Social Learning in Regime Change Games

Chong Huang*

November 17, 2010

*Contact information: hchong@sas.upenn.edu. I am grateful to George Mailath for invaluable guidance in this project and numerous suggestions. I thank Andrew Postlewaite for great comments that improved the paper. I also thank David Dillenberger, Hanming Fang, Itay Goldstein, Qingmin Liu, Steven Matthews, Muhamet Yildiz and micro lunch seminar participants at UPenn for helpful comments. All errors are mine.

Social Learning in Regime Change Games

Abstract

This paper studies social learning effects in dynamic regime change games with a finite number of short-lived players in each period. These games have been usually applied to currency attacks by hedge funds, investments in emerging firms by venture capitalists, and revolutions against dictators. In my model, the state of the status quo is fixed but unobservable to players. Since each short-lived player can observe only one signal about the true state, no individual can individually learn the true state of the status quo. However, I allow players to observe previous play, so the true state may be socially learned. I describe the equilibrium dynamics of attacking and relate the state of the status quo to the likelihood of the regime's eventual fate. This model, in which perfect individual learning is impossible, yields equilibrium properties that differ from earlier results of models in the literature, where perfect individual learning is allowed. First, players may give up attacking even though they don't learn the true state, because extremely informative signals may be ignored when cooperation is required. Second, fundamentals may not determine the eventual fate of the regime, as signals from early periods are important. Third, social learning may lead to either efficiency or inefficiency depending on the state of the status quo.

JEL Classification: C72, C73, D83

Key words: Dynamic global game, social learning, herding, currency attack

1 Introduction

Speculative currency attacks have happened in many countries. Some attacks succeeded immediately; for example, the Thai baht was devaluated after only one attack in May 1997. Some attacks were successful only after several attempts. For example, following an attack against the Russian ruble in November 1997, a further attack in April 1998 triggered a ruble devaluation. And some attacks were unsuccessful, ending with no new attacks. This happened in Hong Kong. During the East Asian financial crisis, the Hong Kong government defended speculative attacks in October 1997, January 1998, and June 1998. After a further attack failed in August 1998, no speculator initiated a new war against the Hong Kong dollar. Did all speculators believe that the Hong Kong dollar was strong enough to survive speculative attacks? Or did all individual speculators believe that they could succeed if all of them shorted the Hong Kong dollar simultaneously but this coordination was very difficult? While most currencies have been attacked at some point, the Dutch guilder has never been attacked. Will this fact suggest to future speculators to never attack the Dutch guilder, even if it was thought to be weak? Similar questions arise when venture capitalists are considering investing in an emerging firm and when potential revolutionaries are planning to fight against a dictator.

These questions share some common features. In particular, past people's behavior affect current agents' decisions. Because an agent's private information about the fundamentals affects that agent's optimal decision, observed past behavior is informative about the fundamentals. As a result, after observing previous plays, current period players update their beliefs about the fundamentals. This process constitutes *social learning*. For example, failed attacks against the Hong Kong dollar in 1997 and 1998 provide good lessons to subsequent speculators. On the one hand, past attacks imply that previous private signals suggested to people in previous times that they should attack; on the other hand, that attacks failed suggests that the Hong Kong dollar was not weak. These inferences affect subsequent speculators' decisions, which in turn will impact future speculators' decisions. Besides learning from public histories, players can learn the fundamentals by collecting private signals. This is usually defined as *individual learning*. However, in an economy consisting of a finite number of agents, individual learning may significantly complicate the analysis of the social learning effects. Therefore, I abstract away individual learning and focus on the social learning effects when addressing the questions above.

This paper studies the social learning effects in the situations above, by analyzing a dynamic regime change game with a finite number of players in each period and synchronous coordination requirement at some states. In each period, there are two new short-lived (one-period-lived) players, each of whom observes previous public plays and one piece of private information about the status quo. Based on these histories, they update their beliefs about the true state of the status quo, which is unknown and 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. 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 regardless of the regime change outcome. The true state of the strength of the status quo is drawn at the beginning of the game from a set consisting of three elements: weak, medium, and strong. If the status quo is weak, one attack changes the regime; for the medium status quo, synchronous coordination is required to trigger the regime change; if the status quo is strong, the status quo can never be beaten. The game continues as long as the status quo is in place.

In my model, given the prior belief about the fundamentals, there are always extremely informative private signals, which make a player have an arbitrarily high posterior belief about some state of the status quo. However, with social learning, players' prior beliefs evolve over time, which generates the dynamics of attacking. When current period players observe that no attack has yet occurred, they infer that previous private signals suggested not to attack, so the true state of the status quo is more likely to be strong. Social learning, in this case, makes current period players pessimistic about triggering the regime change, so it reduces current period players' incentives to attack. The longer the time without any attack, the more pessimistic the players are, and the smaller the ex-ante probability of attacking is. In another scenario, players rule out the possibility of the weak status quo by observing a failed attack by one player in some period before (but there was no period with failed attacks by two players). Now the regime changes only if the status quo is medium and in some period players coordinate synchronously. Suppose no attack happens after the last failed attack. Similar to the first case, social learning makes the ex-ante probability of attacking smaller over time. As a result, though players may be convinced by their private signals that the status quo is medium, they think coordination is very difficult. Hence, players give up attacking, ignoring their private signals.

As the driving force of the dynamics of attacking, social learning plays a critical role in determining eventual regime change outcomes. The fate of a weak regime is an example. After a very long time without attacks since the beginning of the game, players' prior beliefs about the not strong status quo are arbitrarily low. As a result, the ex-ante probability of attacking is arbitrarily small. Therefore, ex-ante, it seems that the weak regime may not change. However, while the ex-ante probability of attacking keeps decreasing, players learn less and less from the no attack history. Such a social learning process prevents the ex-ante probability of attacking from dropping to 0 too fast and thus guarantees the change of the weak regime. Social learning's effect on regime change outcomes is also illustrated when the true state of the status quo is medium. If players cannot observe past plays (so that social learning is impossible), the medium status quo is abandoned eventually. With social learning, a failed attack by one player along with sufficiently long but finite periods of no attack leads players to give up attacking, no matter what their private signals are. Since such an event happens with a positive probability, the medium status quo may not be abandoned.

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

The benchmark dynamic regime change game rules out the possibility of perfect individual learning and focuses on a two-player three-state case. I extend the benchmark model in two directions. In the first extension, players' private signals are increasingly precise and this precision becomes perfect. That is, I allow perfect individual learning. In this extended model, the medium regime will be abandoned almost surely, and I get transitions similar to those in Angeletos, Hellwig and Pavan (2007): the economy transits back and forth between "tranquility" phases (no attack in any equilibrium) and "distress" phases (attack with positive probability in non-trivial equilibria). 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 regime in the benchmark model, while the $(N+1)$th state is like the strong regime in the benchmark 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 the benchmark model.

Abstracting away the dynamic feature, economists have been extensively applying regime change games to currency attacks (Morris and Shin, 2003), bank runs (Goldstein and Pauzner, 2005), debt crises (Morris and Shin, 2004), and political change (Edmond, 2007). These static regime change games follow the static global game literature (Carlsson and Van Damme, 1993), and they are solvable by iterated elimination of strictly dominated strategies (see Morris and Shin, 2003). Hence, fix a prior belief over the parameter space. As the noise of private signals goes to 0, there is a unique Bayesian Nash equilibrium in static global games.

In a recent paper, Angeletos, Hellwig and Pavan (2007) incorporate the learning process to regime change games in a dynamic model.[1] They consider a continuum of long-lived agents, each receiving a piece of private information about the true state in each period. Therefore, an individual eventually learns the true state from her own signals. This individual learning process is the driving force of the dynamics of attacking in their model. But in Angeletos, Hellwig and Pavan (2007), previous plays are not observable. Because the economy consists of a continuum of agents, if previous plays are observable, at the beginning of the second period, players learn perfectly what the true state is.

Different from the static global game literature, Angeletos, Hellwig and Pavan (2007) show the existence of multiple equilibria in a dynamic global game. Such multiplicity originates in the coordination requirement and the private learning process. Since attacking is not a dominant action for all private signals since the second period, an equilibrium in which no attack occurs after the first period always exists. And when the fundamentals that can be beaten are still in the support of players' posterior beliefs, private learning generates subsequent attacks consistent with some equilibrium by making players' private signals highly correlated. Multiple equilibria also exist in my model, but they emerge from two distinct sources. First, similar to Angeletos, Hellwig and Pavan (2007), when the weak regime is ruled out, the coordination requirement makes no more attacks always consistent with an equilibrium; then when players' prior beliefs about the medium regime are high enough, possible subsequent attacks are also consistent with some equilibrium. Second, different from Angeletos, Hellwig and Pavan (2007), when period t players observe that no attack has yet occurred (this is on the path of play in any equilibrium), period t players are playing a static global game with both attacking and not attacking as dominant actions for some private signals. However, since the precision of players' private signals is fixed and players' prior beliefs (after observing previous plays) are evolving over time through social learning, there

[1] Dynamic regime change games are studied as examples of dynamic global games. Other papers contributing to this literature are Dasgupta (2006), Dasgupta, Steiner and Stewart (2009), Giannitsarou and Toxvaerd (2009), and Heidhues and Melissas (2006). Chamley (1999) analyzes a regime switch game, in which a continuum of short-lived players in each period learn the distribution of the private cost in their cohort from previous aggregate actions.

are prior distributions on the fundamentals such that period t players put a sufficiently high probability on the medium regime. As a result, multiple strategies of period t players are all consistent with some equilibrium. This is different from static global games, in which prior beliefs are fixed and the precision of private signals is arbitrarily large.

Social learning models, initialized by Banerjee (1992) and Bikhchandani, Hirshleifer and Welch (1992) and generalized by Smith and Sorenson (2000), have been extensively studied in the herding literature (see also Chamley (2002)). However, there is no strategic uncertainty in these herding models. In a dynamic regime change game, a player's payoff is increasing in her opponent's probability of attack. That is, there is a strategic complementarity in each period of the dynamic regime change game. As shown in the analysis of this paper, this strategic complementarity leads players to join the herd (not to attack ignoring private signals), even though they are convinced by their extremely informative private signals that the status quo can be beaten.

The rest of this paper is organized as follows. Section 2 describes a static regime change game and analyzes it with both interior prior beliefs and prior beliefs ruling out the weak status quo. In section 3, I introduce a dynamic regime change game and provide an algorithm to characterize all equilibria. In section 4, I study both the short-run and the long-run dynamics of attacking, investigate the eventual fate of the status quo conditional on its strength, and analyze the social learning effects in this game. Section 5 is devoted to two extensions of the benchmark model. Section 6 concludes. All proofs are presented in the Appendix.

2 A Two-Player Static Regime Change Game

In this section, I describe a static regime change game. Since the results of this static game will be applied in the dynamic model, I analyze them in detail.

2.1 The Model

Two players $i \in \{1,2\}$ play the static regime change game. Player i can choose one of two actions $a^i \in \{0,1\}$, where $a^i = 0$ is "not attack" while $a^i = 1$ means "attack." Player i's ex-post payoff depends on both his choice and the outcome of the regime change. Let the binary variable R denote the outcome ($R = 1$ means the regime changes, and $R = 0$ means the regime does not change). Player i's ex-post payoff is $u^i = a^i(R - c)$, where $c \in (\frac{1}{2}, 1)$ is the cost of attacking.

