 75(4): 1103-42.

[3] Banerjee, Abhijit V. (1992): “A simple Model of Herd Behavior,” *The Quarterly Journal of Economics*, 107(3):797-817.

[4] Bikhchandani, Sushil, David Hirshleifer and Ivo Welch (1992): “A Theory of Fads, Fashion, Custom, and Culture Change as Informational Cascades,” *Journal of Political Economy*, 100(5): 992-1026.

[5] Bikhchandani, Sushil, David Hirshleifer and Ivo Welch (1998): “Learning from the Behavior of Others: Conformity, Fads, and Informational Cascades,” *Journal of Economic Perspectives*, 12(3): 151-70.

[6] Carlsson, Hans and Eric Van Damme (1993): “Global Games and Equilibrium Selection,” *Econometrica*, 61: 989-1018.

[7] Chamley, Christophe (2004): *Rational Herds*, Cambridge University Press, Cambridge.

[8] Choi, Jay Pil (1997): “Herd Behavior, the “Penguin Effect”, and the Suppression of Informational Diffusion: An Analysis of Informational Externalities and Payoff Interdependency,” *Rand Journal of Economics*, 28(3): 407-425.

[9] Cripps, Martin, Jeffrey Ely, George Mailath and Larry Samuelson (2008): “Common Learning,” *Econometrica*, 76(4): 909-33.

[10] Dasgupta, Amil (2000): “Social Learning and Payoff Complementarities,” available at http://personal.lse.ac.uk/DASGUPT2/research.html.

[11] Dasgupta, Amil (2007): “Coordination and Delay in Global Games,” *Journal of Economic Theory*, 134: 195-225.

[12] Edmond, Chris (2008): “Information Manipulation, Coordination and Regime Change,” Working Paper, Stern School of Business, New York University.

[13] Giannitsarou, Chryssi and Flavio Toxvaerd (2007): “Recursive Global Games,” CEPR Discussion Paper No. DP6470.

[14] Goldstein, Itay and Ady Pauzner (2005): "Demand Deposit Contracts and the Probability of Bank Runs," *Journal of Finance*, 60: 1293-1328.

[15] Heidhues, Paul and Nicolas Melissas (2006): "Equilibria in a Dynamic Global Game: The Role of Cohort Effects," *Economic Theory*, 28, 531-57

[16] Hirshleifer, Jack (1971): "The Private and Social Value of Information and the Reward to Inventive Activity," *American Economic Review*, 61(4): 561-74.

[17] Lee, In Ho (1993): "On the Convergence of Informational Cascades," *Journal of Economic Theory*, 61(2): 396-411.

[18] Monderer, Dov and Dov Samet (1989): "Approximating Common Knowledge with Common Beliefs," *Games and Economic Behavior*, 1: 170-90.

[19] Morris, Stephen and Hyun Song Shin (1998): "Unique Equilibrium in a Model of Self-Fulfilling Currency Attacks," *American Economic Review*, 88:587-97.

[20] Morris, Stephen and Hyun Song Shin (2002): "Social Value of Public Information," *American Economic Review*, 92:1521-34.

[21] Morris, Stephen and Hyun Song Shin (2003): "Global Games: Theory and Applications," in M. Dewatripont, L. Hansen and S. Turnovsky (eds.) *Advances in Economics and Econometrics: Theory and Applications, Eighth World Congress*, vol. 1.

[22] Morris, Stephen and Hyun Song Shin (2004): "Coordination Risk and the Price of Debt," *European Economic Review*, 48: 133-53.

[23] Morris, Stephen and Hyun Song Shin (2007): "Common Belief Foundations of Global Games," working paper, Princeton University.

[24] Smith, Lones and Peter Sørensen (2000): "Pathological Outcomes of Observational Learning," *Econometrica*, 68(2): 371-98.