

INFORMATION DYNAMICS AND MULTIPLICITY IN GLOBAL GAMES OF REGIME CHANGE*

George-Marios Angeletos
MIT and NBER

Christian Hellwig
UCLA

Alessandro Pavan
Northwestern University

October 2004

Abstract

Global games of regime change – that is, coordination games of incomplete information in which a “status quo” is abandoned once a sufficiently large fraction of agents attacks it – have been used to study crises phenomena such bank runs, currency attacks, debt crises, and political change. We extend the static benchmark examined in the literature by allowing agents to accumulate information over time and take actions in many periods. It is shown that dynamics may lead to multiple equilibria under the same information assumptions that guarantee uniqueness in the static benchmark. Multiplicity originates in the interaction between the arrival of information over time and the endogenous change in beliefs induced by the knowledge that the regime survived past attacks. This interaction also generates interesting equilibrium properties, such as the possibility that fundamentals predict the eventual regime outcome but not the timing or the number of attacks, or that dynamics alternate between crises and phases of tranquillity without changes in fundamentals.

JEL Codes: C7, D7, D8, F3.

Keywords: Global games, coordination, crises, multiple equilibria, information dynamics.

**Email addresses:* angelet@mit.edu, chris@econ.ucla.edu, alepavan@northwestern.edu.

1 Introduction

Coordination games in which a “status quo” is abandoned once a sufficiently large number of agents takes an action that favors “regime change” have been widely used to study a variety of socioeconomic phenomena. Prominent examples include currency attacks, bank runs, financial crashes, and debt crises, as well as political reforms, revolutions, and social change. Whereas the earlier contributions in the literature focused on the existence and implications of multiple equilibria when agents can perfectly coordinate with each other, recent work on global games by Morris and Shin (2000, 2001) and others has emphasized the fragility of this multiplicity to perturbations of the information structure away from common knowledge: a unique equilibrium often survives when agents observe a payoff-relevant variable, such as the strength of the status quo, with idiosyncratic noise.

Variants of this uniqueness result have appeared in a variety of applications. For example, Morris and Shin (1998) prove uniqueness in a model of self-fulfilling currency crises; Goldstein and Pauzner (2000) and Rochet and Vives (2003) consider bank runs; Morris and Shin (2003) and Corsetti, Guimaraes and Roubini (2004) study debt crises; Atkeson (2000) uses a riots example; and Edmond (2003) considers political change. These works assume a static game thus abstracting from the possibility that agents may have the option to take multiple shots against the regime. For most of the applications of interest, however, such a possibility seems to be important.

In this paper, we extend the static benchmark examined in the literature by allowing agents to take actions in multiple periods and accumulate information over time. There is a large number of agents and two possible regimes, the status quo and an alternative. In each period, each agent has a binary choice: he may either “attack” the status quo (that is, undertake an action that favors regime change) or refrain from attacking. The status quo is abandoned in any given period if and only if the fraction of agents attacking in that period exceeds a critical threshold θ , which parametrizes the strength of the status quo. θ represents what is commonly referred to in the literature as the “exogenous fundamentals” and is never common knowledge among the agents. Instead, as time passes, agents receive noisy private (and possibly public) signals about θ .

We show that dynamics may sustain multiple equilibria under the same information assumptions that guarantee uniqueness in the static benchmark, namely that the precision of the agents’ private information is sufficiently high relative to the precision of the initial common prior. Multiplicity originates in the interaction between two elements: the knowledge that the regime is still in

place and the arrival of information over time.

The presence of dominant strategies for extreme realizations of the agents' private information ensures that, in any equilibrium, an attack necessarily occurs in the first period of the game and immediately triggers regime change for sufficiently low θ . The observation that the status quo is still in place at any future date then makes it common certainty that it is strong enough to have survived at least one past attack. An unsuccessful attack in the first period thus causes an upward shift (a first-order stochastic-dominance change) in the agents' beliefs about θ . Other things equal, this effect makes further attacks impossible. Indeed, if no new information were to arrive, the equilibrium would be unique with an attack occurring only in the first period, in which case the possibility to take multiple shots against the regime would add nothing to the static analysis.

In contrast, when the agents' private information about θ becomes more precise over time, the influence on posterior beliefs of the knowledge that the regime survived past attacks eventually vanishes. Whether this effect inhibits or favors further attacks however depends on the initial common prior. When the prior mean is sufficiently low, an increase in the precision of private information would make the agents less aggressive even in the static game in which agents do not have the possibility to move in the past. In this case, the arrival of new information contributes in making further attacks impossible after the first period. When instead the prior mean is high so that the size of the attack in the first period is small, the arrival of private information makes new attacks eventually possible by increasing the agents' confidence that the regime may collapse when attacked by a sufficiently large fraction of the population. At the same time, the knowledge that the regime already survived one or more attacks, makes no agent willing to attack in subsequent periods if he expects few other agents to attack again, implying that there also exist equilibria in which attacks cease at any point in time after the first period. A sufficiently high prior mean thus suffices for multiple equilibria to exist.

In the benchmark model, we deliberately assume away the possibility that "fundamentals" change over time. This allows us to isolate the effects of changes in information, as opposed to changes in fundamentals, on the dynamics of coordination. Nevertheless, we also show that the multiplicity result is robust to the introduction of shocks to the fundamentals, in which case dominant actions may exist in every period and for every history of the game, provided that the volatility of these shocks is sufficiently small. In the limit, as the volatility vanishes, any equilibrium in the model with stationary fundamentals can be approximated arbitrarily closely by an equilibrium in the game with shocks. What sustains the multiplicity of equilibria and the

corresponding dynamics is again the combination of the arrival of new information over time with the endogenous pronounced change in posterior beliefs that follows from the fact that the regime survived past attacks. That this change takes the form of a truncation as in the benchmark model simplifies the construction of the equilibrium set, but is not essential for the result.

We also emphasize that multiplicity does not originate in the presence of exogenous public signals or in the observation of the size of past attacks. Indeed, for most of the analysis we assume away any such source of information and concentrate on the case in which agents receive only private signals over time. What the introduction of public news does, is to “smooth out” the strong impact that the initial prior otherwise has on equilibrium dynamics: multiple equilibria then exist for any prior mean and for any relative precision of public and private information. Finally, introducing endogenous signals about the size of past attacks – a form of (noisy) social learning – does not change the equilibrium dynamics, except for the possibility that the endogenous information revealed by the size of an attack may substitute for the arrival of new exogenous information and render further attacks possible immediately after an unsuccessful one.

Finally, it is interesting to note that the existence of multiple equilibria limits, but does not eliminate, the possibility to make predictions. What is more, it may lead to novel properties for equilibrium dynamics. For example, fundamentals may predict the eventual fate of the regime but not the timing of attacks. This seems to be consistent with the view that economic fundamentals may determine the final regime outcome (e.g., whether a currency is eventually devalued), but not when a crisis may occur or when attacks cease. On the contrary, this view is inconsistent with most common-knowledge models, in which fundamentals fail to predict both the timing of attacks and the eventual regime outcome, as well as with unique-equilibrium models like Morris and Shin (1998, 1999), in which both the timing of attacks and the ultimate fate of the regime are uniquely pinned down by fundamentals. Second, dynamics may take the form of cycles in which an initial attack is followed by a phase in which further attacks can not occur and agents only accumulate information, followed by a phase in which an attack is possible but does not take place, eventually resulting in a new attack. These predictions may help understand how an economy may transit from phases of tranquility to crises *without* any change in the underlying economic fundamentals.

Below, we discuss the relation of the paper to the pertinent literature. Section 2 then introduces the model, reviews the static benchmark examined in the literature and introduces dynamics. Section 3 characterizes the set of monotone equilibria for the dynamic game. Section 4 establishes

the multiplicity result and discusses a few properties for equilibrium dynamics. Section 5 examines the robustness of the results to the introduction of shocks in the fundamentals. Section 6 considers the effect of public news and (noisy) social learning. Finally, Section 7 concludes. Proofs omitted in the main text are presented in the Appendix.

Related Literature

This paper contributes to a small, but rapidly growing, literature on dynamic global games. Most of this literature either attempts at generalizing Carlsson and van Damme's uniqueness result to dynamic games, or considers applications of the static global games results to dynamic problems. Levin (2001) considers a global game with overlapping generations of players. Heidhues and Melissas (2003) study a dynamic investment model and identify a condition of *dynamic strategic complementarity* sufficient for uniqueness. This condition is also implicit in Giannitsarou and Toxvaerd's (2003) general uniqueness theorem for dynamic binary action recursive games with finite horizon. Dasgupta (2002) examines the role of social learning in a two-period investment model with irreversible actions. Goldstein and Pauzner (2001), and Goldstein (2002) consider models of contagion.

Closest to our analysis, Morris and Shin (1999) examine a dynamic global game whose stage game is similar to ours, but where fundamentals change over time and where past fundamentals are commonly observed after each period. This reduces the dynamic game to a sequence of one-period games with a unique equilibrium, in which dynamics are driven by changes in fundamentals. A unique equilibrium model with stochastic fundamentals is also examined in Chamley (1999).

This paper departs from the above literature in that it considers dynamics as a natural source of common beliefs. In this respect, it shares with Angeletos, Hellwig and Pavan (2003)¹ the idea that *endogenous* information structures in global games may overturn uniqueness results and lead to predictions that would have not been possible under common knowledge or in unique equilibrium models.

¹That paper examines the signaling effects of active policy interventions in a static global game.

2 A simple game of regime change

In this Section, we first review the static benchmark that has been studied in the literature and then present the dynamic game.

2.1 Static benchmark

Model set-up. There is a continuum of agents of measure one, indexed by i and uniformly distributed over $[0, 1]$. Agents move simultaneously, choosing between two actions: they can either attack the status quo (i.e., take an action that favors regime change) or refrain from attacking. The payoff structure is illustrated in Table 1: the payoff from not attacking ($a_i = 0$) is normalized to 0, whereas the payoff from attacking ($a_i = 1$) is $1 - c > 0$ in case there is regime change ($R = 1$) and $-c < 0$ in case the status quo is maintained ($R = 0$), where $c \in (0, 1)$ parametrizes the relative cost of attacking.

	<i>Regime Change</i> ($R = 1$)	<i>Status Quo</i> ($R = 0$)
<i>Attack</i> ($a_i = 1$)	$1 - c$	$-c$
<i>Not Attack</i> ($a_i = 0$)	0	0

Table 1. Payoffs

Individual payoff are thus $\pi_i = a_i(R - c)$ and hence each agent finds it optimal to attack if and only if his belief about the probability of regime change exceeds c . The status quo is in turn abandoned if and only if the measure of agents attacking, which we denote by A , is greater or equal to a critical threshold $\theta \in \mathbb{R}$, which parametrizes the strength of the status quo.

Coordination and information. The role of coordination is most evident when θ is commonly known by all agents. When $\theta \notin (0, 1]$, each agent finds it dominant either to attack (for $\theta \leq 0$) or not attack (for $\theta > 1$) and therefore the equilibrium is unique. When instead $\theta \in (0, 1]$, there exist two pure-strategy equilibria: in the first, each agent expects all other agents to attack ($A = 1 \geq \theta$) and hence finds it optimal to attack, in which case the status quo is abandoned; in the second, each agent expects no one to attack ($A = 0 < \theta$) and hence also finds it optimal to refrain from attacking, in which case the status quo is maintained.

However, the recent work on global games, pioneered by Carlsson and van Damme (1993) and Morris and Shin (1998, 2000, 2001), has highlighted the fragility of multiple equilibria to

perturbations of the information structure that remove common knowledge. Following Morris and Shin (2000), suppose that nature first draws θ from a normal distribution $\mathcal{N}(z, 1/\beta)$, which defines the initial common prior about θ , and then each agent receives a private signal $x_i = \theta + \xi_i$, where $\xi_i \sim \mathcal{N}(0, 1/\alpha)$ is i.i.d. across agents and independent of θ . The coefficients α and β parametrize, respectively, the precisions of private information and of the common prior. Then, a unique equilibrium survives when α is sufficiently high relative to β , as we show next.