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 $w < m < s$.[2] Whether the regime changes or not depends on both the strength of the status quo θ and the number of attacks $a^1 + a^2$. The following table summarizes these outcomes:

Number of Attacks	$\theta = w$	$\theta = m$	$\theta = s$
0 attack	Fail	Fail	Fail
1 attack	Succeed	Fail	Fail
2 attacks	Succeed	Succeed	Fail

The strength of the status quo θ is unobservable to all players, and players share the common prior belief $\mu \in \Delta(\Theta)$. Before choosing to attack or not to attack, player i receives a private signal $x^i = \theta + \xi^i$, where $\xi^i \sim \mathcal{N}(0, \frac{1}{\beta})$ is independent of θ and independent across players.[3] Therefore, players' signals are independent conditional on θ.

Player i's strategy is a mapping $s^i : \mathbb{R} \to \Delta(\{0,1\})$. A *Bayesian Nash Equilibrium* is a strategy profile $(\hat{s}^1, \hat{s}^2)$, such that given $\hat{s}^j$ and any signal x^i, $u^i(\hat{s}^i, \hat{s}^j) \geq u^i(a^i, \hat{s}^j)$ for all $a^i \in \{0,1\}$. In the rest of this section, I analyze the properties of Bayesian Nash Equilibria when (i) $\mu(\theta) > 0, \forall \theta \in \Theta$ and when (ii) $\mu(w) = 0$,while $\mu(m) > 0$ and $\mu(s) > 0$. In case (i) the static regime change game is a static global game, and in case (ii) it is a coordination

[2] w,m and s denote "weak," "medium" and "strong," respectively.

[3] The additive structure of the signal and the normality assumption of the noisy term are convenient for computing equilibria and comparing equilibrium properties with previous works. In fact, the distribution of the signal can be fairly general, and the only assumptions I have to make are the following: (1) conditional on θ, the two players' private signals are independent and identically distributed; (2) the support of x^i is $\mathbb{R}$, and the conditional pdf $f(x|\theta)$ of the signal is strictly positive for all $x \in \mathbb{R}$ and all $\theta \in \Theta$; (3) unbounded likelihood ratio: $\lim_{x\to-\infty} f(x|\theta)/f(x|\theta') = +\infty$ and $\lim_{x\to+\infty} f(x|\theta)/f(x|\theta') = 0$, whenever $\theta < \theta'$; and (4) monotone likelihood ratio: if $\theta < \theta'$, $f(x|\theta)/f(x|\theta')$ is strictly decreasing in x.

game.

2.2 Interior Prior Beliefs

In this subsection, I focus on the interior prior beliefs case of the static regime change game. Let $\rho(\cdot|x^i)$ denote player i's posterior belief over Θ after receiving signal x^i. Then from Bayes' rule, the posterior belief about θ is:

$$\rho(\theta|x^i) = \frac{\mu(\theta)\phi(\sqrt{\beta}(x^i - \theta))}{\sum\limits_{\theta' \in \Theta} \mu(\theta')\phi(\sqrt{\beta}(x^i - \theta'))},$$

where $\phi(\cdot)$ is the standard normal pdf. Player i's interim payoff from attacking given signal x^i and player j's strategy s^j is:

$$\begin{aligned}
& E_{x^j} u^i(1, s^j|x^i) \\
= \;& \rho(w|x^i) + \Pr(s^j = 1, m|x^i) - c && (2.1)\\
= \;& \rho(w|x^i) + \rho(m|x^i) \Pr(s^j = 1|m) - c && (2.2)\\
= \;& \frac{\mu(w)\phi(\sqrt{\beta}(x^i - w))}{\sum\limits_{\theta' \in \Theta} \mu(\theta')\phi(\sqrt{\beta}(x^i - \theta'))} + \frac{\mu(m)\phi(\sqrt{\beta}(x^i - m))}{\sum\limits_{\theta' \in \Theta} \mu(\theta')\phi(\sqrt{\beta}(x^i - \theta'))} \Pr(s^j = 1|m) - c, && (2.3)
\end{aligned}$$

where $\Phi(\cdot)$ is the standard normal cdf. (2.1) is equal to (2.2) because players' private signals are independent conditional on θ. Note $\rho(w|x^i) \to 1$ as $x^i \to -\infty$; hence, from the regime change rule, attacking is the dominant action for player i. By the continuity of the interim payoff function, there exists an $\underline{x} \in \mathbb{R}$ such that $E_{x^j} u^i(1, s^j|x^i) > E_{x^j} u^i(0, s^j|x^i), \forall x^i \leq \underline{x}$ and $\forall s^j$. I call the set $(-\infty, \underline{x}]$ the *dominant region of attacking*. Similarly, there is an $\bar{x} \in \mathbb{R}$ such that $E_{x^j} u^i(1, s^j|x^i) < E_{x^j} u^i(0, s^j|x^i), \forall x^i \geq \bar{x}$ and $\forall s^j$, so the set $[\bar{x}, +\infty)$ is called the *dominant region of not attacking.*

Proposition 2.1 *In the static regime change game with $\mu \in int(\Delta(\Theta))$, a Bayesian Nash Equilibrium exists. In any Bayesian Nash Equilibrium, players follow a symmetric cutoff strategy with threshold point $x^* \in \mathbb{R}$:*[4]

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

[4]To simplify notation, I denote the strategy not attacking for all signals by $x^* = -\infty$ and the strategy attacking for all signals by $x^* = +\infty$.

Remark 2.1 In the global game literature, a unique equilibrium is usually guaranteed by conditions on the precision of the signal, β. For any fixed prior belief $\mu \in int(\Delta(\Theta))$, as β becomes very large, the uniqueness of the static regime change game can be established by interim iterated elimination of strictly dominated strategies as in Morris and Shin (2003). However, for any fixed β, there is a prior belief $\mu \in int(\Delta(\Theta))$ such that multiple equilibria exist in the static global game, which is discussed in detail in the proof of Proposition 2.1 in the appendix.

♦

Corollary 2.1 *In the static regime change game, when $\mu \in int(\Delta(\Theta))$, the equilibrium ex-ante probability of attacking is positive but less than one.*

2.3 When the Weak State Is Impossible

In this subsection, I analyze the static regime change game when the prior belief is $\mu(w) = 0$ ($\mu(m) > 0$ and $\mu(s) > 0$). Because the weak state is impossible, a dominant region of attacking no longer exists. Therefore, (x^*, x^*)=$(-\infty, -\infty)$ is an equilibrium. Since players attack the regime with zero probability in this equilibrium, I call this a trivial equilibrium. In the remainder of this subsection, I propose a necessary and sufficient condition for the existence of a non-trivial equilibrium, in which players attack the regime with positive probability.

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' common prior belief about $\theta = m$ is high (players are optimistic), cooperation is possible; conversely, when players' common prior 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 of her opponent observing a signal favoring $\theta = m$ is very low. With the sufficiently informative signals assumption given below (I maintain this assumption in the whole paper), Proposition 2.2 formally shows the above intuition.

Assumption 2.1 *Private signals are sufficiently informative:* $\Phi(\frac{\sqrt{\beta}}{2}(s-m)) > c$.

Proposition 2.2 *In the static regime change game with* $\mu(w) = 0$, $\exists\tilde{\mu}(m) \in (0,1)$ *such that*

1. *If* $\mu(m) < \tilde{\mu}(m)$*, there is no non-trivial equilibrium;*

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

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

Remark 2.2 When $\mu(m) > \tilde{\mu}(m)$, the two non-trivial equilibria have different interpretations. The one with the larger threshold point is the most aggressive equilibrium, because players attack the regime with high probability. The one with the smaller threshold point is the lowest possible coordination equilibrium, because if player j's threshold is lower than it, player i will choose not to attack regardless of her private signal.

♦

3 A Dynamic Regime Change Game

I extend the static game to a dynamic regime change game with two new short-lived players in each period given the status quo is in place. Since any individual can get only one piece of information about the state of the status quo, no one can learn the true state individually. However, I assume previous plays (not previous signals) are perfectly observable, so social learning plays an important role in the dynamic model.

Consider the discrete time model and index periods by $t \in \{1, 2, \cdots\}$. At the beginning of the game, $\theta \in \Theta$ is chosen by nature according to a commonly known distribution $\mu_0 \in int(\Delta(\Theta))$. Once picked, θ is fixed forever. 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 is still in place and $R_t = 1$ means the status quo falls by period t. Assume $R_0 = 0$ and if $R_{t-1} = 0$ and $R_t = 1$ for some t, then

$R_\tau = 1$ for all $\tau > t$. In each period t, there are two new short-lived players and each chooses one action from $\{0,1\}$. The regime change rule is as described in the static regime change game. Period t's players' ex-post payoff is $u_t^i = (1-R_{t-1})a_t^i(R_t - c) + R_{t-1}$. Therefore, if the regime changes in period t, period τ ($\tau > t$) players get payoff 1 for sure and do not have strategic behaviors (essentially, the regime change game ends in period t). Before making their decisions, both players observe the number of attacks in each previous period and a private signal $x_t^i = \theta + \xi_t^i$, where $\xi_t^i \sim \mathcal{N}(0, \frac{1}{\beta})$ is independent of θ and independent across agents and across time.

Previous plays are public histories at the beginning of period t. Denote a typical public history $h^t = (b_1, \cdots, b_{t-1})$, where $b_\tau \in \{0,1,2\}$ is the number of attacks in period τ for all $\tau < t$. Letting H^t be the set of all public histories at the beginning of period t, I define a period t player i's strategy by $s_t^i : H^t \times \mathbb{R} \to \Delta(\{0,1\})$.

Definition 3.1 *An assessment* $\left\{(s_t^i)_{t=1,\cdots}^{i=1,2}, (\mu_t)_{t=1,\cdots}\right\}$ *is an equilibrium if*

1. *For any t, given μ_t, (s_t^1, s_t^2) forms a static game equilibrium;*

2. *μ_t is consistent with $(s_t^i)_{t=1,\cdots}^{i=1,2}$.*

The first part of the definition of an equilibrium is a natural requirement of the assumption that players are short-lived. Since period t players have no intertemporal incentive when making decisions, their strategies need to form a static game equilibrium given μ_t. The consistent belief requirement in the definition implies that because h^t is public, period t players have the same prior belief over Θ at the beginning of period t. So for simplicity, in the definition, I denote the belief system $(\mu_t)_{t=1,\cdots}$.

Recall that Proposition 2.2 implies the existence of multiple equilibria for some prior beliefs in the static regime change game. This in turns implies that multiple equilibria exist in the dynamic regime change game. In particular, as analyzed in section 4, the belief $\mu_t(w) = 0$ and $\mu_t(m) \geq \tilde{\mu}(m)$ appears on the path of play. Then, the multiplicity of equilibria in the static regime change game in period t with prior belief μ_t directly implies the multiplicity of 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).[5] However, in the dynamic regime change game, there is another source for multiplicity. If we fix players' strategies when the prior beliefs are not in the interior of the simplex of the state space, remark 2.1 implies that because the precision of private signals is fixed, the prior belief (in the interior of the simplex of the state space) may lead to multiple equilibria in the static regime change game, which in turn implies multiple equilibria in the dynamic regime change game.