Equilibrium analysis. Note that the posterior of an agent about θ is decreasing in his private signal x . Moreover, there exist \underline{x} and \bar{x} such that $\Pr(\theta \leq 0|x) > c$ for all $x < \underline{x}$ and $\Pr(\theta \leq 1|x) < c$ for all $x > \bar{x}$. That is, agents with signal $x < \underline{x}$ find it strictly dominant to attack, whereas agents with signal $x \leq \underline{x}$ find it strictly dominant to refrain from attacking. It is thus natural to look at *monotone* Bayesian Nash equilibria in which the agents' strategy is non-increasing in the private signal. Thus, suppose there is a threshold \hat{x} such that each agent attacks if and only if $x \leq \hat{x}$. The measure of agents attacking is then decreasing in θ and is given by $A(\theta) = \Pr(x \leq \hat{x}|\theta) = \Phi(\sqrt{\alpha}[\hat{x} - \theta])$, where Φ is the c.d.f. of the standard Normal. It follows that the status quo is abandoned if and only if $\theta \leq \hat{\theta}$, where $\hat{\theta}$ solves

$$\hat{\theta} = \Phi(\sqrt{\alpha}(\hat{x} - \hat{\theta})). \tag{1}$$

The posterior probability of regime change for an agent with signal x is then simply $\Pr(R = 1|x) = \Pr(\theta \leq \hat{\theta}|x)$. Since the latter is decreasing in x , each agent finds it optimal to attack if and only if $x \leq \hat{x}$, where \hat{x} solves.

$$\Pr(\theta \leq \hat{\theta}|\hat{x}) = c. \tag{2}$$

A monotone equilibrium is thus identified by a joint solution $(\hat{x}, \hat{\theta})$ to (2) and (1). Such a solution always exists and is unique for all z if and only if $\alpha \geq \beta^2/(2\pi)$. Moreover, iterated elimination of strictly dominated strategies implies that, when the monotone equilibrium is unique, this is also the unique equilibrium of the game. We conclude that

Proposition 1 (Static benchmark) *In the static game, the equilibrium is unique if and only if $\alpha \geq \beta^2/(2\pi)$ and it is in monotone strategies.*

Interpretation. Although the game presented above is highly stylized, it admits a variety of interpretations and possible applications. The most celebrated examples are self-fulfilling *bank runs*, *currency attacks*, and *debt crises*. In these contexts, “regime change” occurs, respectively,

when a large run forces the banking system to suspend its payments, when a large speculative attack forces the central bank to abandon the peg, or when a country/company fails to coordinate its creditors and is forced to bankruptcy.

The model can also be interpreted as one of *political change*, in which a large number of “citizens” decide whether or not to take actions to subvert a repressive dictator or some other political establishment. Atkeson (2000) interprets it as a game of riots, in which the potential rioters face uncertainty about whether they will be able to overwhelm the police force which is in charge of containing social unrest. Edmond (2003) uses a similar model to discuss how the IT revolution influences the stability of autocratic regimes.

2.2 Dynamic game

We modify the static game reviewed above in two ways: first, we allow agents to attack the status quo repeatedly; second, we let agents accumulate information over time.

Model set-up. Time is discrete and indexed by $t \in \{1, 2, \dots\}$. We denote with $R_t = 0$ the event that the status quo is still in place at the beginning of period t , with $R_t = 1$ the alternative event, with $a_{it} \in \{0, 1\}$ the action of agent i at date t , and with $A_t \in [0, 1]$ the measure of agents attacking at date t . The game continues as long as the status quo is in place and it is over once the status quo is abandoned. Conditional on $R_t = 0$, the regime is abandoned in period t ($R_{t+1} = 1$) if and only if $A_t \geq \theta$. Each agent’s payoff in period t is

$$\pi_{it} = a_{it}(R_{t+1} - c),$$

and the payoff in the entire game is $\Pi_i = \sum_{t=1}^{\infty} \rho^{t-1}(1 - R_t)\pi_{it}$, where $\rho \in (0, 1)$ is the common discount factor.

Like in the static model, the critical threshold $\theta \in \mathbb{R}$ represents the strength of the status quo and is never common knowledge among the agents. Unlike the static model, private information evolves over time. Nature continues to draw θ at the beginning of the game from a Normal distribution $\mathcal{N}(z, 1/\beta)$, which defines the initial common prior. Then, in each period $t \geq 1$, every agent i receives a private signal $\tilde{x}_{it} = \theta + \xi_{it}$ about θ , where $\xi_{it} \sim \mathcal{N}(0, 1/\eta_t)$ is i.i.d. across i , independent of θ , and serially uncorrelated. Let $\tilde{x}_i^t = \{\tilde{x}_{i\tau}\}_{\tau=1}^t$ denote agent i ’s history of private signals up to period t . Individual actions and the size of past attacks are *not* observable,² hence

²The possibility that agents observe noisy signals about aggregate past activity is considered in Section ??.

the public history in period t simply consists of the information that the regime is still in place (that is, $R_t = 0$), whereas the private history of an agent is the sequence of own private signals \tilde{x}_i^t and own past actions. Finally, we let $\alpha_t \equiv \sum_{\tau=1}^t \eta_\tau$ and assume that

$$\alpha_1 \geq \beta^2/(2\pi), \quad \alpha_t < \infty \quad \text{for all } t, \quad \text{and} \quad \lim_{t \rightarrow \infty} \alpha_t = \infty.$$

As shown in the next section, α_t parametrizes the precision of private information in period t . The restrictions above ensure that the equilibrium would be unique if the game were static and that information becomes infinitely precise only in the limit.

While this game is highly stylized, it captures two important dimensions that are absent in the static benchmark: first, the possibility of multiple attacks against the status quo; and second, the evolution of beliefs about the probability of regime change. By assuming that per-period payoffs do not depend on past or future actions and by ignoring specific institutional details, the model may of course fail to capture other interesting effects introduced by dynamics, such as the role of wealth accumulation or illiquidity in currency crises. However, abstracting from these other dimensions, allows us to isolate information dynamics as the driving force for the dynamics of coordination.

Equilibrium. In what follows, we limit attention to symmetric Perfect Bayesian equilibria in which the probability an agent attacks in period t , $a_t(\tilde{x}^t)$, is independent of his own past actions and non-increasing in his private signals \tilde{x}_i^t , for all t .³ While this restriction may not be without loss of generality, it suffices to look at this class of equilibria to establish our multiplicity result, as well as to identify interesting properties for the dynamics of attacks. We henceforth refer to this class of equilibria as *monotone equilibria*.

3 Equilibrium characterization

As long as the status quo is in place, no agent can detect out-of-equilibrium play. This simply follows from the fact that neither individual nor aggregate actions are observable, and that $R_t = 0$ is compatible with any strategy profile.⁴ Furthermore, payoffs in one period do not depend on own or other players' actions in any other period. By implication, beliefs are pinned down by Bayes' rule, continuation payoffs do not depend on current own actions, and strategies are sequentially rational

³Note that we do not restrict in any way the set of available strategies; we only look for equilibria in which the equilibrium strategies satisfy these properties.

⁴The latter in turn is true because the regime necessarily survives for $\theta > 1$ and no private signal rules out $\theta > 1$.

if and only if they maximize period-by-period payoffs. This also implies that in *any* equilibrium of the dynamic game, agents follow in period one exactly the same strategy as in the static game in which they can attack only at $t = 1$.

Next, note that $\theta|\tilde{x}^t \sim \mathcal{N}(x_t, \alpha_t^{-1})$, where x_t and α_t are defined recursively as follows:

$$x_t = \frac{\alpha_{t-1}}{\alpha_t}x_{t-1} + \frac{\eta_t}{\alpha_t}\tilde{x}_t \quad \text{and} \quad \alpha_t = \alpha_{t-1} + \eta_t.$$

x_t is thus a sufficient statistic for the history of private signals \tilde{x}^t with respect to θ and therefore with respect to the probability of regime change as well. It follows that *any* perfect Bayesian equilibrium can be described by a sequence $\{a_t(\cdot)\}_{t=1}^{\infty}$ such that at any t

$$a_t(\tilde{x}^t) \in \arg \max_{a \in [0,1]} \{(\Pr[R_{t+1} = 1|x_t, R_t = 0] - c) a\}, \quad (3)$$

and where $\Pr[\cdot]$ is consistent with Bayes' rule, given $\{a_t(\cdot)\}_{t=1}^{\infty}$. The next Lemma then shows that monotone equilibria can be described by sequences of thresholds $\{x_t^*, \theta_t^*\}_{t=1}^{\infty}$ for the sufficient statistics of the agents and the fundamentals, in a fashion similar to the static benchmark. To simplify the notation, we henceforth allow for $x_t^* = -\infty$ and $x_t^* = +\infty$, by which we denote the case where an agent attacks for, respectively, none or every realization of his private information, and let $\overline{\mathbb{R}} = \mathbb{R} \cup \{\pm\infty\}$.

Lemma 1 *In any monotone equilibrium, there is a sequence $\{x_t^*, \theta_t^*\}_{t=1}^{\infty}$ such that: (i) an agent finds it optimal to attack in period t if $x_t < x_t^*$ and not to attack if $x_t > x_t^*$; (ii) the status quo is in place in period $t + 1$ ($R_{t+1} = 0$) if and only if $\theta > \theta_t^*$, where $\theta_t^* \in (0, 1)$.*

Since $\theta_{t-1}^* \geq \theta_1^* > 0$ for any $t \geq 2$, the fact that the status quo is in place in period t makes common certainty among agents that $\theta \geq \theta_{t-1}^* > 0$. This implies that there always exist continuation equilibria in which nobody attacks in period $t \geq 2$, in which case $x_t^* = -\infty$ and $\theta_t^* = \theta_{t-1}^*$. In particular, there exists an equilibrium in which an attack takes place in period 1 and never after. If this were the unique equilibrium, the possibility to take repeated actions against the regime would add nothing to the static analysis and the equilibrium outcome in the dynamic game would coincide with the static one. In what follows, we thus examine under what conditions there also exist equilibria with further attacks.

First, note that from Lemma 1, the size of the attack is decreasing in θ and is given by $A_t(\theta) = \Pr(x \leq x_t^*|\theta) = \Phi(\sqrt{\alpha_t}(x_t^* - \theta))$. It follows that, in any equilibrium in which an attack

occurs in period t , $A_t(\theta) \geq \theta$ if and only if $\theta \leq \theta_t^*$, where θ_t^* solves

$$\theta_t^* = \Phi(\sqrt{\alpha_t}(x_t^* - \theta)). \quad (4)$$

The probability of regime change is then simply $\Pr(R_{t+1} = 1 | x_t, R_t = 0) = \Pr(\theta \leq \theta_t^* | x_t, \theta > \theta_{t-1}^*)$ and is decreasing in x_t . It follows that x_t^* solves

$$\Pr(\theta \leq \theta_t^* | x_t^*, \theta > \theta_{t-1}^*) = c. \quad (5)$$

Note that (5) and (4) are the analogs in the dynamic game of conditions (2) and (1) in the static game. They describe two equations in two unknowns (x_t^*, θ_t^*) which have to be satisfied by any equilibrium in which there is an attack in period t . (5) states that if the regime is still in place in period t , that is $\theta > \theta_{t-1}^*$, an agent is indifferent between attacking and not attacking if and only if his private information is x_t^* ; (4) states when all agents follow a monotone strategy with cut-off x_t^* , the equilibrium size of an attack is equal to the critical size that triggers a regime change exactly when the fundamentals are θ_t^* .

Let $X(\theta^*; \alpha) \equiv \theta^* + \alpha^{-1/2}\Phi^{-1}(\theta^*)$ be the solution to (4) for given $\theta^* \in [0, 1]$ and $\alpha > 0$, that is, the threshold such that, if agents attack if and only if $x \leq X(\theta^*; \alpha)$, regime change occurs if and only if $\theta \leq \theta^*$.⁵ We refer to $X(\theta^*; \alpha)$ as the ‘‘marginal agent that implements regime change for $\theta \leq \theta^*$.’’ Next, define the function $U : [0, 1] \times \overline{\mathbb{R}} \times \mathbb{R}_+^2 \times \mathbb{R} \rightarrow [-c, 1 - c]$ by

$$U(\theta^*; \theta_{-1}, \alpha, \beta, z) \equiv \begin{cases} 1 - \Phi\left(\frac{\sqrt{\alpha}}{\sqrt{\alpha+\beta}}\left[\Phi^{-1}(\theta^*) + \frac{\beta}{\sqrt{\alpha}}(z - \theta^*)\right]\right) - c & \text{if } \theta_{-1} = -\infty \\ 1 - \frac{\Phi\left(\frac{\sqrt{\alpha}}{\sqrt{\alpha+\beta}}\left[\Phi^{-1}(\theta^*) + \frac{\beta}{\sqrt{\alpha}}(z - \theta^*)\right]\right)}{\Phi\left(\frac{\sqrt{\alpha}}{\sqrt{\alpha+\beta}}\left[\Phi^{-1}(\theta^*) + \frac{\beta}{\sqrt{\alpha}}(z - \theta^*)\right] + \sqrt{\alpha+\beta}(\theta^* - \theta_{-1})\right)} - c & \text{if } \theta^* > \theta_{-1} > -\infty \\ -c & \text{if } \theta^* \leq \theta_{-1} \end{cases} \quad (6)$$

$U(\theta^*; \theta_{-1}, \alpha, \beta, z)$ represents the expected net payoff from attacking for the marginal agent that implements regime change for $\theta \leq \theta^*$, conditional on the knowledge that $\theta > \theta_{-1}$, for exogenous information structure given by (α, β, z) . The next theorem then provides a complete characterization of the set of monotone equilibria.

⁵With a slight abuse of notation, we let $\Phi(+\infty) = 1$, $\Phi(-\infty) = 0$, $\Phi^{-1}(1) = \infty$ and $\Phi^{-1}(0) = -\infty$.