Equilibria of the dynamic regime change game can be characterized by the following algorithm:

1. In period 1, compute all solutions to $G(x; \mu_1) = 0$ (note $\mu_1 = \mu_0$). Pick any solution x_1^* to be the threshold point in the first period;

2. After any history h^t, first employ Bayes' rule (whenever possible) to calculate μ_t from μ_{t-1} and x_{t-1}^* ($t > 1$), then compute all equilibria threshold points in period t with prior belief μ_t;

3. If $\hat{h}^{t-1}$ is on the path of play but $\hat{h}^t = (\hat{h}^{t-1}, b_{t-1})$ is off the path of play, then pick any consistent μ_t and compute all equilibria threshold points in period t with prior belief μ_t.

From the above algorithm, multiple equilibria are characterized. Two classes of equilibria have nice properties and are easy to analyze, which I define as follows:

Definition 3.2 $E^{max} = \{(x^*(h^t), x^*(h^t))_{t=1,\cdots}, (\mu_t^*)_{t=1,\cdots}\}$ *is the equilibrium, in which after any* h^t*, for any* $(x'(h^t), x'(h^t))$ *forming a static equilibrium in period* t *with prior belief* μ_t^**,* $x'(h^t) \leq x^*(h^t)$*.*

Definition 3.3 $E^{min} = \{(x^*(h^t), x^*(h^t))_{t=1,\cdots}, (\mu_t^*)_{t=1,\cdots}\}$ *is the equilibrium, in which after any* h^t*, for any* $(x'(h^t), x'(h^t))$ *forming an equilibrium in period* t *with prior belief* μ_t^* *and* $x'(h^t) \in \mathbb{R}$ *(if exists),* $x^*(h^t) \in \mathbb{R}$ *and* $x^*(h^t) \leq x'(h^t)$*; if there is no such* $x'(h^t)$*,* $x^*(h^t) = -\infty$*.*

[5] In their model, after the first period, players lose the dominant region of attacking, so that they are in a coordination game.

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. I will show that the dynamics of attacking and the social learning processes for the same path of histories are different in these two equilibria.

4 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 at the end of this section, social learning may lead to efficiency or inefficiency, depending on the strength of the status quo.

Two facts, which follow directly from the regime change rule and the assumption that the game ends once the regime changes, should be noted before the analysis begins. First, if one attack is observed in period t and $R_t = 0$, players learn immediately that $\theta \neq w$; thus, given $R_{t-1} = 0$ ($t > 1$), if $\theta = w$, the only history players can observe is the one without any attack ($h^t = (0, \cdots, 0)$). Second, if two attacks are observed in period t and $R_t = 0$, players learn immediately that $\theta = s$. This fact has two implications: (1) Given $R_{t-1} = 0$ ($t > 1$), if $\theta = m$, players can only observe histories with at most one attack in each period (h^t with $b_\tau \leq 1, \forall \tau < t$); (2) Since not attacking is the dominant strategy in state s, if two attacks fail to trigger the regime change in some period t, there are no attacks ever again. In the following, I first study the short-run dynamics of attacking and then investigate the long-run dynamics.

4.1 Short-Run Dynamics

Because $\mu_1 = \mu_0 \in int(\Delta(\Theta))$, by Proposition 2.1, players in period 1 follow a symmetric cutoff strategy with threshold point $x_1^* \in \mathbb{R}$.

First, consider the case that there is no attack in period 1 (so $\hat{h}^2 = 0$). Players in the second period (before receiving private signals) learn that $x_1^1 > x_1^*$ and $x_1^2 > x_1^*$. So the prior belief at the beginning of the second period is:

$$\mu_2[\hat{h}^2](\theta) = \frac{\mu_1(\theta)[\Phi(\sqrt{\beta}(\theta - x_1^*))]^2}{\sum\limits_{\theta' \in \Theta} \mu_1(\theta)[\Phi(\sqrt{\beta}(\theta' - x_1^*))]^2}, \quad \forall \theta \in \Theta.$$

Because $x_1^* \in \mathbb{R}$, $\mu_2(\theta) > 0$ for all $\theta \in \Theta$. Therefore, by Proposition 2.1, players in period 2 follow a symmetric cutoff strategy in the equilibrium, with the threshold point $x_2^*(\hat{h}^2) \in \mathbb{R}$. This is true for any period t: if $R_{t-1} = 0$ and $\hat{h}^t = (0, \cdots, 0)$, no attack in period t will lead to players in period $t+1$ following a symmetric cutoff strategy with threshold point $x_{t+1}^*(\hat{h}^{t+1}) \in \mathbb{R}$. So in both E^{max} and E^{min}, players in period $t+1$ attack the status quo with positive probability. However, this probability is lower than that in period t, because no attack in period t makes period $t+1$ players more pessimistic.

Lemma 4.1 *If there is no attack before period $t+1$, in both E^{max} and E^{min}, players in period $t+1$ attack the status quo with strictly lower probability than period t players.*

Now consider that there is exactly one attack in period 1 ($\bar{h}^2 = 1$). Period 2 players learn two facts: (1) $\theta \neq w$ and (2) $x_1^i > x_1^*$ and $x_1^j \leq x_1^*$. From the first fact, players know that cooperation is needed for the regime change. The latter fact shifts the belief about $\{m, s\}$. Formally, the prior belief at the beginning of period 2 is:

$$\mu_2[\bar{h}^2](\theta) = \begin{cases} 0, & \theta = w; \\ \frac{\mu_1(\theta)\Phi(\sqrt{\beta}(\theta - x_1^*))\Phi(\sqrt{\beta}(x_1^* - \theta))}{\sum\limits_{\theta' \in \{m,s\}} \mu_1(\theta')\Phi(\sqrt{\beta}(\theta' - x_1^*))\Phi(\sqrt{\beta}(x_1^* - \theta'))}, & \theta = m, s. \end{cases}$$

From Proposition 2.2, second period players attack the status quo with positive probability in the equilibrium if and only if $\mu_2[\bar{h}^2](m) \geq \tilde{\mu}(m)$. This is also true for any period t. Fix any non-trivial equilibrium. If there is no attack before period t, players attack the status quo with positive probability (threshold point $x_t^*(\hat{h}^t) \in \mathbb{R}$). Suppose one of the period t players chooses to attack but the other does not. Then period $t+1$ players attack the status quo with positive probability if and only if $\mu_{t+1}[(\hat{h}^t, 1)](m) \geq \tilde{\mu}(m)$. If $\mu_{t+1}[(\hat{h}^t, 1)](m) < \tilde{\mu}(m)$, in any non-trivial equilibrium, the probability of attacking in period $t+1$ is 0, which of course is less than the probability of attacking in period t. Hence, when comparing the probabilities of attacking in period t and in period $t+1$ along this history $h^{t+1} = (0, \cdots, 0, 1)$, it is more interesting to assume $\mu_t[\hat{h}^t]$ and exactly one attack leading to $\mu_{t+1}[(\hat{h}^t, 1)](m) \geq \tilde{\mu}(m)$. One failed attack in period t makes period $t+1$ players understand that they cannot trigger the regime change unilaterally. But whether exactly one attack increases or decreases period $t+1$ players' beliefs about the medium status quo depends on $x_t^*(\hat{h}^t)$. If $x_t^*(\hat{h}^t) < \frac{m+s}{2}$, exactly one attack in period t makes

$$\mu_{t+1}[(\hat{h}^t, 1)](m) > \frac{\mu_t[\hat{h}^t](m)}{\mu_t[\hat{h}^t](m) + \mu_t[\hat{h}^t](s)},$$

so period $t+1$ players put more probability on the medium status quo.[6] Conversely, if $x_t^*(\hat{h}^t) > \frac{m+s}{2}$, exactly one attack in period t makes period $t+1$ players put more probability on the strong status quo.

Lemma 4.2 *Suppose $\mu_t[\hat{h}^t]$ and exactly one attack lead to $\mu_{t+1}[(\hat{h}^t, 1)](m) \geq \tilde{\mu}(m)$. In E^{max}, when $\mu_t[\hat{h}^t](w)$ is close to 1 or $x_t^*(\hat{h}^t) \geq \frac{m+s}{2}$, the probability of attacking in period $t+1$ is lower than that in period t. But there is $\mu_t[\hat{h}^t]$ with $\mu_t[\hat{h}^t](w)$ close to 0 and $x_t^*(\hat{h}^t) < \frac{m+s}{2}$.[7] The probability of attacking in period $t+1$ is higher than that in period t.*

Now consider the third period players' behavior when there is one attack in the first period and no attack in the second period. Obviously, because of the failed attack by one player in the first period, players know $\theta \neq w$. When there are only two possible states, m and s, if second period players attack the status quo with positive probability, no attack in the second period decreases third period players' shared prior belief about $\theta = m$. Hence, in the third period, a non-trivial equilibrium exists if and only if $\mu_3(m) \geq \tilde{\mu}(m)$. This also can be generalized to any period t. Consider the history $h^t = (0, \cdots, 0, 1)$ with no attack in period t (therefore $h^{t+1} = (0, \cdots, 0, 1, 0)$). There is an equilibrium in which the probability of attacking in period $t+1$ is positive if and only if $\mu_{t+1}[(0, \cdots, 0, 1, 0)](m) \geq \tilde{\mu}(m)$.

[6]Note $x < \frac{m+s}{2}$ implies $\Phi(\sqrt{\beta}(m-x))\Phi(\sqrt{\beta}(x-m)) > \Phi(\sqrt{\beta}(s-x))\Phi(\sqrt{\beta}(x-s))$.

[7]This is possible. When $\tilde{x}$ (the unique solution to $H(x; \tilde{\mu})$) is less than $\frac{m+s}{2}$, let $\mu_t[\hat{h}^t](w)$ go to 0, and there exists μ_t such that in all equilibria, $x_t^*(\hat{h}^t) < \frac{m+s}{2}$.

Now suppose period $t+1$ players attack the status quo with positive probability. Is this probability higher or lower than that in period t? In E^{max}, the most aggressive equilibrium, because period $t+1$ players are more pessimistic than period t players, the highest probability of attacking in period $t+1$ is lower than the highest probability of attacking in period t. In contrast, in E^{min}, the lowest possible coordination equilibrium, period t players can cooperate at signals lower than the lowest cooperation point in period $t+1$, because period $t+1$ players are more pessimistic. These intuitions are verified in the following lemma:

Lemma 4.3 *Suppose that $h^t = (0, \cdots, 0, 1)$, there is no attack in period t, and $\mu_{t+1}[(h^t, 0)](m) \geq \tilde{\mu}(m)$. In E^{max}, period $t+1$ players attack the status quo with lower probability than period t players; in E^{min}, period $t+1$ players attack the status quo with higher probability than period t players.*

Finally, consider the history with one attack in both period $t-1$ and period t. Because period $t+1$ players can rule out the case that $\theta = w$, there is an equilibrium in which the probability of attacking in period $t+1$ is strictly positive after the history under consideration if and only if $\mu_{t+1}(m) \geq \tilde{\mu}(m)$. Suppose in this equilibrium period $t+1$ players attack the status quo with positive probability under the history, then will they attack more or less frequently than period t players? Obviously, this question can be answered by comparing $\mu_t(m)$ with $\mu_{t+1}(m)$. From footnote 6 and the proof of Lemma 4.3, the following lemma can easily be shown.