Theorem 1 (Equilibrium characterization) $\{a_t(\cdot)\}_{t=1}^\infty$ is a monotone equilibrium if and only if there exists a sequence $\{x_t^*, \theta_t^*\}_{t=1}^\infty$ such that:

- (i) for all t , $a_t(\cdot) = 1$ if $x_t < x_t^*$ and $a_t(\cdot) = 0$ if $x_t > x_t^*$;
- (ii) for $t = 1$, θ_1^* solves

$$U(\theta_1^*; -\infty, \alpha_1, \beta, z) = 0 \quad (7)$$

and $x_1^* = X(\theta_1^*, \alpha_1)$.

(iii) for any other $t \geq 2$, either $\theta_t^* = \theta_{t-1}^* > 0$, in which case $x_t^* = -\infty$, or $\theta_t^* > \theta_{t-1}^*$ is a solution to

$$U(\theta_t^*; \theta_{t-1}^*, \alpha_t, \beta, z) = 0, \quad (8)$$

in which case $x_t^* = X(\theta_t^*, \alpha_t)$.

The above characterization is independent of whether the horizon is finite or infinite: it is clearly valid even if the game ends exogenously at an arbitrary period $T < \infty$. It also provides a simple algorithm for constructing the entire equilibrium set. Start with $t = 1$ and let θ_1^* be the unique solution to (7). Proceed to period $t = 2$ and let either $\theta_2^* = \theta_1^*$ or θ_2^* be a solution to (8), provided the latter exists. Repeat the same step for all $t \geq 3$. The set of sequences $\{\theta_t^*\}_{t=1}^\infty$ constructed this way, together with the associated sequences $\{x_t^*\}_{t=1}^\infty$, identifies the set of monotone equilibria.

The next lemma completes the characterization of monotone equilibria by examining under what conditions equation (8) admits a solution.

Lemma 2 (i) $U(\theta^*; \theta_{-1}, \alpha, \beta, z)$ is continuous in all its arguments, non-monotonic in θ^* when $\theta_{-1} \in (0, 1)$, and strictly decreasing in θ_{-1} and z . Furthermore, for all $\theta_{-1} < 1$ and $\theta^* > \theta_{-1}$, $\lim_{\alpha \rightarrow \infty} U(\theta^*; \theta_{-1}, \cdot) = \theta_\infty - \theta^*$, where $\theta_\infty \equiv 1 - c$.

(ii) Let $\hat{\theta}_t$ be the unique solution to $U(\hat{\theta}_t, -\infty; \alpha_t, \beta, z) = 0$. A solution to (8) exists only if $\theta_{t-1}^* < \hat{\theta}_t$ and is necessarily bounded from above by $\hat{\theta}_t$.

(iii) If $\theta_{t-1}^* > \theta_\infty$, a solution to (8) does not exist for α_t sufficiently high.

(iv) If $\theta_{t-1}^* < \theta_\infty$, a solution to (8) necessarily exists for α_t sufficiently high.

(v) If θ_t^* is the highest solution to (8) for period t , then there exists $\underline{\alpha} > \alpha_t$ such that (8) admits no solution for any period $\tau > t$ such that $\alpha_\tau < \underline{\alpha}$.

The non-monotonicity of $U(\theta^*, \theta_{t-1}^*, \cdot)$ with respect to θ^* in any period $t \geq 2$ (where $\theta_{t-1}^* \in (0, 1)$) is a direct consequence of the fact that for $\theta^* < \theta_{t-1}^*$, the marginal agent attaches probability

zero to regime change. For $\theta^* > \theta_{t-1}^*$, on the other hand, his belief about regime change is necessarily positive, but converges to zero as $\theta^* \rightarrow 1$, for then $X(\theta^*, \alpha) \rightarrow \infty$ and hence he attaches probability one to $\theta > 1$. His payoff is thus initially increasing and eventually decreasing, as illustrated by the solid line in Figure 1.⁶ Note how this contrasts with the payoff in the static game, which is represented by the dotted line. Whereas the monotonicity of the payoff of the marginal agent in the static game ensures uniqueness, the non-monotonicity in any period $t \geq 2$ of the dynamic game leaves open the possibility for multiple equilibria.

[Insert Figure 1 here]

The other properties identified in Lemma 2 are also quite intuitive. An increase in the mean of the common prior implies a first-order stochastic-dominance change in the posterior beliefs about θ , which explains why U decreases with z . To understand why U is also decreasing in θ_{-1} , note that the more aggressive the attacks in the past, the stronger the regime must have been in order to have survived these attacks, and therefore the lower the expected payoff from attacking in the current period. For the same reason, at any $t \geq 2$, the payoff of the marginal agent is always lower than the payoff $U(\theta^*, -\infty, \alpha_t, \beta, z)$ in the corresponding static game, which explains why the static-game threshold $\hat{\theta}_t$ is an upper bound for any solution to (8). Nevertheless, the impact of the information that $\theta > \theta_{t-1}^*$ on the posterior of the marginal agent who implements any $\theta^* > \theta_{t-1}^*$ vanishes as the precision of his private information increases,⁷ and hence, for any $\theta^* > \theta_{t-1}^*$, the difference between $U(\theta^*, \theta_{t-1}^*, \alpha_t, \beta, z)$ and $U(\theta^*, -\infty, \alpha_t, \beta, z)$ vanishes as $\alpha_t \rightarrow \infty$. Combined with the fact that $U(\theta^*, -\infty, \alpha_t, \beta, z) \rightarrow \theta_\infty - \theta^*$ as $\alpha_t \rightarrow \infty$, this in turn implies that, for α_t sufficiently high, (8) admits a solution if and only if $\theta_{t-1}^* < \theta_\infty$.

Finally, to understand (v), note that any unsuccessful attack causes an upward shift (a first-order stochastic-dominance change) in the posterior beliefs about the strength of the regime, which other things equal reduces the expected payoff from attacking. It follows that, if the largest possible attack (that is, the highest solution of (8)) is played in one period and no new information arrives thereafter, no attack is possible in any subsequent period. By continuity then, further attacks remain impossible as long as the change in the precision of private information is not large enough.

⁶The proof that U is strictly quasi concave is available upon request.

⁷This follows from the fact that, for any $x > \theta_{-1}$ and any θ' , $\Pr(\theta \leq \theta'|x, \theta > \theta_{-1}) - \Pr(\theta \leq \theta'|x) \rightarrow 0$ as $\alpha \rightarrow \infty$.

4 Multiplicity and Dynamics

Part (v) of Lemma 2 highlights that the arrival of new private information is necessary for further attacks to become possible after period 1. Whether this is also sufficient depends on the initial prior. This is because an increase in the precision of private information leads agents to “discount”, not only the information conveyed by the fact that the regime survived past attacks, but also their initial prior beliefs.

When the prior mean z is low, an increase in the precision of information would make the agents less aggressive even in the static game. The knowledge that the regime survived past attacks then only reinforces this effect. Therefore, for z sufficiently low, there exists a unique equilibrium, in which either the regime is abandoned in the first period or it survives forever.

When instead z is high, the arrival of new private information makes agents more aggressive, and may eventually offset the incentives not to attack induced by the knowledge that the regime survived past attacks. Indeed, if z is high enough so that $\theta_1^* < \theta_\infty$, Lemma 2 ensures that a second attack necessarily becomes possible once α_t is sufficiently high. If this attack leads to a threshold $\theta_t^* < \theta_\infty$, a third attack becomes possible at some future date, and so on. Hence, there also exist equilibria with multiple attacks.

Theorem 2 (Multiplicity) *There exist thresholds $\underline{z} \leq \bar{z} \leq \bar{\bar{z}}$ such that:*

- (i) *If $z \leq \underline{z}$, the equilibrium is **unique** and an attack occurs only in period 1.*
- (ii) *If $z \in (\underline{z}, \bar{z})$, there are at most finitely many equilibria; moreover, there exist T such that, in any equilibrium, no attack occurs after period T .*
- (iii) *If $z > \bar{z}$, there are **infinitely many** equilibria; if in addition $z > \bar{\bar{z}}$, for any t and N , there is an equilibrium in which N attacks occur after period t .*

Finally, $\underline{z} = \bar{z} = \bar{\bar{z}}$ when $c \leq 1/2$, whereas $\underline{z} \leq \bar{z} < \bar{\bar{z}}$ when $c > 1/2$.

Proof. Let $\hat{\theta}_t = \hat{\theta}(\alpha_t, z, \beta)$ be the unique solution to $U(\hat{\theta}; -\infty, z, \alpha_t, \beta) = 0$, that is, the equilibrium threshold of the static game corresponding to $\alpha = \alpha_t$. Recall that the first period is isomorphic to the static game and therefore $\theta_1^* = \hat{\theta}_1$. Moreover, since the knowledge that the regime survived past attacks reduces the payoff from attacking relative to the static game, any solution to (8) is bounded from above by $\hat{\theta}_t$ and therefore $\theta_t^* < \hat{\theta}_t$ for $t \geq 2$. Next, as we show in Lemma A1 in the Appendix, there exist thresholds $\underline{z} \leq \bar{z} \leq \bar{\bar{z}}$ with the following properties: $\hat{\theta}_t \leq \hat{\theta}_1$ for all t if $z \leq \underline{z}$; $\hat{\theta}_1 \leq (\geq) \theta_\infty$ if and only if $z \geq (\leq) \bar{z}$; and $\hat{\theta}_t < \theta_\infty$ for all t if and only if $z > \bar{\bar{z}}$.

(i) Consider first $z \leq \underline{z}$. Then, $\hat{\theta}_t \leq \hat{\theta}_1 = \theta_1^*$ for all t , and hence, by Lemma 2, (8) admits no solution at any $t \geq 2$. The unique equilibrium is thus $\theta_t^* = \theta_1^*$ for all t .

(ii) Next, consider $z \in (\underline{z}, \bar{z})$, in which case $\max_{t \geq 1} \hat{\theta}_t > \hat{\theta}_1 = \theta_1^* > \theta_\infty$. Then, we can not rule out the possibility that there exists a period $t \geq 2$ such that $U(\theta; \theta_1^*, z, \alpha_t, \beta) = 0$ admits a solution. Nevertheless, since $\hat{\theta}_t \rightarrow \theta_\infty$ as $t \rightarrow \infty$, there exists T such that $\hat{\theta}_t < \theta_1^* \leq \theta_{t-1}^*$ for all $t \geq T$. It follows that (8) admits no solution for $t \geq T$. Moreover, U is strictly quasi-concave in $\theta^* > \theta_{-1}$, which implies that (8) has at most two solutions for $t < T$. Hence, there are at most finitely many equilibria, and in any equilibrium no attack occurs after period T .

(iii) Finally, consider $z > \bar{z}$, in which case $\theta_1^* < \theta_\infty$. Then, by Lemma 2, there exists a $T < \infty$ such that $U(\theta; \theta_1^*, z, \alpha_t, \beta) = 0$ admits a solution for all $t \geq T$. Hence, for any $t \geq T$, there is an equilibrium in which $\theta_\tau^* = \theta_1^*$ for $\tau < t$, θ_t^* solves $U(\theta_t^*; \theta_1^*, z, \alpha_t, \beta) = 0$, and $\theta_\tau^* = \theta_t^*$ for all $\tau > t$. That is, there exist (countably) infinitely many equilibria, indexed by the time at which the second attack occurs.

When $z \in (\bar{z}, \bar{\bar{z}})$, the second attack may lead to a threshold $\theta_t^* > \theta_\infty$, in which case a third attack might be impossible. If however $z > \bar{\bar{z}}$, then $\hat{\theta}_t < \theta_\infty$ for all t , and therefore any attack leads to a threshold $\theta_t^* < \theta_\infty$, thus eventually permitting further attacks. But then by Lemma 2, for any $t \geq 1$ and any $N \geq 1$, there exist sequences $\{t_2, \dots, t_N\}$ and $\{\theta_2, \dots, \theta_N\}$ such that $U(\theta_2; \theta_1^*, z, \alpha_{t_2}, \beta) = 0$, $U(\theta_3; \theta_2, z, \alpha_{t_3}, \beta) = 0$, and so on. The following is then an equilibrium: $\theta_t^* = \theta_1^*$ for $t < t_2$, $\theta_t^* = \theta_j$ for $t \in \{t_j, \dots, t_{j+1} - 1\}$ and $j \in \{2, \dots, N - 1\}$, and $\theta_t^* = \theta_N$ for $t \geq t_N$. That is, for any $t \geq 1$ and any $N \geq 1$, there exists an equilibrium in which N attacks occur after date t . ■

The existence of infinitely many equilibria in the case $z > \bar{z}$ relies on the assumption that the game continues for ever as long as the status quo is in place: if the game ended at a finite $T < \infty$, there would be only finitely many equilibria. Nevertheless, as long as $z > \bar{z}$, for any M , there exists a finite T such that the game with horizon T has at least M equilibria. Moreover, even if $T = 2$, α_2 high enough suffices for the game to have multiple equilibria when $z > \bar{z}$.

We note that the prior mean has a strong impact on equilibrium multiplicity relies on the assumption that the exogenous public information does not change over time. As we show in Section 6, the introduction of public signals naturally “smooths out” the role of the initial prior and ensures the existence of multiple equilibria possible for *any* level of the prior mean and *any* precision of private and public information.

In the remainder of this section, we discuss a few interesting properties of the equilibrium set.