Lemma 4.4 *Suppose $h^t = (0, \cdots, 0, 1)$. There is one attack in period t, and $\mu_{t+1}[(h^t, 1)](m) \geq \tilde{\mu}(m)$. In E^{max}, period $t+1$ players attack the status quo with lower probability than period t players do if and only if $x_t^* > \frac{m+s}{2}$; in E^{min}, period $t+1$ players attack the status quo with higher probability than period t players do if and only if $x_t^* > \frac{m+s}{2}$.*

Before moving on to the long-run dynamics, I summarize the analysis about short-run dynamics in this subsection. First, if there has never been an attack, period $t+1$ players attack the status quo with positive probability lower than that in period t. Second, if there is one attack in some period τ and the status quo is still in place in period t $(t \geq \tau)$, there is an equilibrium, in which players in period $t' > t$ attack the status quo with positive probability, if and only if $\mu_{t'}(m) \geq \tilde{\mu}(m)$. Third, if there is one attack before period t and period $t+1$ players attack the status quo with positive probability after no attack in period t, then the probability of attacking is lower in period $t+1$ than in period t in E^{max}, and it is higher in period $t+1$ than in period t in E^{min}. Finally, if period $t+1$ players attack the status quo with positive probability after one attack in period t, the probability of attacking in period $t+1$ is larger than that in period t if $x_t^* < (m+s)/2$ in E^{max} and the probability of attacking in period $t+1$ is smaller than that in period t if $x_t^* > (m+s)/2$ in E^{max}. The converse claim is true in E^{min}.

4.2 Long-Run Dynamics and Regime Change

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

Along $\hat{h}^\infty$, in both E^{max} and E^{min}, players attack the status quo with positive probability in every period t, because $\mu_t \in int(\Delta(\Theta))$. Since the probability of attacking is positive in any period t, period $t+1$ players learn something about θ from no attack in period t. Lemma 4.1 implies that along $\hat{h}^\infty$, players become pessimistic, so the probability of attacking is strictly decreasing over time in both E^{max} and E^{min}.

Proposition 4.1 *In both E^{max} and E^{min}, along $\hat{h}^\infty$, $\mu_t(w) \to 0$ and $\Pr($attack in period $t|\hat{h}^t) \to 0$.*[8]

When $\theta = w$, one attack in any period t (given $R_{t-1} = 0$) triggers the regime change. So the probability that the status quo does not fall at the beginning of period t (conditional on $\theta = w$) is $\prod_{\tau=1}^{t-1}[1 - \Phi(\sqrt{\beta}(x_\tau^*(\hat{h}^\tau) - w))]^2$. For any finite period t, this probability is positive.

[8]If other equilibria exist, $\Pr(\text{attack in period } t|\hat{h}^t)$ may not be monotonic, but the convergence result is still true.

The asymptotic probability that the status quo does not fall when $\theta = w$ is not clear at first glance, because the probability of attacking is strictly decreasing and converges to 0. Therefore, the asymptotic probability of the regime change when $\theta = w$ depends on the speed of the probability of attacking converging to 0. If it is too fast, the status quo may not fall. But while players continue to learn, they learn slowly. This slow social learning process determines that the probability of attacking cannot converge to 0 too fast. As a result, the following proposition holds.

Proposition 4.2 *If $\theta = w$, in any equilibrium, the status quo is abandoned almost surely.*

Now, let's turn to the analysis conditional on $\theta = m$. Since first period players attack the status quo with positive probability, two attacks in the first period happen with positive probability. As a result, when $\theta = m$, the status quo falls with positive probability. Then does the status quo fall almost surely conditional on $\theta = m$? Because players' private signals have the unbounded likelihood ratio property, for any prior beliefs, there are signals making their posterior beliefs about $\theta = m$ arbitrarily close to 1. Therefore, it seems that in a non-trivial equilibrium (in which attacking with positive probability if possible), with probability 1, there is one period in which both players in that period receive signals informing $\theta = m$. As a result, conditional on $\theta = m$, the status quo should fall almost surely. However, this is not true. Consider the history $\bar{h}^t \equiv (1, 0, \cdots, 0)$, in which there is one attack in the first period but no attacks afterward. Note for any finite period t, $\bar{h}^t$ is on the path of play.

From the failed first period attack, players learn that $\theta > w$. So by Proposition 2.2, whether players attack the status quo with positive probability in period t depends on whether their prior belief at the beginning of period t, $\mu_t(m)[\bar{h}^t]$, is greater than or equal to $\tilde{\mu}(m)$. Then the observation of "no attack" after the first period indicates that players are increasingly pessimistic, so social learning leads to no attack for sure after some period. This analysis is summarized in the following proposition.

Proposition 4.3 *Along the history $\bar{h}^t$, in E^{max}, the probability of attacking is strictly decreasing over time until some period $T_1 \geq 1$. From period $T_1 + 1$ on, the probability of* *attacking is* 0*. Along this history, in* E^{min}*, the probability of attacking is strictly increasing over time until period* $T_2 \geq 1$*. From period* $T_2 + 1$ *on, the probability of attacking is* 0*.*

This result holds in all equilibria. Fixing any equilibrium, along the history $\bar{h}^t$, there is a T such that players stop attacking from period $T + 1$ on. As a result, $\bar{h}^\infty$ occurs with the same probability as $\bar{h}^T$. Since the status quo does not fall under $\bar{h}^\infty$, and $\bar{h}^T$ occurs with positive probability, the following proposition is obvious.

Proposition 4.4 *If* $\theta = m$*, in any equilibrium the probability of the status quo being abandoned is strictly between* 0 *and* 1*.*

So when $\theta = m$, the fundamental itself cannot predict the eventual outcome of the regime change; early period signals play important roles in determining the regime change. This is a herding result; that is, players in later periods join the herd by ignoring their own private signals, no matter how informative such signals are.

Why is the previous argument that the medium regime is almost surely abandoned wrong? Or what is the intuition of the dynamics in Proposition 4.3? This is related to the global game literature (see Carlsson and Van Damme (1993) and Morris and Shin (2007)). Fix any finite precision of private signals. Even though both players can believe $\theta = m$ by their private signals, this is never a common belief. In particular, one player, although she believes $\theta = m$, puts a bounded (away from 1) probability on the other player also receiving an extreme private signal indicating $\theta = m$. Then when the prior belief about $\theta = m$ is sufficiently small, this player's belief that her opponent will choose to attack is sufficiently low. As a result, when social learning drives the prior belief about $\theta = m$ below $\tilde{\mu}(m)$, no attack is the only outcome consistent with an equilibrium.

Finally, when $\theta = s$, the status quo never falls. However, it is still interesting to examine the dynamics of attacking in this case. One interesting and typical history is that with exactly one attack in each period, $\tilde{h}^t \equiv (1, \cdots, 1)$. Since the second period, a failed attack by one player in period t can be evidence favoring $\theta = m$ or $\theta = s$. Lemma 4.4 indicates that if $x_t^*[\tilde{h}^t] < (>) \frac{m+s}{2}$, one attack in period t leads period $t+1$ players to increase

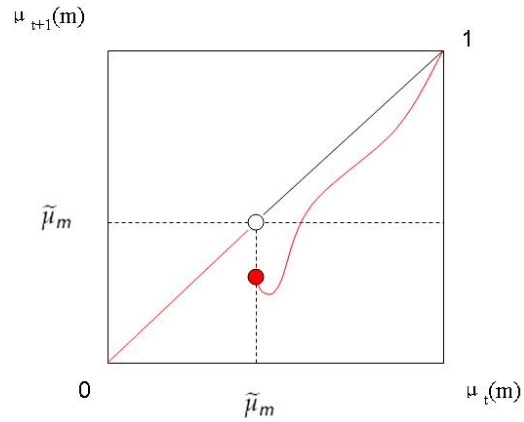

Figure 4.1: $\tilde{x} > \frac{m+s}{2}$

(decrease) their beliefs about $\theta = m$. Since players behave quite differently along $\tilde{h}^t$ in E^{max} and in E^{min}, I analyze the dynamics of attacking in E^{max} and E^{min} separately.

From Proposition 2.2, when $\mu_t(m) = \tilde{\mu}(m)$, there exists a unique $\tilde{x} \in \mathbb{R}$ such that $h(\tilde{x}; \tilde{\mu}) = 0$. If $\tilde{x} > \frac{m+s}{2}$, in period t, the largest solution to $h(x; \mu_t) = 0$ is greater than or equal to $\tilde{x}$. As a result, in E^{max}, $x_t^*[\tilde{h}^t] > \frac{m+s}{2}$. This implies that $\mu_{t+1}(m) < \mu_t(m)$ after exactly one attack in period t. Hence, in this case, if period t players attack with positive probability, period $t+1$ players attack less frequently than period t players do. Because $\tilde{x} > \frac{m+s}{2}$, one attack in period t is always evidence in favor of $\theta = s$, so players always learn from having one attack in the previous period, although players learn less over time. Therefore, beginning in some period $T_3 + 1$, players' beliefs about $\theta = m$ drop below $\tilde{\mu}(m)$, and players stop attacking. This intuition is illustrated in Figure 4.1.

When $\tilde{x} = \frac{m+s}{2}$, in E^{max}, the dynamic of attacking along $\tilde{h}^t$ is similar to that in the case of $\tilde{x} > \frac{m+s}{2}$. But the information players learn from one attack in the previous period is not bounded away from zero. That is, one attack in period t becomes neutral in the limit because $x_t^*[\tilde{h}^t]$ is decreasing toward $\tilde{x}$ over time (and if $x_t^*[\tilde{h}^t] = \tilde{x}$, one attack in period t is neutral). However, as shown in Proposition 4.5, if $x_t^*[\tilde{h}^t]$ is in the neighborhood of $\tilde{x}$, μ_t is so close to $\tilde{\mu}$ that the discrete adjustment of the prior belief after one attack in period t makes $\mu_{t+1}(m) < \tilde{\mu}(m)$. Hence, players stop attacking from that period onward. Figure 4.2 shows this intuition.

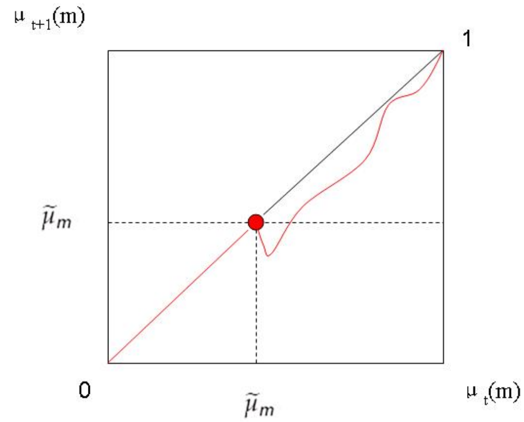

Figure 4.2: $\tilde{x} = \frac{m+s}{2}$

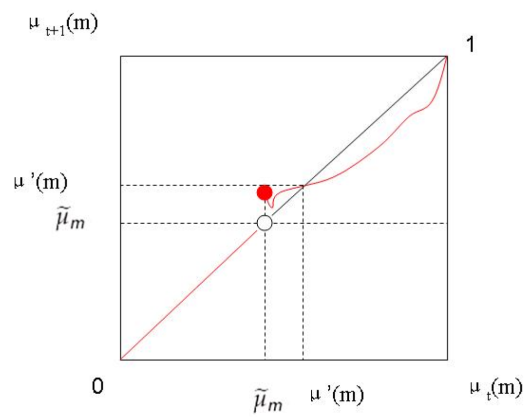

Figure 4.3: $\tilde{x} < \frac{m+s}{2}$

The last case is $\tilde{x} < \frac{m+s}{2}$. When $\mu_t(m) > (<)\mu'(m)$ (where μ' is the prior belief such that $\frac{m+s}{2}$ is the biggest solution to the equation $h(x;\mu') = 0$), one attack in period t is evidence favoring $\theta = s$ ($\theta = m$). When β is very large and $x_t^*[\tilde{h}^t]$ is in the neighborhood of $\tilde{x}$, the adjustment of the prior from one attack in period t is very small, so that $\mu_{t+1}(m)$ is between $\mu_t(m)$ and $\mu'(m)$. Therefore, the probability of attacking converges to $\Phi(\frac{\sqrt{\beta}}{2}(s-m))$ along $\tilde{h}^t$. See Figure 4.3 for this case.