The existence of multiple equilibria limits, but does not necessarily eliminate, the predictive power of fundamentals: if $z > \bar{z}$, there are infinitely many equilibria, nevertheless the regime outcome is unique for all $\theta \notin (\theta_1^*, \theta_\infty)$. Interestingly, even when fundamentals predict final outcomes, they may however fail to predict both the timing and the number of attacks.

Corollary 1 *Suppose $\theta > \theta_\infty$ and $z > \bar{z}$. The status quo survives in any equilibrium. Nevertheless, there exists \underline{t} such that an arbitrary number of attacks may occur starting at any date $t \geq \underline{t}$.*

This seems to square well with the common view that economic fundamentals may predict the final regime outcome (e.g., whether a currency is eventually devalued) but not when a crisis may occur or when attacks cease. On the contrary, it is important to note that this view is inconsistent with the common-knowledge version of our model, in which case fundamentals fail to predict both the timing of attacks and the eventual regime outcome,⁸ as well as with a unique-equilibrium model like Morris and Shin (1998, 1999), in which both the timing of attacks and the ultimate fate of the regime are uniquely pinned down.

Another interesting property of our equilibria relates to how information affects the possibility of consecutive attacks. Recall that an unsuccessful attack causes a first-order stochastic-dominance change in the agents' posterior beliefs about θ . In other words, the regime “builds up reputation” by surviving attacks. This in turn implies that

Corollary 2 *After the most aggressive attack for a given period occurs, the game may enter a phase of tranquillity, during which no attack is possible. This phase is longer the slower the arrival of private information.*

In the model we have examined so far, the game may exit a phase of tranquillity only via an exogenous increase in α . Once we introduce public signals, a new attack may be triggered by the arrival of sufficiently bad public news. Alternatively, social learning, by allowing agents to receive noisy private or public signals about the size of past attacks, may lead to endogenous changes in (α, β, z) . Nevertheless, an attack remains impossible as long as the change in (α, β, z) is small enough.

This property may generate interesting dynamics. For example, when $\theta > \theta_\infty$ and $z > \bar{z}$, dynamics may take the form of cycles in which an initial attack is followed by a phase of tranquillity,

⁸Of course, this is true provided $\theta \in [0, 1]$. If instead $\theta \notin [0, 1]$, fundamentals predict both the timing and the final outcome, as for $\theta < 0$ it is dominant to attack immediately and for $\theta > 1$ it is dominant never to attack.

eventually leading to a phase where an attack is possible but does not take place, finally resulting into an attack. These predictions may help understand how an economy may transit from a phase of tranquillity to a phase in which a crisis is possible, *without* any change in the underlying economic fundamentals. This would not be possible either under common knowledge or with a unique equilibrium.

Finally, note that the existence of phases during which necessarily no agent attacks relies on the property that the lower dominance region is eliminated after the first period. The existence of multiple equilibria, on the other hand, does not rely on this property. As we show next, multiplicity survives even when shocks to fundamentals re-introduce the lower dominance region in every period, in which case a positive measure of agents may attack after any history. What is essential for the multiplicity result is that dynamics generate a non-monotonicity in the payoff of the marginal agent, which is a consequence of the fact that unsuccessful attacks result in a common posterior that assigns sufficiently low measure to low fundamentals. That beliefs take the extreme form of a truncation, is however not necessary.

5 Shocks in fundamentals

In this section, we examine the possibility that the fundamentals change over time as the result of shocks. To this aim, we modify the game as follows. Agents continue to receive private signals about θ , but the regime is abandoned in period t if and only if $A_t \geq \theta + \delta\omega_t$. The variable θ continues to denote the component of the fundamentals for which agents receive information, whereas the shock ω_t can be interpreted either as a shock to the strength of the regime, or as the impact of “noise traders.” It is independent of both θ and the noise in the agents’ signals and has a continuous c.d.f. F .⁹ The scalar $\delta > 0$ parametrizes the volatility in the fundamentals: the baseline model is nested as the limit when $\delta \rightarrow 0$. We assume shocks are not observed. This is important, not only because it is a realistic scenario, but also because unobservable shocks “smooth out” the information generated by the fact that the regime is still in place and may reintroduce the lower-dominance region. We nevertheless show that our multiplicity result continues to hold provided the volatility of the shocks is small enough.

In this modified game, we loose the ability to characterize equilibrium strategies recursively

⁹We will consider both the case where ω_t has unbounded support, as well as the case where ω_t has a small bounded support; for that reason we do not restrict ω to be normal.

by a sequence of threshold $\{\theta_t^*\}_{t=1}^\infty$, for the regime outcome now depends not only on θ , but also on the shock ω_t . Nevertheless, we can still characterize monotone equilibria through a sequence of thresholds $\{\bar{x}_t\}_{t=1}^\infty$ such that an agent attacks in period t if and only if $x_t \leq \bar{x}_t$, where $\bar{x}_t \in \overline{\mathbb{R}}$.

Consider first the regime outcomes induced by a particular $\{\bar{x}_t\}_{t=1}^\infty$. Given θ , the size of the attack is $A_t(\theta) = \Phi(\sqrt{\alpha_t}(\bar{x}_t - \theta))$ and hence the status quo is abandoned if and only if $\omega_t \leq \bar{\omega}_t(\theta)$, where

$$\bar{\omega}_t(\theta; \bar{x}_t) \equiv \frac{1}{\delta} [\Phi(\sqrt{\alpha_t}(\bar{x}_t - \theta)) - \theta] \quad (9)$$

It follows that the probability of regime change in period t conditional on θ and given that agents use the threshold \bar{x}_t , is

$$p_t(\theta; \bar{x}_t) \equiv \Pr(\omega_t \leq \bar{\omega}_t(\theta; \bar{x}_t)) = F(\bar{\omega}_t(\theta; \bar{x}_t)).$$

Next, consider the beliefs induced by a given $\{\bar{x}_t\}_{t=1}^\infty$. For $t \geq 2$, let $\Psi_t(\theta) = \Psi_t(\theta; \bar{x}^{t-1})$ denote the *common* posterior distribution about θ in period t and $\psi_t(\theta) = \psi_t(\theta; \bar{x}^{t-1})$ the associated density function, where $\bar{x}^t = \{\bar{x}_1, \dots, \bar{x}_t\}$. By Bayes' rule,

$$\psi_t(\theta; \bar{x}^{t-1}) = \frac{\Pr(R_t = 0 | \theta, R_{t-1} = 0) \psi_{t-1}(\theta; \bar{x}^{t-2})}{\int_{-\infty}^{+\infty} \Pr(R_t = 0 | \theta', R_{t-1} = 0) \psi_{t-1}(\theta'; \bar{x}^{t-2}) d\theta'} = \frac{\prod_{s=1}^{t-1} [1 - p_s(\theta; \bar{x}_s)] \psi_1(\theta)}{\int_{-\infty}^{+\infty} \prod_{s=1}^{t-1} [1 - p_s(\theta'; \bar{x}_s)] \psi_1(\theta') d\theta'},$$

where ψ_1 is the density of the initial prior. We similarly let $\Psi_t(\theta|x) = \Psi_t(\theta|x; \bar{x}^{t-1})$ denote the posterior distribution of an agent with *private* information summarized by x in period t , and $\psi_t(\theta|x) = \psi_t(\theta|x; \bar{x}^{t-1})$ the corresponding density.

Finally, consider payoffs. For any $t \geq 1$, $x \in \mathbb{R}$, and $\bar{x}^t \in \overline{\mathbb{R}}^t$, let

$$v_t(x; \bar{x}^t) \equiv \int_{-\infty}^{+\infty} p_t(\theta; \bar{x}_t) \psi_t(\theta|x; \bar{x}^{t-1}) d\theta - c$$

denote the net expected payoff from attacking in period t for an agent with sufficient statistics x when all other agents attack in period $\tau \leq t$ if and only if their sufficient statistic in τ is less than \bar{x}_τ . Note that $v_t(x; \bar{x}^t)$ depends on both the contemporaneous threshold \bar{x}_t and the sequence of past thresholds \bar{x}^{t-1} ; the former determines the probability of regime change conditional on θ , whereas the latter determines the posterior beliefs about θ . Next, for any $t \geq 1$ and $\bar{x}^t \in \overline{\mathbb{R}}^t$, let

$$V_t(\bar{x}^t) = V_t(\bar{x}^t; \delta) \equiv \begin{cases} \lim_{x \rightarrow +\infty} v_t(x; \bar{x}^t) & \text{if } \bar{x}_t = +\infty \\ v_t(\bar{x}_t; \bar{x}^t) & \text{if } \bar{x}_t \in \mathbb{R} \\ \lim_{x \rightarrow -\infty} v_t(x; \bar{x}^t) & \text{if } \bar{x}_t = -\infty \end{cases}.$$

V_t is the analogue of the function U in the benchmark model: it represents the net payoff from attacking in period t for the marginal agent with sufficient statistic \bar{x}_t . (For $\bar{x}_t = -\infty(+\infty)$, $V_t(\bar{x}^t)$ is defined as the payoff from attacking for an agent with arbitrarily low (respectively, high) x when no (every) other agent attacks in that period.)

Denoting by $\Gamma(\delta)$ the game corresponding to a given volatility δ , we then have the following

Proposition 2 $\{a_t(\cdot)\}_{t=1}^\infty$ is a monotone equilibrium for $\Gamma(\delta)$ if and only if there is a sequence $\{\bar{x}_t\}_{t=1}^\infty = \{\bar{x}_t(\delta)\}_{t=1}^\infty$ such that:

- (i) for all t , $a_t(x_t) = 1$ if $x_t < \bar{x}_t(\delta)$ and $a_t(x_t) = 0$ if $x_t > \bar{x}_t(\delta)$;
- (ii) for $t = 1$, $V_1(\bar{x}_1(\delta); \delta) = 0$;
- (iii) for any $t \geq 2$, either $\bar{x}_t(\delta) = -\infty$ and $V_t(\bar{x}^t(\delta)) \leq 0$; or $-\infty < \bar{x}_t < +\infty$ and $V_t(\bar{x}^t(\delta); \delta) = 0$.

We next examine the properties of V_t , and consequently of the equilibrium set of $\Gamma(\delta)$, as $\delta \rightarrow 0$. Let $\theta_t(\bar{x}_t)$ be the solution to $\theta = \Phi(\sqrt{\alpha_t}(\bar{x}_t - \theta))$, with $\theta_t(-\infty) = 0$ and $\theta_t(+\infty) = 1$, and $\theta_t(\bar{x}^t) = \max_{\tau \leq t} \{\theta_\tau(\bar{x}_t)\}$. As $\delta \rightarrow 0$, the dependence of the regime outcome on the shock ω_t vanishes and therefore the probability $p_t(\theta, \bar{x}_t)$ that the regime is abandoned in period t converges to either 0 or 1, depending on whether θ is lower or higher than $\theta_t(\bar{x}_t)$, like in the benchmark model. It follows that the common posterior in period t converges to the truncated Normal generated by the knowledge that $\theta > \theta_{t-1}$,

$$\Psi_t(\theta; \bar{x}^{t-1}, \delta) \rightarrow \Psi_t(\theta; \bar{x}^{t-1}, 0) = \begin{cases} 0 & \text{if } \theta \leq \theta_{t-1}(\bar{x}^{t-1}) \\ \frac{\Phi(\sqrt{\beta}(\theta-z)) - \Phi(\sqrt{\beta}(\theta_{t-1}-z))}{1 - \Phi(\sqrt{\beta}(\theta_{t-1}-z))} & \text{if } \theta > \theta_{t-1}(\bar{x}^{t-1}) \end{cases}$$

Similarly for the private posteriors, $\Psi_t(\theta|x; \bar{x}^{t-1}, \delta) \rightarrow \Psi_t(\theta|x; \bar{x}^{t-1}, 0)$ as $\delta \rightarrow 0$. It follows that, for any \bar{x}^{t-1} and any $\bar{x}_t > -\infty$,

$$V_t(\bar{x}^t; \delta) \rightarrow U_t(\bar{x}^t),$$

where $U_t(\bar{x}^t) \equiv U(\theta_t(\bar{x}_t); \theta_{t-1}(\bar{x}^{t-1}), z, \alpha_t, \beta)$. That is, the payoff for the marginal agent in the perturbed game converges pointwise to the payoff of the marginal agent in the benchmark game. Using this result, we then prove that, for δ small enough, “essentially” all equilibria in the benchmark game $\Gamma(0)$ can be approximated by equilibria in the perturbed game $\Gamma(\delta)$. The qualification “essential” is needed because convergence may fail in the knife-edge case that U_t reaches its maximum at the equilibrium threshold x_t^* .

Theorem 3 (Shocks in fundamentals) *For any $\varepsilon > 0$, and any $T \geq 1$, there exists $\delta(\varepsilon, T) > 0$ such that, for all $\delta < \delta(\varepsilon, T)$, the following is true:*

For any equilibrium $\{x_t^\}_{t=1}^\infty$ of $\Gamma(0)$ such that $x_t^* \neq \arg \max_x U(x^{*t-1}, x)$ for all $t \leq T$, there exists an equilibrium $\{\bar{x}_t(\delta)\}_{t=1}^\infty$ of $\Gamma(\delta)$ such that, for all $t \leq T$, either $|x_t^* - \bar{x}_t(\delta)| < \varepsilon$, or $\max\{x_t^*, \bar{x}_t(\delta)\} < -1/\varepsilon$.*

We conclude that our results are robust to the introduction of small shocks in the fundamentals.