The analyses above are summarized in the following proposition:

Proposition 4.5 *In E^{max}, along the history $\tilde{h}^t$,*

1. *If $\tilde{\mu}(m) > \frac{m+s}{2}$, the probability of attacking is decreasing over time after period 2; and there is a period T_3, such that the probability of attacking is 0 from period $T_3 + 1$ on.*

2. If $\tilde{\mu}(m) = \frac{m+s}{2}$, the probability of attacking is decreasing over time after period 2; and there is a period T_4 such that from period T_4+1, the probability of attacking is constant at either $\Phi(\frac{\sqrt{\beta}}{2}(s-m))$ or 0.

3. If $\tilde{\mu}(m) < \frac{m+s}{2}$, the probability of attacking is decreasing (increasing) over time if $\mu_2(m) > (<)\mu'(m)$; and the probability of attacking converges to $\Phi(\frac{\sqrt{\beta}}{2}(s-m))$.

Now, let's move to the dynamics of attacking along $\tilde{h}^t$ in E^{min}. Lemma 4.4 implies that if $x_2^*[\tilde{h}^2] < (>)\frac{m+s}{2}$, the probability of attacking is decreasing (increasing) over time. The following proposition describes the asymptotic behavior in E^{min} along $\tilde{h}^t$.

Proposition 4.6 *In E^{min}, along $\tilde{h}^t$, if $x_2^*[\tilde{h}^2] > \frac{m+s}{2}$, there exists a period T_5 such that the probability of attacking is 0 from period T_5+1 on; if $x_2^*[\tilde{h}^2] < \frac{m+s}{2}$, the probability of attacking converges to c, the cost of attacking.*

It is interesting to compare part 3 of Proposition 4.5 with the second case of Proposition 4.6. Along $\tilde{h}^t$, in both E^{max} and E^{min} under some conditions, the probability of attacking converges. However, the social learning processes, which drive these two convergence results, are different. In E^{max}, because one attack in the previous period becomes neutral evidence, players learn more slowly over time, and in the limit they learn nothing. This leads to convergence in E^{max}. In E^{min}, players learn increasingly more from one attack in the previous period, and in the limit they are convinced that $\theta = m$. So players in the limit follow the strategy, which purifies the mixed strategy 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.)

4.3 Social Welfare

In this section, I analyze the effects 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 2.1, in any equilibrium, players attack the status quo with positive probability (bounded away from 0) in each period. Consider the ex-post social welfare $\mathcal{W} = (1-\delta)\sum_{t=1}^{\infty}\delta^t(u_t^1 + u_t^2)$, 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 4.7 *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 a $\delta_m \in (0,1)$, such that social learning is inefficient for any $\delta \in (\delta_m, 1)$;*

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

The intuition behind the first part of this proposition is that social learning delays the regime change. However, when the discount factor is sufficiently large, such inefficiency disappears. In the second part, when $\theta = m$, social learning causes players to stop attacking with positive probability, so the status quo may not fall with positive probability, which is inefficient because positive utilities after the regime change cannot be collected. In the third case, the lower the number of attacks the better, because the regime cannot be beaten. Since social learning can prevent infinitely many attacks, it is more efficient.

5 Extensions

In this section, I consider two extensions of the dynamic regime change game of the benchmark model. In the first extension, I allow the precision of signals, β_t, to change over time deterministically. In particular, $\{\beta_t\}_t$ is assumed to be an increasing sequence and $\lim_{t\to\infty}\beta_t = \infty$. This extension captures the idea of information technology innovation and demonstrates how the benchmark model differs from models with private learning. The second extension models a N-player $(N+1)$-state dynamic regime change game. Restricted to the monotone equilibrium, I get results very similar to the benchmark model.

5.1 Private Learning

In the benchmark model, 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 $\tilde{\mu}(m)$ below which players choose not to attack, ignoring their private signals, is constant over time (m, s and c are given). As a result, along the history $\bar{h}^t \equiv (1, 0, \cdots, 0)$, there is a period T such that $\mu_{T+1}(m) < \tilde{\mu}(m)$. Since period $T+1$ players choose not to attack with probability 1, period $T+2$ players could not update their beliefs about $\theta = m$ from the no attack outcome in period $T+1$. So $\mu_{T+2}(m) = \mu_{T+1}(m) < \tilde{\mu}(m)$, and period $T+2$ players will also choose not to attack for any private signals. This analysis leads to Proposition 4.3 and Proposition 4.4.

Then, what happens if private learning is allowed? In particular, does the herding result in Proposition 4.3 occur with private learning? And can players with private learning beat the medium status quo almost surely? Because information technology innovates over time, a straightforward way to incorporate private learning in the benchmark model is to assume that players in the later period have more precise private information than players in the earlier periods and that private signals are accurate in the limit. That is, the precision of period t's private signals, β_t, increases over time and converges to ∞.

The most interesting case is when $\theta = m$. Consider the history $\bar{h}^t$. A failed attack by one player in the first period rules out $\theta = w$, so players lose the dominant region of attacking. While the prior belief about $\theta = m$, $\mu_t(m)$, keeps decreasing over time, the critical belief $\tilde{\mu}_t(m)$ changes over time. From Proposition 2.2, attacks can occur with positive probability in period t, in which $\mu_t(m) \geq \tilde{\mu}_t(m)$, and no attack is the unique equilibrium outcome in period t', in which $\mu_{t'}(m) < \tilde{\mu}_{t'}(m)$. Since $\tilde{\mu}_t(m)$ goes to 0 (by Lemma 7.5 in the appendix), unless the status quo is abandoned, attacks occur in infinitely many periods. Therefore, when $\theta = m$, the status quo falls almost surely.

The reason why the herding result in Proposition 4.3 disappears is due to the increasing without bound precision of private signals. As players' precision of 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 common prior belief (if, in equilibrium, period t's players' strategies are to not attack for all signals, then period $t+1$ players' prior beliefs are the same as period t's players'), there is a sufficiently large precision of private signals such that attacking with positive probability is consistent with an equilibrium.

This extended model with private learning shows the possibility of one transition in Angeletos, Hellwig and Pavan (2007). That is, an economy can transit back and forth between "tranquility" phases (no attack in any equilibrium) and "distress" phases (attack with positive probability in non-trivial equilibria). Without private learning, this transition cannot happen.

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

The benchmark 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 benchmark 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 $\Theta = \{m_1, \cdots, m_{N+1}\}$ with $N > 2$ and $m_1 < m_2 < \cdots < m_{N+1}$, is the set of states. When $\theta = m_k$, at least k attacks are needed to trigger the regime change. Hence, when $\theta = m_1$, attacking is the dominant action for players, because one attack is enough to make the status quo be abandoned. 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 2.1 and Proposition 2.2, given some strategy profile of other players S_{-i}, the best response of player i may not be a cutoff strategy. In particular, 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 7.6 in the appendix shows that if all other players are following this kind of strategy, 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 2.1, 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), only a trivial equilibrium exists. Hence, restricted to monotone equilibria, the static game analysis is similar to that of the benchmark model.

For any fixed monotone equilibrium, the eventual outcome of regime change in this extended model is similar to that of the benchmark model. When $\theta = m_1$, an $N + 1$ state version of Lemma 7.4 and the fact that after an infinitely long history without any attack players' belief about $\theta = m_1$ is 0 together imply that the status quo falls almost surely. When $\theta = m_{N+1}$, by assumption, the status quo never falls. For any state, which requires coordination to trigger the regime change, the state itself cannot predict the eventual outcome of the regime change. This is so for two reasons: (1) the history $\bar{h}^t = (1, 0, \cdots, 0)$ happens with positive probability for any period t, and (2) as t becomes large, players' beliefs about $\theta = m_{N+1}$ are larger than $\tilde{\mu}(m_{N+1})$.

6 Conclusion

This paper examined how social learning influences the dynamics of attacking and the eventual outcome of regime change in a dynamic regime change game without perfect individual learning. I show that when the regime is weak, although players become pessimistic and eventually choose not to attack, the social learning process guarantees that the regime falls almost surely. When the regime is in the medium state, social learning leads to a herding result. That is, when players are sufficiently pessimistic before observing their private signals (because they have updated their beliefs from previous plays), they choose not to attack despite their private signals. Therefore, in this case, both the status quo and the early signals determine the eventual outcome of the regime change. My result also shows that although the dynamics of attacking are similar in different equilibria, they are driven by different social learning processes.

Social learning may lead to more or less efficient outcomes, compared with the scenario without social learning. In particular, when the status quo is weak or medium, social learning delays the regime change time, which is inefficient. But when the status quo is strong, social learning can prevent infinitely many attacks, which leads to a higher discounted social welfare.

My results contribute to the growing literature on dynamic global games by analyzing social learning instead of individual learning and to the herding literature by incorporating a strategic complementarity environment in each period. From a methodological perspective, multiple equilibria emerge in my setting, as in Angeletos, Hellwig, and Pavan (2007). However, the multiplicity is due not only to the pure coordination property after the weak status quo is ruled out but also from the updating of priors over time. In particular, the uniqueness established in Morris and Shin (2003) comes from fixing a prior belief and letting private signals become very precise. However, in my model, the precision of private signals is fixed while the prior is being updated over time. So multiplicity may appear under the history without any attack (a global game).

7 Appendix: Proofs

PROOF OF PROPOSITION 2.1: I first show that if a Bayesian Nash Equilibrium exists, it is in cutoff strategies. Because signals are conditionally independent, in equation 2.3, 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 $E_{x^j} u^i(1, s^j|x^i)$ is strictly decreasing in x^i. Note also that since $\lim_{x^i \to -\infty} E_{x^j} u^i(1, s^j|x^i) = 1 - c > 0$ (dominant region of attacking) and $\lim_{x^i \to +\infty} E_{x^j} u^i(1, s^j|x^i) = -c < 0$ (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 not, 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 $E_{x^j} u^i(1, s^j|x^i)$ is strictly supermodular and $E_{x^j} u^i(1, \hat{x}^j|\hat{x}^i) = 0$, $\hat{x}^i$ is strictly increasing in $\hat{x}^j$. Thus $\hat{\hat{x}}^2 = \hat{x}^1 > \hat{x}^2$ implies $\hat{x}^2 = \hat{\hat{x}}^1 > \hat{x}^1$. Contradiction.

Now consider any symmetric cutoff strategy profile (x, x). Fix any prior μ, the interim payoff from attacking given the signal x and the opponent's cutoff strategy with threshold point x can be written as:

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

Because $G(x;\mu)$ is continuous in x, the dominant region of attacking and 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 $\mu \in int(\Delta(\Theta))$ such that multiple equilibria exist in this static regime change game. To show this claim, I just need to show that there exists a $\mu \in int(\Delta(\Theta))$ such that there are more than one solutions to $G(x;\mu) = 0$. Note that $\lim_{\mu(m) \to 1} G(\frac{w+m}{2};\mu) = \Phi(\frac{\sqrt{\beta}}{2}(w-m)) - c < 0$ (by $c > \frac{1}{2}$) and $\lim_{\mu(m) \to 1} G(\frac{m+s}{2};\mu) = \Phi(\frac{\sqrt{\beta}}{2}(s-m)) - c > 0$ (by Assumption 2.1). 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 $(\frac{m+s}{2}, +\infty)$. *Q.E.D.*

PROOF OF PROPOSITION 2.2: First when $\mu(w)=0$, for a fixed s^j, $E_{x^j}u^i(1,s^j|x^i)$ is strictly decreasing in x^i, and the regime change game is supermodular. So similar to the proof of Proposition 2.1, if a non-trivial equilibrium exists, it is symmetric and in cutoff strategies. Denote a symmetric cutoff strategy profile by (x,x). Let $H(x;\mu)=G(x;\mu(w)=0)$, then (x^*,x^*) is a non-trivial equilibrium of the regime change game if and only if $H(x^*;\mu)=0$. Therefore, conditions for the existence of a non-trivial equilibrium are equivalent to those for the existence of a solution to $H(x;\mu)=0$. Note $H(x;\mu)$ can be equivalently written as $H(x;\mu)=f(x;\mu)h(x;\mu)$, where

$$g(x;\mu)=\frac{\mu(m)\phi(\sqrt{\beta}(x-m))}{\mu(m)\phi(\sqrt{\beta}(x-m))+(1-\mu(m))\phi(\sqrt{\beta}(x-s))}$$

and

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

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

Lemma 7.1 *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 $h(x;\mu)=0$; and for all $\mu(m)\in[\bar{\mu}(m),1)$, there is $x^*\in\mathbb{R}$ such that $h(x^*;\mu)=0$.*

Proof.

First consider the case where $\mu(m)$ is close to 1. Since $\Phi(\frac{\sqrt{\beta}}{2}(s-m))>c$, $h(\frac{s+m}{2};\mu)>0$. Note that for all $\mu(m)\in(0,1)$, $h(m;\mu)<0$ and $\lim_{x\to+\infty}h(x,\mu)<0$, so by continuity of $h(x;\mu)$ in x, there exist $\hat{x}\in(m,\frac{s+m}{2})$ and $\hat{\hat{x}}\in(\frac{s+m}{2},+\infty)$ such that $h(\hat{x},\mu)=0$ and $h(\hat{\hat{x}},\mu)=0$. Therefore, there exists $\bar{\mu}(m)\in(0,1)$ such that solutions to $h(x;\mu)=0$ exist for all $\mu(m)\in[\bar{\mu}(m),1)$. Now consider $\mu(m)$ is close to 0. The last term of $h(x;\mu)$ is very negative for any x larger than m, so $h(x;\mu)<0$ for all $x>m$. Combined with the fact that $h(x;\mu)<0$ for all $x\le m$, there exists $\underline{\mu}(m)\in(0,1)$ such that for all $\mu(m)\in(0,\underline{\mu}(m)]$,

$h(x;\mu) < 0, \forall x \in \mathbb{R}$. 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 7.2 *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 $h(x;\mu) = 0$; and for all $\mu(m) \in (\tilde{\mu}(m), 1)$, there are two solutions to $h(x;\mu) = 0$. Therefore, claims (1) and (2) in Proposition 2.2 are true.*

Proof.

Suppose $1 > \mu'(m) > \mu''(m) > 0$ and $\exists x'' \in (m, +\infty)$ such that $h(x''; \mu'') = 0$ (because all $x \leq m$ cannot be a solution to $h(x;\mu'') = 0$). Since $h(x;\mu)$ is strictly increasing in $\mu(m)$ for any fixed $x \in \mathbb{R}$, $h(x'';\mu') > h(x'';\mu'') = 0$. Then by the continuity of $h(x;\mu')$ and $\lim_{x \to +\infty} h(x, \mu') < 0$, there exists $x' \in (x'', +\infty)$ such that $h(x';\mu') = 0$. Similarly, if $1 > \mu'(m) > \mu''(m) > 0$ and $h(x;\mu') < 0$ for all $x \in \mathbb{R}$, $h(x;\mu'') < 0$ for all $x \in \mathbb{R}$. Define $\tilde{\mu}(m) = \inf\{\mu(m) \in (0,1) : \exists x \in \mathbb{R} \text{ such that } h(x;\mu) = 0\} = \sup\{\mu(m) \in (0,1) : h(x;\mu) < 0 \quad \forall x \in \mathbb{R}\}$ (since for a given μ, $h(x;\mu)$ either has a solution or doesn't 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 $\frac{\partial^2 h}{\partial x^2} < 0$ for all $x \geq m$ and $h(x;\mu)$ has a single peak in $(m, +\infty)$. Therefore, when $\mu(m) \in (\tilde{\mu}(m), 1)$, there are two solutions to $h(x;\mu) = 0$. ■

Lemma 7.3 *There exists a unique $\tilde{x} \in (m, +\infty)$ such that $h(\tilde{x}, \tilde{\mu}) = 0$. Therefore, claim (3) in Proposition 2.2 is true.*

Proof.

Suppose $\forall x \in \mathbb{R}$, $h(x;\tilde{\mu}) < 0$. Recall that because $\mu(m) < 1$, for any $x \in (\bar{x}(\tilde{\mu}(m)), +\infty)$, $h(x;\tilde{\mu}(m)) < 0$ (dominant region of not attacking), where $\bar{x}(\mu(m)) = \inf\{x \in \mathbb{R} : E_{x^j} u^i(1, s^j | x^i, \mu) < 0 \text{ for all } s^j\}$. Since $E_{x^j} u^i(1, s^j | x^i)$ is increasing in $\mu(m)$, $\bar{x}(\mu(m))$ is an increasing function in $\mu(m)$. As a result, $h(x;\tilde{\mu}(m)) < 0$ for all $x > \bar{x}(\bar{\mu}(m))$, because $\tilde{\mu}(m) < \bar{\mu}(m)$. Since $c > \frac{1}{2}$, for any μ, $h(x;\mu) < 0$ for all $x < m$. Now consider the compact set $[m, \bar{x}(\bar{\mu}(m))]$. From the continuity of $h(x;\mu)$ in x, $\exists \hat{x} \in [m, \bar{x}(\bar{\mu}(m))]$ such that $h(x;\tilde{\mu}) \leq h(\hat{x};\tilde{\mu}) < 0$. 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 $h(x;\mu)$ is continuous in $\mu(m)$ for any $x \in [m, \bar{x}(\bar{\mu}(m))]$, $\lim_{k \to +\infty} h(x;\mu^k) = h(x;\tilde{\mu})$. Defining $M^k = \sup_{x \in [m, \bar{x}(\bar{\mu}(m))]} |h(x;\mu^k) - h(x;\tilde{\mu})|$, it can be calculated that

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

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 $h(x;\mu^k)$ converges to $h(x;\tilde{\mu})$ uniformly. So there exists K' such that for all $k > K'$, $h(x;\mu^k) - h(x;\tilde{\mu}) < \frac{|h(\hat{x},\tilde{\mu}(m))|}{2}$, so $h(x;\mu^k) < -\frac{|h(\hat{x},\tilde{\mu}(m))|}{2} < 0$ for all $x \in [m, \bar{x}(\bar{\mu}(m))]$. Note for any $x < m$ and $x > \bar{x}(\bar{\mu}(m))$, $h(x;\mu^k) < 0$, so for all $x \in \mathbb{R}$, $h(x;\mu^k) < 0$. But by the definition of $\tilde{\mu}(m)$, there must be some $x' \in \mathbb{R}$ such that $h(x';\mu^k) = 0$. Therefore, when $\mu(m) = \tilde{\mu}(m)$, there exists $\tilde{x}$ such that $h(\tilde{x};\tilde{\mu}) = 0$.

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

The following graph of $h(x,\mu)$ is helpful for further analysis. In the graph, $\mu''(m) > \tilde{\mu}(m) > \mu'(m)$, so $h(x;\mu'') = 0$ has two solutions, $h(x;\tilde{\mu}) = 0$ has unique solution $\tilde{x}$, and $h(x;\mu') < 0$ for all $x \in \mathbb{R}$.

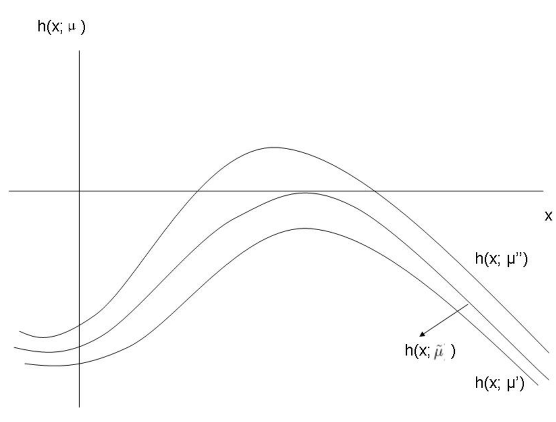

Figure 7.1: $\mu''(m) > \tilde{\mu}(m) > \mu'(m)$

Q.E.D.

PROOF OF LEMMA 4.1: Because there is no attack before period $t+1$ and any finite history without any attack is on the path of play conditional on any $\theta \in \Theta$, Bayes' rule implies that $\mu_t, \mu_{t+1} \in int(\Delta(\Theta))$ in any equilibrium. Fix any x and recall $\hat{h}^t \equiv (0, \cdots, 0)$,

$$
\begin{aligned}
& G(x;\mu_{t+1}) \\
= \; & \frac{\mu_{t+1}(w)\phi(\sqrt{\beta}(x-w))}{\sum\limits_{\theta'\in\Theta}\mu_{t+1}(\theta')\phi(\sqrt{\beta}(x-\theta'))} + \frac{\mu_{t+1}(m)\phi(\sqrt{\beta}(x-m))\Phi(\sqrt{\beta}(x-m))}{\sum\limits_{\theta'\in\Theta}\mu_{t+1}(\theta')\phi(\sqrt{\beta}(x-\theta'))} - c \\
= \; & \frac{\begin{bmatrix} \mu_t(w)[\Phi(\sqrt{\beta}(w-x_t^*(\hat{h}^t)))]^2\phi(\sqrt{\beta}(x-w)) \\ + \mu_t(m)[\Phi(\sqrt{\beta}(m-x_t^*(\hat{h}^t)))]^2\phi(\sqrt{\beta}(x-m))\Phi(\sqrt{\beta}(x-m)) \end{bmatrix}}{\begin{bmatrix} \mu_t(w)[\Phi(\sqrt{\beta}(w-x_t^*(\hat{h}^t)))]^2\phi(\sqrt{\beta}(x-w)) \\ + \mu_t(m)[\Phi(\sqrt{\beta}(m-x_t^*(\hat{h}^t)))]^2\phi(\sqrt{\beta}(x-m)) \\ + \mu_t(s)[\Phi(\sqrt{\beta}(s-x_t^*(\hat{h}^t)))]^2\phi(\sqrt{\beta}(x-s)) \end{bmatrix}} - c \\
< \; & \frac{\begin{bmatrix} \mu_t(w)[\Phi(\sqrt{\beta}(m-x_t^*(\hat{h}^t)))]^2\phi(\sqrt{\beta}(x-w)) \\ + \mu_t(m)[\Phi(\sqrt{\beta}(m-x_t^*(\hat{h}^t)))]^2\phi(\sqrt{\beta}(x-m))\Phi(\sqrt{\beta}(x-m)) \end{bmatrix}}{\begin{bmatrix} \mu_t(w)[\Phi(\sqrt{\beta}(m-x_t^*(\hat{h}^t)))]^2\phi(\sqrt{\beta}(x-w)) \\ + \mu_t(m)[\Phi(\sqrt{\beta}(m-x_t^*(\hat{h}^t)))]^2\phi(\sqrt{\beta}(x-m)) \\ + \mu_t(s)[\Phi(\sqrt{\beta}(m-x_t^*(\hat{h}^t)))]^2\phi(\sqrt{\beta}(x-s)) \end{bmatrix}} - c \\
= \; & \frac{\mu_t(w)\phi(\sqrt{\beta}(x-w))}{\sum\limits_{\theta'\in\Theta}\mu_t(\theta')\phi(\sqrt{\beta}(x-\theta'))} + \frac{\mu_t(m)\phi(\sqrt{\beta}(x-m))\Phi(\sqrt{\beta}(x-m))}{\sum\limits_{\theta'\in\Theta}\mu_t(\theta')\phi(\sqrt{\beta}(x-\theta'))} - c \\
= \; & G(x;\mu_t).
\end{aligned}
$$