6 Public news and social learning

6.1 The effect of public news

To capture the effect of public news, we now modify the game as follows. In addition to their private signals, agents observe in each period $t \geq 1$ a public signal $\tilde{z}_t = \theta + \varepsilon_t$, where ε_t is common noise, independent over time, independent of θ , and normally distributed with zero mean and precision $\eta_t^z > 0$. These signals may represent, for example, the information generated by news analysis in the media, publication of government statistics, announcements by policy makers, etc.. The common posterior about θ conditional on z and the history of public signals $\tilde{z}^t \equiv \{\tilde{z}_\tau\}_{\tau=1}^t$ is normal with mean z_t and precision β_t , where z_t and β_t are defined recursively by

$$z_t = \frac{\beta_{t-1}}{\beta_t} z_{t-1} + \frac{\eta_t^z}{\beta_t} \tilde{z}_t, \quad \beta_t = \beta_{t-1} + \eta_t^z,$$

with $(z_0, \beta_0) = (z, \beta)$. We maintain the assumption that $\alpha_t > \beta_t^2/(2\pi)$ in all t , which ensures equilibrium uniqueness in the static game.

The set of monotone equilibria can be constructed in a similar way as in the benchmark model, except for two minor adjustments. First, the posteriors in period t now depend on the random realization of the statistics z_t rather than the fixed initial prior z . Second, although (x_t, z_t) are sufficient statistics for the exogenous public and private information with respect to θ , they are not sufficient statistics with respect to regime outcomes when agents use past public signals as correlation devices. We thus allow agents to condition their actions on the entire sequence \tilde{z}^t , or equivalently on $z^t \equiv \{z_\tau\}_{\tau=1}^t$, which means that the thresholds x_t^* and θ_t^* are now functions of the exogenous history of public signals. Following then the same steps as in Section 3, we conclude

Proposition 3 $\{a_t(\cdot)\}_{t=0}^\infty$ is an equilibrium if and only if there is a sequence $\{x_t^*(z^t), \theta_t^*(z^t)\}_{t=1}^\infty$, with $x_t^*(z^t) \in \overline{\mathbb{R}}$ and $\theta_t^*(z^t) \in (0, 1)$, such that for all t :

- (i) $a_t(\cdot) = 1$ if $x_t < x_t^*(z^t)$ and $a_t(\cdot) = 0$ if $x_t > x_t^*(z^t)$;
- (ii) at $t = 1$, $\theta_1^*(z_1)$ solves $U(\theta_1^*; -\infty, \alpha_1, \beta_1, z_1) = 0$ and $x_1^*(z_1) = X(\theta_1^*(z_1), \alpha_1)$
- (iii) at $t \geq 2$, either $\theta_t^*(z^t)$ solves

$$U(\theta_t^*; \theta_{t-1}^*, \alpha_t, \beta_t, z_t) = 0 \tag{10}$$

and $x_t^*(z^t) = X(\theta_t^*(z^t), \alpha_t)$, or $\theta_t^*(z^t) = \theta_{t-1}^*(z^{t-1})$ and $x_t^*(z^t) = -\infty$.

Consider now the impact of public news on the possibility of an attack. Since a higher z_t reduces the posterior probability of regime change for any given (monotone) strategy followed by the other agents, $U(\theta; \theta_{t-1}, \alpha_t, \beta_t, z_t)$ decreases with z_t . It follows that, in every period t , there exists a threshold \bar{z}_t such that (10) admits a solution if and only if $z_t \leq \bar{z}_t$; that is, an attack is possible if and only if the exogenous public information conveyed by z_t is sufficiently negative. Moreover, the threshold \bar{z}_t decreases with θ_{t-1}^* since, other things equal, a higher θ_{t-1}^* implies a lower posterior probability of regime change. Finally, an unsuccessful attack, other things equal, causes a discrete reduction in U for all z and therefore a discrete increase in the probability that the game enters a phase of inaction during which no attack is possible.

Like in the benchmark model without public news, there always exist equilibria in which attacks cease after any arbitrary period. For example, there is an equilibrium with no action after period 1. However, (10) necessarily admits a solution for z_t sufficiently low, no matter the historical threshold θ_{t-1}^* and the precisions α_t and β_t . Therefore, unlike the benchmark model, there always exist equilibria in which N attacks occur with *positive* probability after any period $t \geq 1$, for any $N \geq 1$. We conclude:

Theorem 4 (Public news) *In the game with public news, there always exist multiple equilibria, no matter the mean of the prior and the precision of private and public information.*

The above result extends and reinforces the multiplicity result of Theorem 2: multiplicity now arises for any z , since the arrival of new public information “smooths out” the strong impact of the initial prior and makes further attacks possible with positive probability.

6.2 Social learning

Apart from the knowledge that the regime survived past attacks, the dynamics of information were assumed to be exogenous. However, one can endogenize part of this information by letting agents receive noisy signals – either private or public – about past aggregate activity. This also illustrates the robustness of our results to a form of *social learning*.

We modify the game as follows. In addition to the exogenous private and public signals about θ , agents observe in each period $t \geq 2$ a private and/or a public signal about the size of the attack in the previous period. These signals are, respectively,

$$\tilde{X}_{it} = S(A_{t-1}, u_{it}) \quad \text{and} \quad \tilde{Z}_t = S(A_{t-1}, v_t),$$

where u_{it} is idiosyncratic noise, v_t is common noise, and $S : (0, 1) \times \mathbb{R} \rightarrow \mathbb{R}$. To preserve Normality of the information structure, we adopt a similar specification as Dasgupta (2002): $S(A, u) = \Phi^{-1}(A) + u$, $u_t^i \sim \mathcal{N}(0, 1/\gamma_t^x)$, and $v_t \sim \mathcal{N}(0, 1/\gamma_t^z)$, where λ_t^x and λ_t^z parametrize the level of exogenous noise in social learning.

Using the property that, in any monotone equilibrium, $A_{t-1} = \Phi(\sqrt{\alpha_{t-1}}(x_{t-1}^* - \theta))$, and following similar steps as in Section 3, we can again construct sufficient statistics x_t and z_t with precisions α_t and β_t such that

$$\theta | \left(\tilde{x}^t, \tilde{z}^t, \tilde{X}^t, \tilde{Z}^t \right) \sim \mathcal{N} \left(\frac{\alpha_t}{\alpha_t + \beta_t} x_t + \frac{\beta_t}{\alpha_t + \beta_t} z_t, \frac{1}{\alpha_t + \beta_t} \right).$$

The only difference is that these statistics are now defined by

$$\begin{aligned} x_t &= \frac{\alpha_{t-1}}{\alpha_t} x_{t-1} + \frac{\eta_t^x}{\alpha_t} \tilde{x}_t + \frac{\mathbf{1}_{t-1} \alpha_{t-1} \gamma_t^x}{\alpha_t} \left\{ x_{t-1}^* - \tilde{X}_t / \sqrt{\alpha_{t-1}} \right\}, \\ z_t &= \frac{\beta_{t-1}}{\beta_t} z_{t-1} + \frac{\eta_t^z}{\beta_t} \tilde{z}_t + \frac{\mathbf{1}_{t-1} \alpha_{t-1} \gamma_t^z}{\beta_t} \left\{ x_{t-1}^* - \tilde{Z}_t / \sqrt{\alpha_{t-1}} \right\}, \\ \alpha_t &= \alpha_{t-1} + \eta_t^x + \mathbf{1}_{t-1} \alpha_{t-1} \gamma_t^x \quad \text{and} \quad \beta_t = \beta_{t-1} + \eta_t^z + \mathbf{1}_{t-1} \alpha_{t-1} \gamma_t^z, \end{aligned}$$

where $\mathbf{1}_t$ is an indicator variable that takes value 1 if an attack occurs in period t ($A_t > 0$) and 0 otherwise, and the terms $x_{t-1}^* - \tilde{X}_t / \sqrt{\alpha_{t-1}}$ and $x_{t-1}^* - \tilde{Z}_t / \sqrt{\alpha_{t-1}}$ are just linear transformations of X_t and Z_t .

Although the realizations of (x_t, z_t) and similarly the dynamics of (α_t, β_t) are now partly endogenous, the equilibrium characterization is otherwise the same as in the model with only exogenous signals. By implication, our multiplicity results directly extend to the game with social

learning. Moreover, the structure of equilibrium dynamics remains as before, except for the property that an unsuccessful attack now does not necessarily reduce the incentives for further attacks. This is because an unsuccessful attack also generates new private and public signals, which in some cases may offset the impact of the information that the status quo has survived the attack.

To see this, consider first the case that all signals are private ($\gamma_t^x > 0, \eta_t^x > 0, \gamma_t^z = \eta_t^z = 0$). Then the only effect of social learning is to contribute to an endogenous increase in α_t following any attack. A new attack may now become possible right after an unsuccessful one even without any exogenous arrival of information. If the signal about past attacks is public rather than private ($\gamma_t^z > 0 = \gamma_t^x$), then the dynamics of α_t remain exogenous whereas z_t and β_t change only due to an unsuccessful attack. In this case, an unsuccessful attack may make a further attack possible if the endogenous public signal about the size of the attack is sufficiently negative, similarly to the case of exogenous public news.

7 Conclusions

This paper examined how the dynamics of information influences the dynamics of coordination in a stylized global game of regime change. It has been shown that dynamics introduce multiple equilibria under the same informational assumptions that guarantee uniqueness in the static benchmark examined in the literature. As a result, the occurrence of coordinated attacks and the timing of regime change depend, not only on the underlying fundamentals and the evolution of information, but also on when agents switch from lenient to aggressive behavior. Despite this multiplicity, the model generates a few distinctive predictions for equilibrium dynamics, such as the succession of phases of tranquility, during which agents accumulate information and no attacks are possible, and periods of crises, during which attacks may occur but do not necessarily take place. Although the model is highly stylized, we believe these predictions may help understand phenomena such as currency attacks, financial crashes, or political change.

References

- [1] Angeletos, George-Marios, Christian Hellwig, and Alessandro Pavan (2003) “Coordination and Policy Traps,” MIT/UCLA/Northwestern working paper.
- [2] Atkeson, Andrew (2000), “Discussion on Morris and Shin,” *NBER Macroeconomics Annual*.
- [3] Carlsson, Hans, and Eric van Damme (1993), “Global Games and Equilibrium Selection,” *Econometrica* 61, 5, 989-1018.
- [4] Chamley, Christophe (1999), “Coordinating Regime Switches,” *Quarterly Journal of Economics* 114, 3, 869-905.
- [5] Corsetti, Giancarlo, Bernardo Guimaraes and Nouriel Roubini, (2004), “International Lending of Last Resort and Moral Hazard: A Model of IMF’s Catalytic Finance,” University of Rome, working paper.
- [6] Dasgupta, Amil (2002), “Coordination, Learning and Delay,” working paper, London School of Economics
- [7] Edmond, Chris (2003), “Information and the Limits to Autocracy,” working paper, UCLA.
- [8] Giannitsarou, Chryssi, and Flavio Toxvaerd (2003), “Recursive Global Games,” working paper, Universidade Nova de Lisboa and Hebrew University of Jerusalem.
- [9] Goldstein, Itay (2002), “Strategic Complementarities and the Twin Crises,” working paper, Duke University.
- [10] Goldstein, Itay, and Ady Pauzner (2000), “Demand Deposit Contracts and the Probability of bank Runs,” working paper, Duke University and Tel Aviv University.
- [11] Goldstein, Itay, and Ady Pauzner (2001), “Contagion of Self-Fulfilling Financial Crises due to Diversification of Investment Portfolios,” working paper, Duke University and Tel Aviv University.
- [12] Heidhues, Paul, and Nicolas Melissas (2003), “Equilibria in a Dynamic Global Game: The Role of Cohort Effects”, working paper, WZB Berlin and University of Leicester.

- [13] Morris, Stephen and Hyun Song Shin (1998), “Unique Equilibrium in a Model of Self-Fulfilling Currency Attacks,” *American Economic Review*, 88, 3, 587-597.
- [14] Morris, Stephen and Hyun Song Shin (1999), “A Theory of the Onset of Currency Attacks,” in Agenor, Vines and Weber, eds., *Asian Financial Crisis: Causes, Contagion and Consequences*, Cambridge, UK: Cambridge University Press.
- [15] Morris, Stephen and Hyun Song Shin (2000), “Rethinking Multiple Equilibria in Macroeconomics,” *NBER Macro Annual*.
- [16] Morris, Stephen and Hyun Song Shin (2001), “Global Games - Theory and Applications,” in M. Dewatripont, L. Hansen, and S. Turnovsky, eds., *Advances in Economics and Econometrics (8th World Congress of the Econometric Society)*, Cambridge, UK: Cambridge University Press.
- [17] Morris, Stephen and Hyun Song Shin (2004), “Coordination Risk and the Price of Debt,” *European Economic Review*.