In E^{max}, let $x_{t+1}^*(\hat{h}^{t+1})$ be the largest solution to $G(x;\mu_{t+1}) = 0$, so $G(x_{t+1}^*(\hat{h}^{t+1});\mu_t) > G(x_{t+1}^*(\hat{h}^{t+1});\mu_{t+1}) = 0$. Since $\lim\limits_{x\to+\infty} G(x;\mu_t) < 0$ (dominant region of not attacking), there exists $x' \in (x_{t+1}^*(\hat{h}^{t+1}), +\infty)$ such that $G(x';\mu_t) = 0$. Therefore, $x_t^*(\hat{h}^{t+1}) \geq x' > x_{t+1}^*(\hat{h}^{t+1})$, which implies that if there is no attack before period $t+1$, period $t+1$ players attack the status quo with lower probability than period t players in E^{max}. In the similar way, it can be shown that in E^{min}, if there is no attack in the first period, second period players attack the status quo with lower probability. *Q.E.D.*

PROOF OF LEMMA 4.2: To simplify notation, I ignore all history entries. Just keep in mind that in this lemma, there is no attack before period t and there is exactly one attack in period t. First note that if $\mu_t(w)$ is very close to 1, $G(\frac{m+s}{2};\mu_t) > 0$. Therefore, in E^{max}, $x_t^* > \frac{m+s}{2}$. So I just need to show the case that $x_t^* \geq \frac{m+s}{2}$. Fix any x, let $A = \Phi(\sqrt{\beta}(m - x_t^*))\Phi(\sqrt{\beta}(x_t^* - m))$ and $B = \Phi(\sqrt{\beta}(s - x_t^*))\Phi(\sqrt{\beta}(x_t^* - s))$, from footnote 6,

$$
\begin{aligned}
& H(x;\mu_{t+1}) \\
= \;& \frac{\mu_t(m)A\phi(\sqrt{\beta}(m-x))\Phi(\sqrt{\beta}(x-m))}{\mu_t(m)A\phi(\sqrt{\beta}(m-x)) + \mu_t(s)B\phi(\sqrt{\beta}(s-x))} - c \\
\leq \;& \frac{\mu_t(m)\phi(\sqrt{\beta}(m-x))\Phi(\sqrt{\beta}(x-m))}{\mu_t(m)\phi(\sqrt{\beta}(m-x)) + \mu_t(s)\phi(\sqrt{\beta}(s-x))} - c \\
< \;& \frac{\mu_t(w)\phi(\sqrt{\beta}(w-x)) + \mu_t(m)\phi(\sqrt{\beta}(m-x))\Phi(\sqrt{\beta}(x-m))}{\mu_t(w)\phi(\sqrt{\beta}(w-x)) + \mu_t(m)\phi(\sqrt{\beta}(m-x)) + \mu_t(s)\phi(\sqrt{\beta}(s-x))} - c \\
= \;& G(x;\mu_t).
\end{aligned}
$$

Since in E^{max}, x_t^* is the largest solution to $G(x;\mu_t) = 0$ and $\lim\limits_{x\to+\infty} G(x;\mu_t) < 0$, $G(x;\mu_t) < 0, \quad \forall x > x_t^*$, which implies that $H(x;\mu_{t+1}) < 0, \quad \forall x > x_t^*$. Therefore, all threshold points in period $t+1$ are smaller than x_t^*, so the probability of attacking in period $t+1$ is lower than that in period t.

Now consider the case that in E^{max}, $\mu_t(w)$ is close to 0 and $x_t^* < \frac{m+s}{2}$. Define

$$
\begin{aligned}
& D_1 \\
= \;& G(x_t^*;\mu_t) - [\frac{\mu_t(m)\phi(\sqrt{\beta}(m-x_t^*))\Phi(\sqrt{\beta}(x_t^*-m))}{\mu_t(m)\phi(\sqrt{\beta}(m-x_t^*)) + \mu_t(s)\phi(\sqrt{\beta}(s-x_t^*))} - c] \\
= \;& \frac{\begin{bmatrix} \mu_t(w)\phi(\sqrt{\beta}(x_t^*-w)) \\ \quad * [\mu_t(m)\phi(\sqrt{\beta}(x_t^*-m))(1-\Phi(\sqrt{\beta}(m-x_t^*))) + \mu_t(s)\phi(\sqrt{\beta}(x_t^*-s))] \end{bmatrix}}{\begin{bmatrix} (\sum_{\theta'\in\Theta}\mu(\theta')\phi(\sqrt{\beta}(x_t^*-\theta'))) \\ \quad * (\sum_{\theta'\in\{m,s\}}\mu(\theta')\phi(\sqrt{\beta}(x_t^*-\theta'))) \end{bmatrix}},
\end{aligned}
$$

and

$$
\begin{aligned}
& D_2 \\
= \;& G(x_t^*;\mu_{t+1}) - [\frac{\mu_t(m)\phi(\sqrt{\beta}(m-x_t^*))\Phi(\sqrt{\beta}(x_t^*-m))}{\mu_t(m)\phi(\sqrt{\beta}(m-x_t^*))+\mu_t(s)\phi(\sqrt{\beta}(s-x_t^*))} - c] \\
= \;& \frac{\mu_t(m)\phi(\sqrt{\beta}(m-x_t^*))\Phi(\sqrt{\beta}(x_t^*-m))\mu_t(s)\phi(\sqrt{\beta}(s-x_t^*))(A-B)}{\begin{bmatrix} (\mu_t(m)\phi(\sqrt{\beta}(m-x_t^*))A+\mu_t(s)\phi(\sqrt{\beta}(s-x_t^*))B) \\ *\,(\mu_t(m)\phi(\sqrt{\beta}(m-x_t^*))+\mu_t(s)\phi(\sqrt{\beta}(s-x_t^*))) \end{bmatrix}}.
\end{aligned}
$$

Note that when $x_t^* < \frac{m+s}{2}$, $A > B$ and $A-B$ is decreasing in x_t^*. Consider $\mu_t = (\epsilon, \tilde{\mu}(m), 1-\tilde{\mu}(m)-\epsilon)$. It can be shown that, x_t^* is decreasing with ϵ, but x_t^* is bounded below by $\tilde{x}$, the unique solution to $G(x;\tilde{\mu}) = 0$. Therefore, as $\epsilon \to 0$, $A-B$ becomes larger, so D_2 is positive and bounded away from 0. Since D_1 converges to 0 as $\epsilon \to 0$, in this case, when $\mu_t(w)$ is very close to 0, $G(x_t^*;\mu_{t+1}) > G(x_t^*;\mu_t) = 0$. Therefore, in E^{max}, the threshold point in period $t+1$ is larger than that in period t, which implies that the probability of attacking is larger in period $t+1$ than that in period t. *Q.E.D.*

PROOF OF LEMMA 4.3: First note $\mu_{t+1}(m) < \mu_t(m)$, so for any x, $H(x;\mu_t) > H(x;\mu_{t+1})$. In E^{max}, if x_{t+1}^* is the largest threshold point in period $t+1$, then $H(x_{t+1}^*;\mu_t) > H(x_{t+1}^*;\mu_{t+1}) = 0$. Recall that $\lim_{x\to+\infty} H(x;\mu_t) < 0$, $\exists x' \in (x_{t+1}^*, +\infty)$ such that $H(x';\mu_t) = 0$. Let x_t^* be the largest threshold point in period t, then $x_t^* \geq x' > x_{t+1}^*$. So in E^{max}, the probability of attacking in period $t+1$ is lower than that in period t. Now consider E^{min}. Let x_{t+1}^* be the smallest threshold point in period $t+1$, then $H(x_{t+1}^*;\mu_t) > H(x_{t+1}^*;\mu_{t+1}) = 0$. Because $c > \frac{1}{2}$, for all $x \leq m$, $0 > H(x;\mu_t) > G(x;\mu_{t+1})$. So, $x_{t+1}^* > m$, and $x_t^* \in (m, x_{t+1}^*)$. As a result, in E^{min}, the probability of attacking in period $t+1$ is higher than that in period t. *Q.E.D.*

PROOF OF PROPOSITION 4.1: From Bayes' rule,

$$
\mu_{t+1}(w) = \frac{\mu_t(w)[\Phi(\sqrt{\beta}(w-x_t^*))]^2}{\sum_{\theta'\in\Theta} \mu_t(\theta')[\Phi(\sqrt{\beta}(w-x_t^*))]^2}. \tag{7.1}
$$

By induction, $\mu_{t+1} \in int(\Delta(\Theta))$ and $x_t^* \in \mathbb{R}$. Therefore, $\mu_{t+1}(w) < \mu_t(w)$, so $\mu_t(w) \to \mu_\infty(w) \geq 0$, since μ_t is a probability measure. Lemma 4.1 shows that in both E^{max} and

E^{min}, x_t^* is strictly decreasing, so $x_t^* \to x_\infty^*$. Suppose $\mu_\infty(w) > 0$, then because of the dominant region of attacking, $x_\infty^* \in \mathbb{R}$. Then from the updating function 7.1, $\mu_\infty(w) = 0$. Contradiction. So $\mu_t(w) \to 0$.

Because Pr(attack in period $t|\hat{h}^t$) is decreasing in t and bounded, Pr(attack in period $t|\hat{h}^t$) $\to$ $p \geq 0$. Suppose $p > 0$, then $x_t^* \to x_\infty^* \in \mathbb{R}$. Consequently, $\lim_{t\to\infty} G(x_t^*; \mu_t) = -c$, that is, for sufficiently large t, $G(x_t^*; \mu_t) < 0$, which contradicts the definition of x_t^*. *Q.E.D.*

PROOF OF PROPOSITION 4.2: This proposition is a consequence of Proposition 4.1 and the following lemma, which states that conditional on $\theta = w$, Bayes' rule leads prior belief about $\theta = w$ to converge to a random variable larger than 0 almost surely. Since along $\hat{h}^\infty$, $\mu_t(w) \to 0$, conditional on $\theta = w$, $\hat{h}^\infty$ is a zero measure event. Therefore, when $\theta = w$, the status quo falls almost surely.