8 Appendix

Proof of Proposition (1). Solving (1) for \hat{x} gives $\hat{x} = \hat{\theta} + \alpha^{-1/2}\Phi^{-1}(\hat{\theta})$. Substituting this into (2) gives a single equation in $\hat{\theta}$:

$$U^{st}(\hat{\theta}; \alpha, \beta, z) = 0, \tag{11}$$

where

$$U^{st}(\theta; \alpha, \beta, z) \equiv 1 - \Phi\left(\frac{\sqrt{\alpha}}{\sqrt{\alpha+\beta}}\left[\Phi^{-1}(\theta) + \frac{\beta}{\sqrt{\alpha}}(z - \theta)\right]\right) - c. \tag{12}$$

Note that $U^{st}(\theta; \cdot)$ is continuous in $\theta \in (0, 1)$, with $\lim_{\theta \rightarrow 0} U^{st}(\theta) = 1 - c > 0$ and $\lim_{\theta \rightarrow 1} U^{st}(\theta) = -c < 0$. A solution to (11) therefore always exists. Next, note that

$$\frac{dU^{st}(\theta; \cdot)}{d\theta} = -\frac{\sqrt{\alpha}}{\sqrt{\alpha+\beta}}\phi\left(\frac{\sqrt{\alpha}}{\sqrt{\alpha+\beta}}\left[\Phi^{-1}(\theta) + \frac{\beta}{\sqrt{\alpha}}(z - \theta)\right]\right)\left[\frac{1}{\phi(\Phi^{-1}(\theta))} - \frac{\beta}{\sqrt{\alpha}}\right].$$

Since $\min_{\theta \in [0,1]} [1/\phi(\Phi^{-1}(\theta))] = \sqrt{2\pi}$, the condition $\alpha \geq \beta^2/(2\pi)$ is both necessary and sufficient for U^{st} to be monotonic in θ , in which case the monotone equilibrium is unique. Finally, for the proof that only this equilibrium survives iterated deletion of dominated strategies, see Morris and Shin (2000). ■

Proof of Lemma (1). We prove the claim by induction. At $t = 1$, the result clearly holds since the equilibrium play is the same as in the static benchmark. Consider next any $t \geq 2$ and suppose that the result holds for any $\tau \leq t - 1$. Since $a_t(x^t)$ is non-increasing in x^t , $A_t(\theta)$ is non-increasing in θ , implying that either $A_t(\theta) \leq \theta_{t-1}^*$ for all θ , in which case $\theta_t^* = \theta_{t-1}^*$, or there exists $\theta_t^* > \theta_{t-1}^*$ such that $A_t < \theta$ if and only if $\theta > \theta_t^*$. It follows that agent i 's posterior that a regime change will take place in period t is given by $\Pr(R_{t+1} = 1 | x_t, R_t = 0) = \Pr(\theta \leq \theta_t^* | x_t, \theta > \theta_{t-1}^*)$. Since the distribution of x_t given θ satisfies the MLRP, it follows that $\Pr(\theta \leq \theta_t^* | x_t, \theta > \theta_{t-1}^*)$ is strictly decreasing in x_t . This implies that either $\Pr(\theta \leq \theta_t^* | x_t, \theta > \theta_{t-1}^*) < c$ for all x_t , in which case $x_t^* = -\infty$; or $\Pr(\theta \leq \theta_t^* | x_t, \theta > \theta_{t-1}^*) > c$ for all x_t , in which case $x_t^* = +\infty$; or there exists $x_t^* \in \mathbb{R}$ such that $\Pr(\theta \leq \theta_t^* | x_t, \theta > \theta_{t-1}^*) \geq c$ if and only if $x_t \leq x_t^*$. Finally, since the status quo is never abandoned for $\theta \geq 1$,

$$\Pr(R_{t+1} = 1 | x_t, R_t = 0) \leq \Pr(\theta \leq 1 | x_t, \theta > \theta_{t-1}^*) \leq \Pr(\theta \leq 1 | x_t) = \Phi\left(\sqrt{\alpha_t + \beta} \left[1 - \frac{\alpha_t x_t + \beta z}{\alpha_t + \beta}\right]\right)$$

and therefore it is dominant not to attack for x_t high enough, which rules out $x_t^* = +\infty$. But then $A_t(\theta) < 1$ for all θ and therefore $\theta_t^* < 1$, which together with $\theta_t^* \geq \theta_1^* > 0$ completes the result. ■

Proof of Theorem (1). Necessity. For $t = 1$, the result follows from the fact that $U(\theta; -\infty, \alpha, \beta, z) = U^{st}(\theta; \alpha, \beta, z)$ for all $(\theta, \alpha, \beta, z)$ and hence, in any equilibrium of the dynamic game, $a_1(\cdot)$ is necessarily the same as in the unique equilibrium of the static game where agents move only in period 1. For $t \geq 2$, by Lemma 1, $R_t = 0$ implies that $\theta > \theta_{t-1}^* > 0$ with probability one. Hence, there always exist monotone equilibria in which no agent attacks in period t , in which case $x_t^* = -\infty$ and $\theta_t^* = \theta_{t-1}^*$. Finally, suppose $\theta_t^* > \theta_{t-1}^*$. Then,

$$\Pr(R_{t+1} = 1 | x_t, R_t = 0) = \Pr(\theta \leq \theta_t^* | x_t, \theta > \theta_{t-1}^*) = 1 - \frac{\Phi\left(\sqrt{\alpha_t + \beta} \left[\frac{\alpha_t x_t + \beta z}{\alpha_t + \beta} - \theta_t^*\right]\right)}{\Phi\left(\sqrt{\alpha_t + \beta} \left[\frac{\alpha_t x_t + \beta z}{\alpha_t + \beta} - \theta_{t-1}^*\right]\right)},$$

which is continuous and decreasing in x_t , with $\lim_{x_t \rightarrow \infty} \Pr(\theta \leq \theta_t^* | x_t, \theta > \theta_{t-1}^*) = 0$. Therefore, for an agent to find it optimal to attack if and only if $x_t \leq x_t^*$, it must be that x_t^* solves $\Pr(\theta \leq \theta_t^* | x_t^*, \theta > \theta_{t-1}^*) - c = 0$, or equivalently

$$1 - \frac{\Phi\left(\sqrt{\alpha_t + \beta} \left[\frac{\alpha_t}{\alpha_t + \beta} x_t^* + \frac{\beta}{\alpha_t + \beta} z - \theta_t^*\right]\right)}{\Phi\left(\sqrt{\alpha_t + \beta} \left[\frac{\alpha_t}{\alpha_t + \beta} x_t^* + \frac{\beta}{\alpha_t + \beta} z - \theta_{t-1}^*\right]\right)} - c = 0.$$

Substituting $x_t^* = X(\theta_t^*; \alpha_t) \equiv \theta_t^* + \alpha_t^{-1/2} \Phi^{-1}(\theta_t^*)$ and rearranging gives (8).

Sufficiency. Take any sequence $\{x_t^*, \theta_t^*\}_{t=1}^\infty$ that satisfies conditions (i) – (iii), and let $a_t(\cdot) = 1$ if $x_t \leq x_t^*$ and $a_t(\cdot) = 0$ if $x_t > x_t^*$, for all t . By construction, when all other agents follow $\{a_t(\cdot)\}_{t=1}^\infty$, and for every t , the net payoff from attacking is positive for $x_t < x_t^*$ and negative for $x_t > x_t^*$, which implies that the sequence of strategies $\{a_t(\cdot)\}_{t=1}^\infty$ satisfies the rationality condition (3) and therefore constitutes an equilibrium. ■

Proof of Lemma (2). Consider part (i). Continuity is obvious. Non-monotonicity in θ^* for $\theta_{-1} \in (0, 1)$ follows from the following facts: when $\theta^* \leq \theta_{-1}$, $\Pr(\theta \leq \theta^* | x, \theta > \theta_{-1}) = 0$ and therefore $U(\theta^*; \cdot) = -c$; when instead $\theta^* \in (\theta_{-1}, 1)$, $\Pr(\theta \leq \theta^* | x, \theta > \theta_{-1}) > 0$ and therefore $U(\theta^*; \cdot) > -c$; and finally $U(1; \cdot) = \lim_{\theta^* \rightarrow 1} U(\theta^*; \cdot) = -c$, since $\Pr(\theta \leq \theta^* | x, \theta > \theta_{-1}) \leq \Pr(\theta \leq 1 | x)$, $\Pr(\theta \leq 1 | x) \rightarrow 0$ as $x \rightarrow \infty$, and $X(\theta^*, \alpha) \rightarrow \infty$ as $\theta^* \rightarrow 1$.

Next, the monotonicity of posterior beliefs about θ in x and z (which follows from standard representation theorems) implies that U is decreasing in z . Monotonicity in θ_{-1} follows from the fact that, for any x ,

$$\Pr(\theta \leq \theta^* | x, \theta > \theta_{-1}) = \frac{\Pr(\theta_{-1} < \theta \leq \theta^* | x)}{\Pr(\theta > \theta_{-1} | x)} = 1 - \frac{\Pr(\theta > \theta^* | x)}{\Pr(\theta > \theta_{-1} | x)}$$

and $\Pr(\theta > \theta_{-1} | x)$ is decreasing in θ_{-1} .

As $\alpha \rightarrow \infty$, for any $x > \theta_{-1}$ and any θ^* , $\Pr(\theta > \theta_{-1} | x) \rightarrow 1$ and therefore $\Pr(\theta \leq \theta^* | x, \theta > \theta_{-1}) - \Pr(\theta \leq \theta^* | x) \rightarrow 0$. Moreover, for any $\theta^* \in (\theta_{-1}, 1)$, $X(\theta^*; \alpha) \rightarrow \theta^* > \theta_{-1}$. It follows that, for any $\theta^* \in (\theta_{-1}, 1)$, $U(\theta^*; \theta_{-1}, \cdot) - U^{st}(\theta^*; \cdot) \rightarrow 0$ as $\alpha \rightarrow \infty$. Finally, $U^{st}(\theta^*; \cdot) \rightarrow 1 - c - \theta^*$ follows directly from (12)

For (iii), take any $\theta_{t-1}^* > \theta_\infty$. Since $\lim_{\alpha \rightarrow \infty} U(\theta^*, \theta_{-1}, \alpha, \beta, z) = \theta_\infty - \theta^*$ for all $\theta^* > \theta_{-1}$, there exists $\bar{\alpha}$ such that, for all $\alpha > \bar{\alpha}$, $U(\theta^*, \theta_{t-1}^*, \alpha, \beta, z) < 0$ for all $\theta^* > \theta_{t-1}^* > \theta_\infty$, and therefore (8) admits no solution for $\alpha_t > \bar{\alpha}$.

For (iv), take any $\theta_{t-1}^* < \theta_\infty$. Since $\lim_{\alpha \rightarrow \infty} U(\theta^*, \theta_{-1}, \alpha, \beta, z) = \theta_\infty - \theta^*$ for all $\theta^* > \theta_{-1}$, there exist $\theta' \in (\theta_{t-1}^*, \theta_\infty)$ and $\bar{\alpha}$ such that, for any $\alpha > \bar{\alpha}$, $U(\theta'; \theta_{t-1}^*, \alpha, \beta, z) > 0$. By the continuity of $U(\theta; \theta_{-1}, z, \alpha, \beta)$ in θ and the fact that $\lim_{\theta \rightarrow 1} U(\theta; \theta_{t-1}^*, \alpha, \beta, z) = -c$, it follows then that (8) admits a solution for $\alpha_t > \bar{\alpha}$.