Lemma 7.4 *Fix an equilibrium, the strategy profile and the prior belief induce a probability measure* $\mathbb{P} \in \Delta(X^\infty)$, *where* $X = \{0, 1, 2\}$. *Suppose* $\mathbb{P}_w$ *and* $\hat{\mathbb{P}}$ *are the probability measures induced on* X^∞ *by* $\mathbb{P}$, *conditional on* $\theta = w$ *and* $\theta \in \{m, s\}$ *respectively. Hence,* $\mathbb{P} = \mu_0(w)\mathbb{P}_w + (1 - \mu_0(w))\hat{\mathbb{P}}$. *Then* $\mu_t(w) \to \mu_\infty(w), \mathbb{P}_w -$ *almost surely and* $\mu_\infty(w) > 0, \mathbb{P}_w -$ *almost surely.*

Proof.

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 $\mathbb{P}$, $\mu_t(w) \to \mu_\infty(w), \mathbb{P}_w -$ almost surely.

Now suppose there is a set $A \in X^\infty$ such that $\mu_\infty(w)[a] = 0, \forall a \in A$ and $\mathbb{P}_w(A) > 0$. It is easy to show that from Bayes' rule, the odds ratio $\{(1-\mu_t(w))/\mu_t(w)\}_t$ is a $\mathbb{P}_w-$martingale, so $E[\frac{1-\mu_t(w)}{\mu_t(w)}] = \frac{1-\mu_0(w)}{\mu_0(w)}$ for all t. However, $E[\frac{1-\mu_t(w)}{\mu_t(w)}] = E[\frac{1-\mu_t(w)}{\mu_t(w)}\chi(A)] + E[\frac{1-\mu_t(w)}{\mu_t(w)}(1-\chi(A))]$, 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 lead to a contradiction. ■ *Q.E.D.*

PROOF OF PROPOSITION 4.3: The monotonicity property follows directly from lemma 4.3. To show the existence of T_1, suppose in an equilibrium, for any period t, players attack the status quo with positive probability (this probability is bounded below by 1/2 because $c > 1/2$). Then from Bayes' rule, $\mu_t(m)[\bar{h}^t] \to 0 < \tilde{\mu}(m)$. That is, there is T_1 such that for all $t > T_1$, $\mu_t(m)[\bar{h}^t] < \tilde{\mu}(m)$, so for all $t > T_1$, players don't attack the status quo for sure. This leads to the contradiction. In a similar way, the existence of T_2 can be demonstrated. *Q.E.D.*

PROOF OF PROPOSITION 4.5: Part 1: $\tilde{x} > \frac{m+s}{2}$. If $\mu_t(m) \geq \tilde{\mu}(m)$, then $x_t^*[\tilde{h}^t] > \frac{m+s}{2}$. So from Bayes' rule, one attack in period t makes $\mu_{t+1}(m) < \mu_t(m)$, which implies the attacking probability is decreasing along $\tilde{h}^t$ in E^{max}. Suppose there is no T_3 such that the probability of attacking is 0 from $T_3 + 1$ on in E^{max}, then Bayes' rule implies that $\mu_t(m) \to 0 < \tilde{\mu}(m)$. That is, there is T_3 such that for all $t > T_3$, $\mu_t(m)[\bar{h}^t] < \tilde{\mu}(m)$, so for all $t > T_3$, players don't attack the status quo for sure. This leads to the contradiction.

Part 2: $\tilde{x} = \frac{m+s}{2}$. The decreasing attacking probability part is similar to that in part 1. I show the existence of T_4 as follows. First note that, when $\theta = w$ has been ruled out, in any period, given the previous period prior belief $\mu(m)$ and one attack in the previous period, Bayes' rule implies that the prior belief about $\theta = m$ at the beginning of the current period, $T\mu(m)$ is:

$$T\mu(m) = \frac{\mu(m)\Phi(\sqrt{\beta}(x^* - m))\Phi(\sqrt{\beta}(m - x^*))}{\mu(m)\Phi(\sqrt{\beta}(x^* - m))\Phi(\sqrt{\beta}(m - x^*)) + (1 - \mu(m))\Phi(\sqrt{\beta}(x^* - s))\Phi(\sqrt{\beta}(s - x^*))},$$

where $x^* \in \mathbb{R}$ is the largest solution to $H(x; \tilde{\mu}) = 0$. So at the point $\mu(m) = \tilde{\mu}(m)$,

$$\lim_{\mu(m)\to\tilde{\mu}(m)^+} \frac{dT\mu(m)}{d\mu(m)}|_{\tilde{\mu}(m),\tilde{x}} = \frac{\partial T\mu(m)}{\partial\mu(m)}|_{\tilde{\mu}(m),\tilde{x}} + \frac{\partial T\mu(m)}{\partial x^*}|_{\tilde{\mu}(m),\tilde{x}} \lim_{\mu(m)\to\tilde{\mu}(m)^+} \frac{\partial x^*}{\partial\mu(m)}.$$

The first term above equals to 1 and $\frac{\partial T\mu(m)}{\partial x^*}|_{\tilde{\mu}(m),\tilde{x}}$ is negative and bounded. Since $\tilde{x}$ is the unique solution to $H(x; \tilde{\mu}) = 0$, the term $\frac{d}{dx}H(\tilde{x}; \tilde{\mu}) = 0$, which leads to $\lim_{\mu(m)\to\tilde{\mu}(m)^+} \frac{\partial x^*}{\partial\mu(m)} = +\infty$. As a result, $\lim_{\mu(m)\to\tilde{\mu}(m)^+} \frac{dT\mu(m)}{d\mu(m)}|_{\tilde{\mu}(m),\tilde{x}} = -\infty$, so there is $\epsilon > 0$ such that for all $\mu(m) \in (\tilde{\mu}(m), \tilde{\mu}(m) + \epsilon)$, $T\mu(m) < \tilde{\mu}(m)$. Therefore, $\{\mu_t(m)\}_t$ cannot be a strictly decreasing sequence converging to $\tilde{\mu}(m)$. If given μ_0, there is a period T_4 such that $\mu_{T_4+1}(m) = \tilde{\mu}$, then from period T_4+1 the probability of attacking is constant at $\Phi(\frac{\sqrt{\beta}}{2}(s-m))$. Such μ_0 exists, because in the updating function, as $\mu(m)\to 1$, $T\mu(m)\to 1$, and $T\mu(m)$ is continuous in $\mu(m)$, there is some $\hat{\mu}(m)\in(\tilde{\mu}(m),1)$ such that $T\mu(m)=\tilde{\mu}(m)$. So for any finite T_4, μ_0 can be found. If given μ_0, there is no T_4 such that $\mu_{T_4+1}(m)=\tilde{\mu}$, then in any period t, if players attack the status quo with positive probability, the threshold point is bigger than $\tilde{x}=\frac{m+s}{2}$. Since $\{\mu_t(m)\}_t$ cannot be a strictly decreasing sequence converging to $\tilde{\mu}(m)$, there must be a period T_4 such that from period T_4+1, the probability of attacking is constant at 0.

Part 3: From the updating function, at $\mu(m)=\tilde{\mu}(m)$, when β is very large, $\left|\frac{dT\mu(m)}{d\mu(m)}\Big|_{\tilde{\mu}(m),\tilde{x}}\right|<1$. Therefore, the probability of attacking converges to $\Phi(\frac{\sqrt{\beta}}{2}(s-m))$. *Q.E.D.*

PROOF OF PROPOSITION 4.6: The proof of the case $x_2^*[\tilde{h}^2]>\frac{m+s}{2}$ is similar to that of Proposition 4.3, by noting that $\mu_t(m)$ is strictly decreasing over time since period 2.

Now consider the case $x_2^*[\tilde{h}^2]<\frac{m+s}{2}$. Since $\mu_t(m)$ is strictly increasing over time since period 2, and $x_t^*[\tilde{h}^t]$ is decreasing over time, then Bayes' rule implies $\mu_t(m)\to 1$. Because $H(x_t^*[\tilde{h}^t];\mu_t)=0$ for all t, $\Phi(\sqrt{\beta}(x_t^*-m))\to c$. *Q.E.D.*

PROOF OF PROPOSITION 4.7: 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 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 $t=1,2,\cdots$. 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. Define $\mathcal{V}_t$ 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 $\delta \in (\delta_w, 1)$, $\left|\sum_{t=1}^{\infty} \mathbb{P}_L(\kappa = t|\theta = w)\mathcal{V}_t - \sum_{t=1}^{\infty} \mathbb{P}_N(\kappa = t|\theta = w)\mathcal{V}_t\right| < \epsilon$. 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 Proposition 4.4. On one hand, because $\mathbb{P}_N(\text{regime changes}|\theta = m) = 1$, as δ goes to 1, $\mathcal{W}^N$ converges to 2. On the other hand, Proposition 4.4 implies that $\mathbb{P}_L(\text{regime changes}|\theta = m) < 1$, which in turns implies that $\mathcal{W}^L$ is strictly less than 2. Note that infinitely many attacks are prevented with or without social learning and that the discounted value of the cost from finite 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 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 less 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 $\left|\sum_{t=T}^{\infty} \mathbb{P}_L(\text{no attack after period } t|\theta = s)V_t^L\right| < \epsilon$. Because $\mathbb{P}_N(\text{attack, i.o.}|\theta = s) > 0$, for any T', there is a $\epsilon > 0$ such that $\left|\sum_{t=T}^{\infty} \mathbb{P}_N(\text{no attack after period } t|\theta = s)V_t^N\right| > \epsilon'$. 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.*

Lemma 7.5 *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 $H(\tilde{x};\tilde{\mu}) = 0$ (where $\tilde{\mu}(w) = 0$). Since $H(x;\tilde{\mu}) < 0$ for all $x \neq \tilde{x}$, $H'(\tilde{x};\tilde{\mu}) = 0$. As in Proposition 2.2, instead of studying $H(x;\tilde{\mu})$ directly, it is easier to study the function $h(x;\tilde{\mu}) = [\Phi(\sqrt{\beta}(x-m)) - c] - c(\frac{1}{\tilde{\mu}(m)} - 1)\exp(\frac{\beta}{2}(s-m)(2x-s-m))$. Since $\tilde{x}$ is also the unique solution to $h(x;\tilde{\mu}) = 0$, $h'(\tilde{x};\tilde{\mu}) = 0$. 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 $h'(\tilde{x};\tilde{\mu}) > 0$, which leads to the contradiction. As a result, as $\beta \to +\infty$, $\tilde{\mu}(m) \to 0$. ■

Lemma 7.6 *In a 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 $k = 1, 2, \cdots, N$. Then the interim payoff of player i when she observes private signal x^i and choose to attack is:

$$u^i(x^i|S^{-i}) = \frac{\sum_{k=1}^{N} \mu_k L_k(x^i) Z_k}{\sum_{k=1}^{N+1} \mu_k L_k(x^i)}.$$

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

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

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 $x^i > \bar{x}^i$.

■ *Q.E.D.*

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

[3] 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.

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

[5] Chamley, Christophe (1999): "Coordinating Regime Switches," *Quarterly Journal of Economics*, 114(3): 869-905.

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

[7] Dasgupta, Amil, Jakub Steiner and Colin Stewart (2009): "Efficient Dynamic Coordination with Individual Learning," Working Paper, Department of Economics, University of Toronto.

[8] Giannitsarou, Chryssi and Flavio Toxvaerd (2003): "Recursive Global Games," working paper, Universidade Nova de Lisboa and Hebrew University of Jerusalem.

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

[10] Heidhues, Paul and Nicolas Melissas (2006): "Equilibria in a Dynamic Global Game: The Role of Cohort Effects," *Economic Theory*, 28, 531-557

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

[12] ——— (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.

[13] ——— (2007): "Common Belief Foundations of Global Games," working paper, Princeton University.

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