Finally, consider (v). Fix $\theta_{t-1}^*, z, \beta, \alpha_t$ and suppose that θ_t^* is the highest solution to (8), which means that $U(\theta^*, \theta_{t-1}^*, \alpha_t, \beta, z) < 0$ for all $\theta^* > \theta_{t-1}^*$. This, together with the properties that U is decreasing in θ_{-1} , is continuous in θ^* , and satisfies $U(\theta_{-1}, \theta_{-1}, z, \alpha, \beta) = -c$, implies that there exists $\Delta > 0$ such that, for all $\theta^* > \theta_t^*$, $U(\theta^*, \theta_t^*, z, \alpha_t, \beta) < -\Delta$. By continuity of U in α , then,

there exists $\underline{\alpha} > \alpha_t$ such that $U(\theta^*, \theta_t^*, z, \alpha, \beta) < 0$ for all $\alpha \in [\alpha_t, \underline{\alpha}]$ and all $\theta \in [\theta_t^*, 1]$ and therefore (8) admits no solution in any period $\tau > t$ for which $\alpha_\tau < \underline{\alpha}$. ■

Lemma A1 *There exist thresholds $\underline{z} \leq \bar{z} \leq \bar{\bar{z}}$ such that: $\hat{\theta}_t \leq \hat{\theta}_1$ for all t if $z \leq \underline{z}$; $\hat{\theta}_1 \leq (\geq) \theta_\infty$ if and only if $z \geq (\leq) \bar{z}$; and $\hat{\theta}_t < \theta_\infty$ for all t if and only if $z > \bar{\bar{z}}$. These thresholds satisfy $\underline{z} = \bar{z} = \bar{\bar{z}}$ when $c \leq 1/2$ and $\underline{z} \leq \bar{z} < \bar{\bar{z}}$ when $c > 1/2$.*

Proof. Let $\hat{\theta} = \hat{\theta}(\alpha)$ the static-game threshold by $U(\hat{\theta}; -\infty, z, \alpha, \beta) = 0$ and let

$$\begin{aligned}\tilde{z}(\alpha) &\equiv \theta_\infty + \frac{\sqrt{\alpha + \beta} - \sqrt{\alpha}}{\beta} \Phi^{-1}(\theta_\infty) \\ \hat{z}(\alpha) &\equiv \Phi\left(\frac{\sqrt{\alpha}}{\sqrt{\alpha + \beta}} \Phi^{-1}(\theta_\infty)\right) + \frac{1}{\sqrt{\alpha + \beta}} \Phi^{-1}(\theta_\infty)\end{aligned}$$

The threshold $\tilde{z}(\alpha)$ is defined so that $\hat{\theta}(\alpha) \geq (\leq) \theta_\infty$, if and only if $z \leq (\geq) \tilde{z}(\alpha)$. The threshold $\hat{z}(\alpha)$, on the other hand, is defined so that $z > \hat{z}(\alpha)$ implies that $\partial\hat{\theta}/\partial\alpha > 0$, whereas $z < \hat{z}(\alpha)$ implies that $\partial\hat{\theta}/\partial\alpha < 0$.

First, consider $c = 1/2$, in which case $\hat{z}(\alpha) = \tilde{z}(\alpha) = 1/2$ for all α . When $z < 1/2$, $\hat{\theta}(\alpha) > \theta_\infty$ for all α and $\partial\hat{\theta}/\partial\alpha < 0$, and therefore $\hat{\theta}_1 = \max_{t \geq 1} \hat{\theta}_t > \theta_\infty$. When instead $z = 1/2$, $\hat{\theta}(\alpha) = \theta_\infty$ and $\partial\hat{\theta}/\partial\alpha = 0$, and therefore $\hat{\theta}_1 = \max_{t \geq 1} \hat{\theta}_t = \theta_\infty$. Finally, when $z > 1/2$, $\hat{\theta}(\alpha) < \theta_\infty$ and $\partial\hat{\theta}/\partial\alpha > 0$, and therefore $\hat{\theta}_1 \leq \hat{\theta}_t < \theta_\infty$ for all t . The result thus holds with $\underline{z} = \bar{z} = \bar{\bar{z}} = 1/2$.

Next, consider $c < 1/2$, in which case $\tilde{z}(\alpha)$ and $\hat{z}(\alpha)$ are both decreasing in α , satisfy $\hat{z}(\alpha) > \tilde{z}(\alpha) > \theta_\infty$ for all α , and converge to θ_∞ as $\alpha \rightarrow \infty$. When $z \leq \theta_\infty$, then clearly $z < \tilde{z}(\alpha) < \hat{z}(\alpha)$ for all α and therefore $\hat{\theta}(\alpha)$ is always higher than θ_∞ and decreasing in α , which implies that $\hat{\theta}_1 = \max_{t \geq 1} \hat{\theta}_t > \theta_\infty$. When $z \in (\theta_\infty, \tilde{z}(\alpha_1))$, there are $\alpha'' > \alpha' > \alpha_1$ such that $\tilde{z}(\alpha') = \hat{z}(\alpha'') = z$. For $\alpha \in [\alpha_1, \alpha')$, $\hat{\theta}(\alpha)$ is higher than θ_∞ and decreases with α . As soon as $\alpha \in (\alpha', \alpha'')$, $\hat{\theta}(\alpha)$ becomes lower than θ_∞ and continues to decrease with α . Once $\alpha \geq \alpha''$, $\hat{\theta}(\alpha)$ starts increasing with α , but never exceeds θ_∞ . Hence, once again, $\hat{\theta}_1 = \max_{t \geq 1} \hat{\theta}_t > \theta_\infty$. When $z = \tilde{z}(\alpha_1)$, $\hat{\theta}_1 = \max_{t \geq 1} \hat{\theta}_t = \theta_\infty$. Finally, when $z > \tilde{z}(\alpha_1)$, $\hat{\theta}(\alpha) < \theta_\infty$ for all α , and therefore $\hat{\theta}_t < \theta_\infty$ for all t . We conclude that the result holds for $c < 1/2$ with $\underline{z} = \bar{z} = \bar{\bar{z}} = \tilde{z}(\alpha_1)$.

Finally, consider $c > 1/2$, in which case $\hat{z}(\alpha)$ and $\tilde{z}(\alpha)$ are both increasing in α , satisfy $\hat{z}(\alpha) < \tilde{z}(\alpha) < \theta_\infty$, and converge to θ_∞ as $\alpha \rightarrow \infty$. When $z \leq \hat{z}(\alpha_1)$, then clearly $z < \hat{z}(\alpha) < \tilde{z}(\alpha)$ for all $\alpha > \alpha_1$ and therefore $\hat{\theta}(\alpha)$ is always higher than θ_∞ and decreasing in α , which implies that $\hat{\theta}_1 = \max_{t \geq 1} \hat{\theta}_t > \theta_\infty$. When $z \in (\hat{z}(\alpha_1), \tilde{z}(\alpha_1))$, there is $\alpha' > \alpha_1$ such $\hat{z}(\alpha') = z$. For $\alpha \in (\alpha_1, \alpha')$, $\hat{\theta}(\alpha)$ is higher than θ_∞ and increasing in α , whereas for $\alpha > \alpha'$, $\hat{\theta}(\alpha)$ decreases with α , converging

to θ_∞ from above. It follows that $\max_{t \geq 1} \hat{\theta}_t \geq \hat{\theta}_1 > \theta_\infty$. When $z = \tilde{z}(\alpha_1)$, $\max_{t \geq 1} \hat{\theta}_t \geq \hat{\theta}_1 = \theta_\infty$. When $z \in (\tilde{z}(\alpha_1), \theta_\infty)$, there are $\alpha'' > \alpha' > \alpha_1$ such that $\tilde{z}(\alpha') = \hat{z}(\alpha'') = z$. For $\alpha \in (\alpha_1, \alpha')$, $\hat{\theta}(\alpha)$ is lower than θ_∞ and increasing in α . For $\alpha \in (\alpha', \alpha'')$, $\hat{\theta}(\alpha)$ is higher than θ_∞ and increases with α . And for $\alpha > \alpha''$, $\hat{\theta}(\alpha)$ decreases with α and asymptotes to θ_∞ from above. Hence, $\max_{t \geq 1} \hat{\theta}_t > \theta_\infty > \hat{\theta}_1$. Finally, when $z \geq \theta_\infty$, then clearly $z > \tilde{z}(\alpha) > \hat{z}(\alpha)$ for all α and therefore $\hat{\theta}(\alpha)$ is always lower than θ_∞ , increases with α , and asymptotes to θ_∞ from below. Hence, $\hat{\theta}_t < \theta_\infty$ for all t . We conclude that the result holds for $c > 1/2$ with $\underline{z} = \hat{z}(\alpha_1)$, $\bar{z} = \tilde{z}(\alpha_1)$, and $\bar{\bar{z}} = \theta_\infty$. ■

Proof of Proposition (2). We first rule out the possibility that in any equilibrium, $\bar{x}_t = +\infty$ for some t . Indeed, for any $\hat{\theta}$

$$\begin{aligned} V_t(x, \bar{x}^{t-1}) &\leq \int_{-\infty}^{+\infty} p_t(\theta; +\infty) \psi_t(\theta|x; \bar{x}^{t-1}) d\theta - c \\ &= \int_{-\infty}^{+\infty} F\left(\frac{1}{\delta}(1-\theta)\right) \psi_t(\theta|x; \bar{x}^{t-1}) d\theta - c \\ &\leq \Psi_t(\hat{\theta}|x; \bar{x}^{t-1}) + F\left(\frac{1}{\delta}(1-\hat{\theta})\right) \left[1 - \Psi_t(\hat{\theta}|x; \bar{x}^{t-1})\right] - c. \end{aligned}$$

Furthermore,

$$\lim_{x \rightarrow +\infty} \Psi_t(\hat{\theta}|x; \bar{x}^{t-1}) \leq \lim_{x \rightarrow +\infty} \Phi\left(\sqrt{\alpha_t + \beta} \left(\hat{\theta} - \frac{\alpha_t x + \beta z}{\alpha_t + \beta}\right)\right) = 0,$$

so that $\lim_{x \rightarrow \infty} V_t(x, \bar{x}^{t-1}) \leq F\left(\frac{1}{\delta}(1-\hat{\theta})\right) - c$. Choosing $\hat{\theta}$ s.t. $F\left(\frac{1}{\delta}(1-\hat{\theta})\right) < c$ then gives the result.

Sufficiency. Now, suppose the sequence $\{\bar{x}_t\}_{t=1}^\infty$ satisfies conditions (ii) and (iii) in the proposition. The monotonicity of $v_t(x; \bar{x}^t)$ w.r.t. x guarantees that $v_t(x; \bar{x}^t) \geq V_t(\bar{x}^t)$ if and only if $x \leq \bar{x}_t$, for all t and \bar{x}_t , which implies that $a_t(\cdot) = 1$ for $x_t < \bar{x}_t(\delta)$ and $a_t(\cdot) = 0$ for $x_t > \bar{x}_t(\delta)$ is sequentially optimal and hence $\{a_t(\cdot)\}_{t=1}^\infty$ constitutes a monotone equilibrium.

Necessity. Conversely, suppose that $\{a_t(\cdot)\}_{t=1}^\infty$ is a monotone equilibrium. Since in any such equilibrium $\Pr(R_{t+1} = 1 | \theta, R_t = 0)$ is a decreasing function in θ , the expected payoff from attacking is also necessarily decreasing in the sufficient statistics x_t , implying that agents will follow cut-off strategies at every period. Sequential rationality then requires that at any cut-off, $V_t(\bar{x}^t) = 0$ if

$-\infty < \bar{x}_t < +\infty$, $V_t(\bar{x}^t) \leq 0$ if $\bar{x}_t = -\infty$, and $V_t(\bar{x}^t) \geq 0$ if $\bar{x}_t = +\infty$. Next, note that at $t = 1$,

$$\begin{aligned} V_1(-\infty) &= \lim_{x_1 \rightarrow -\infty} \int_{-\infty}^{+\infty} F\left(\frac{1}{\delta}(-\theta)\right) \psi_1(\theta|x) d\theta - c \\ &\geq \lim_{x_1 \rightarrow -\infty} \Psi_1(\hat{\theta}|x) F\left(\frac{1}{\delta}(-\hat{\theta})\right) - c \\ &= F\left(\frac{1}{\delta}(-\hat{\theta})\right) - c \end{aligned}$$

for some arbitrary $\hat{\theta}$. Then choosing $\hat{\theta}$ such that $F\left(\frac{1}{\delta}(-\hat{\theta})\right) > c$, proves that necessarily $\bar{x}_1 > -\infty$. Along with the result that $\bar{x}_t = +\infty$ can never be part of an equilibrium, for any t , this proves that (i) – (iii) must necessarily hold in any monotone equilibrium. ■

Proof of Theorem (3). With some abuse of notation, let

$$U_1(\bar{x}_1) \equiv U(\theta_1(\bar{x}_1); -\infty, z, \alpha_1, \beta)$$

and

$$U_t(\bar{x}^t) \equiv U\left(\theta_t(\bar{x}_t); \max_{\tau \leq t-1} \{\theta_\tau(\bar{x}_\tau)\}, z, \alpha_t, \beta\right)$$

denote the payoff for the marginal agent in the unperturbed game, where $\theta_t(\bar{x}_t)$ is implicitly defined by $\Pr(x_t \leq \bar{x}_t | \theta; \alpha_t) = \theta$, with $\theta_t(+\infty) = 1$ and $\theta_t(-\infty) = 0$.

We prove the proposition in 5 steps. Step 1 establishes some continuity properties of U_t and V_t and Step 2 pointwise convergence of V_t to U_t as $\delta \rightarrow 0$. Step 3 proves existence and shows how to construct equilibria for the perturbed game. Finally, Steps 4 proves how equilibria in $\Gamma(0)$ can be approximated by equilibria in $\Gamma(\delta)$

Step 1: Continuity of U_t and V_t .¹⁰

Consider first the payoff of the marginal agent in the unperturbed game, $U_t(\bar{x}^t)$. Note that the function $\theta_t(\bar{x}_t)$ is continuous in $\bar{x}_t \in \overline{\mathbb{R}}$. It follows that the function $\theta_{t-1}(\bar{x}^{t-1}) \equiv \max_{\tau \leq t-1} \{\theta_\tau(\bar{x}_\tau)\}$ is also continuous in $\bar{x}^{t-1} \in \overline{\mathbb{R}}^t$. But then, since $U(\theta; \theta_{-1}, \alpha_t, \beta, z)$ is continuous in both θ and θ_{-1} , this implies that the function $U_t(\bar{x}^t)$ is continuous in $\overline{\mathbb{R}}^t$.

¹⁰For any function f defined over $\overline{\mathbb{R}}$, we will say that f is continuous at $-\infty$ if $\lim_{x \rightarrow -\infty} f(x) = f(-\infty)$, and similarly for $+\infty$. Furthermore, we will say that the function f defined over $\overline{\mathbb{R}}^t$ is continuous at $x^t \in \overline{\mathbb{R}}^t$, if and only if for any $\varepsilon > 0$ there exists a $\delta > 0$ such that for all $y^t \in \overline{\mathbb{R}}^t$ such that for any $\tau \leq t$ either $|x_\tau - y_\tau| < \delta$, or $\max\{x_\tau, y_\tau\} < -1/\delta$, we have that $|f(x^t) - f(y^t)| < \varepsilon$.

Next, consider the payoff of the marginal agent in the perturbed game, $V_t(\bar{x}^t) = V_t(\bar{x}^{t-1}, \bar{x}_t)$. Note that the common posterior $\Psi_t(\theta; \bar{x}^{t-1})$ is continuous in both $\theta \in \mathbb{R}$ and $\bar{x}^{t-1} \in \bar{\mathbb{R}}^{t-1}$. From Bayes' rule, the private posterior are

$$\begin{aligned} \psi_t(\theta|x; \bar{x}^{t-1}) &= \frac{\phi(\sqrt{\alpha_t}(x - \theta)) \psi_t(\theta; \bar{x}^{t-1})}{\int_{-\infty}^{+\infty} \phi(\sqrt{\alpha_t}(x - \theta')) \psi_t(\theta'; \bar{x}^{t-1}) d\theta'} \\ &= \frac{\prod_{s=1}^{t-1} [1 - p_s(\theta; \bar{x}_s)] \phi\left(\sqrt{\alpha_t + \beta}\left(\theta - \frac{\alpha_t x + \beta z}{\alpha_t + \beta}\right)\right)}{\int_{-\infty}^{+\infty} \prod_{s=1}^{t-1} [1 - p_s(\theta'; \bar{x}_s)] \phi\left(\sqrt{\alpha_t + \beta}\left(\theta' - \frac{\alpha_t x + \beta z}{\alpha_t + \beta}\right)\right) d\theta'} \end{aligned}$$

Hence, $\Psi_t(\theta|x; \bar{x}^{t-1}) = \int_{-\infty}^{\theta} \psi_t(\tilde{\theta}|x; \bar{x}^{t-1})$ is also continuous in all arguments. Moreover, since the signal distribution $\Phi(\sqrt{\alpha_t}(x - \theta))$ satisfies the MLRP and is continuous in both $\theta \in \mathbb{R}$ and $x \in \mathbb{R}$, it follows from standard representation theorems that $\Psi_t(\theta|x; \bar{x}^{t-1})$ is decreasing in x . Finally, $p_t(\theta; \bar{x}_t)$ is continuous in θ and \bar{x}_t , increasing in \bar{x}_t , decreasing in θ , and for any \bar{x}_t , $\lim_{\theta \rightarrow \infty} p_t(\theta; \bar{x}_t) = 0$, and $\lim_{\theta \rightarrow -\infty} p_t(\theta; \bar{x}_t) = 1$. It follows that $v_t(x; \bar{x}^t)$ is continuous in both $x \in \mathbb{R}$ and $\bar{x}^t \in \bar{\mathbb{R}}^t$, and decreasing in x . Moreover, $v_t(x; \bar{x}^t)$ is bounded in $[-c, 1 - c]$, which together with the monotonicity of $v_t(x; \bar{x}^t)$ in x ensures the existence of $\lim_{x \rightarrow -\infty} v_t(x; \bar{x}^t)$ and $\lim_{x \rightarrow +\infty} v_t(x; \bar{x}^t)$ for any $\bar{x}^t \in \bar{\mathbb{R}}^t$; this in turn implies that $V_t(\bar{x}^{t-1}, \bar{x}_t)$ is well-defined for $\bar{x}_t = \pm\infty$. Now note that since $v_t(x; \bar{x}^t)$ is continuous in $\bar{x}^{t-1} \in \bar{\mathbb{R}}^{t-1}$, so is $V_t(\bar{x}^t)$. And since $v_t(x; \bar{x}^t) = v_t(x; \bar{x}^{t-1}, \bar{x}_t)$ is continuous in $(x, \bar{x}_t) \in \mathbb{R}^2$, $V_t(\bar{x}^t) = v_t(\bar{x}_t; \bar{x}^{t-1}, \bar{x}_t)$ is also continuous in $\bar{x}_t \in \mathbb{R}$. Moreover,

$$\lim_{\bar{x}_t \rightarrow \infty} V_t(\bar{x}^t) = V_t(\bar{x}^{t-1}, \infty) \equiv \lim_{x \rightarrow \infty} v_t(x; \bar{x}^{t-1}, \infty) = -c$$

for all \bar{x}^{t-1} , which hence proves that $V_t(\bar{x}^{t-1}, \bar{x}_t)$ is continuous in $\bar{x}_t \in \mathbb{R} \cup \{+\infty\}$ ¹¹ On the other hand, $V_t(\bar{x}^{t-1}, \bar{x}_t)$ is not necessarily continuous at $\bar{x}_t = -\infty$, and we can not (and do not need to) rule out the possibility that in equilibrium $\bar{x}_t = -\infty$ for some $t \geq 2$. Nevertheless, note that, by supermodularity,

$$V_t(\bar{x}^{t-1}, -\infty) \equiv \lim_{x_t \rightarrow -\infty} v_t(x; \bar{x}^{t-1}, -\infty) \leq \lim_{x_t \rightarrow -\infty} \sup_{x \in (-\infty, x_t)} V_t(\bar{x}^{t-1}, x),$$

a result that we will use in the rest of the proof.

¹¹To see that $\lim_{x \rightarrow +\infty} v_t(x; \bar{x}^t) = -c$, note that $\Psi_t(\theta|x; \bar{x}^{t-1}, z, \alpha_t, \beta) \leq \Phi\left(\sqrt{\alpha_t + \beta}\left(\theta - \frac{\alpha_t x + \beta z}{\alpha_t + \beta}\right)\right)$ by first-order stochastic dominance. Since $p_t(\theta; \bar{x}_t)$ is decreasing in θ with $\lim_{\theta \rightarrow +\infty} p_t(\theta; \bar{x}_t) = 0$, it follows that $\lim_{x \rightarrow +\infty} v_t(x; \bar{x}^t) \leq \lim_{x \rightarrow +\infty} \int_{-\infty}^{+\infty} p_t(\theta; \bar{x}_t) d\Phi\left(\sqrt{\alpha_t + \beta}\left(\theta - \frac{\alpha_t x + \beta z}{\alpha_t + \beta}\right)\right) - c = -c$.

Step 2: Convergence of V_t to U_t

To prove that for any $\bar{x}^{t-1} \in \overline{\mathbb{R}}^{t-1}$, and any $\bar{x}_t > -\infty$, $V_t(\bar{x}^{t-1}, \bar{x}_t; \delta)$ converges pointwise to $U_t(\bar{x}^{t-1}, \bar{x}_t)$ as $\delta \rightarrow 0$, note that $p_t(\theta; \bar{x}_t) \rightarrow 1$, if $\theta < \theta_t(\bar{x}_t)$, and $p_t(\theta; \bar{x}_t) \rightarrow 0$, if $\theta > \theta_t(\bar{x}_t)$. Therefore, if $t = 1$,

$$\begin{aligned} \lim_{\delta \rightarrow 0} V_1(\bar{x}_1; \delta) &= \int_{-\infty}^{\theta_1(\bar{x}_1)} \psi_1(\theta | \bar{x}_1) d\theta - c = \Psi_1(\theta_1(\bar{x}_1) | \bar{x}_1) - c \\ &= \Phi\left(\sqrt{\alpha_1 + \beta} \left(\theta_1(\bar{x}_1) - \frac{\alpha_1 \bar{x}_1 + \beta z}{\alpha_1 + \beta}\right)\right) - c = U_1(\bar{x}_1) \end{aligned}$$

For $t > 1$, when \bar{x}_t satisfies $\theta_t(\bar{x}_t) \geq \theta_{t-1}(\bar{x}^{t-1}) \equiv \max_{\tau \leq t-1} \{\theta_\tau(\bar{x}_\tau)\}$, then

$$\begin{aligned} \lim_{\delta \rightarrow 0} V_t(\bar{x}^{t-1}, \bar{x}_t; \delta) &= \lim_{\delta \rightarrow 0} \int_{-\infty}^{\theta_t(\bar{x}_t)} \psi_t(\theta | \bar{x}_t; \bar{x}^{t-1}, \delta) d\theta - c = \lim_{\delta \rightarrow 0} \Psi_t(\theta_t(\bar{x}_t) | \bar{x}_t; \bar{x}^{t-1}, \delta) - c \\ &= \frac{\Phi\left(\sqrt{\alpha_t + \beta} \left(\theta_t(\bar{x}_t) - \frac{\alpha_t \bar{x}_t + \beta z}{\alpha_t + \beta}\right)\right) - \Phi\left(\sqrt{\alpha_t + \beta} \left(\theta_{t-1}(\bar{x}^{t-1}) - \frac{\alpha_t \bar{x}_t + \beta z}{\alpha_t + \beta}\right)\right)}{1 - \Phi\left(\sqrt{\alpha_t + \beta} \left(\theta_{t-1}(\bar{x}^{t-1}) - \frac{\alpha_t \bar{x}_t + \beta z}{\alpha_t + \beta}\right)\right)} - c \\ &= U_t(\theta_t(\bar{x}_t); \theta_{t-1}(\bar{x}^{t-1}), \alpha_t, \beta, z) \equiv U_t(\bar{x}^t) \end{aligned}$$

On the other hand, for any $\theta < \theta_{t-1}(\bar{x}^{t-1})$, $\lim_{\delta \rightarrow 0} \Psi_t(\theta | \bar{x}_t; \bar{x}^{t-1}, \delta) = 0$, and therefore

$$\lim_{\delta \rightarrow 0} V_t(\bar{x}^{t-1}, \bar{x}_t; \delta) = U_t(\theta_t(\bar{x}_t); \theta_{t-1}(\bar{x}^{t-1}), \alpha_t, \beta, z) \equiv U_t(\bar{x}^t) = -c$$

for all \bar{x}_t s.t. $\theta_t(\bar{x}_t) < \theta_{t-1}(\bar{x}^{t-1})$. For $\bar{x}_t = -\infty$, however, $V_t(\bar{x}^{t-1}, -\infty; \delta)$ need not converge to $U_t(\bar{x}^t)$ in general, but as we will see below, this is not required for the result.

Step 3: Existence and construction of equilibria for $\Gamma(\delta)$

For any δ , $\lim_{x \rightarrow -\infty} v_1(x; -\infty) > 0$ and $\lim_{x \rightarrow \infty} v_1(x; +\infty) < 0$. This immediately rules out $\bar{x}_1(\delta) = \pm\infty$ in equilibrium. Moreover, together with complementarity, it implies that there exist x'_1 and x''_1 such that $V_1(x) \geq v_1(x; -\infty) > 0$ for any $x < x'_1$ and $V_1(x) \leq v_1(x; +\infty) < 0$ for any $x > x''_1$, which in turn, by continuity of $V_1(x)$, ensures existence of a solution $\bar{x}_1(\delta) \in \mathbb{R}$ to $V_1(\bar{x}_1) = 0$. Next, for every $t \geq 2$ and every \bar{x}^{t-1} , either there exists $x' \in \mathbb{R}$ such that $V_t(\bar{x}^{t-1}, x') \geq 0$, or $V_t(\bar{x}^{t-1}, x_t) < 0$ for all $x_t \in \mathbb{R}$. In the former case, since $\lim_{x_t \rightarrow \infty} V_t(\bar{x}^{t-1}, x_t) = V_t(\bar{x}^{t-1}, +\infty) < 0$, we can also find an $x'' \in \mathbb{R}$ such that $V_t(\bar{x}^{t-1}, x'') < 0$; but then continuity of $V_t(\bar{x}^{t-1}, x_t)$ in x_t ensures the existence of $\bar{x}_t \in (x', x'')$ such that $V_t(\bar{x}^{t-1}, \bar{x}_t) = 0$. In the latter case, $\lim_{x_t \rightarrow -\infty} \sup_{x \in (-\infty, x_t)} V_t(\bar{x}^{t-1}, x) \leq 0$; and since $V_t(\bar{x}^{t-1}, -\infty) \leq \lim_{x_t \rightarrow -\infty} \sup_{x \in (-\infty, x_t)} V_t(\bar{x}^{t-1}, x)$ (by complementarity), we have $V_t(\bar{x}^{t-1}, -\infty) \leq 0$, so that $\bar{x}_t = -\infty$ can be part of an equilibrium.

Step 4: Proof of Convergence

We prove the result by induction. Consider first $T = 1$. From the strict monotonicity of $U_1(x_1)$,¹² for any $\varepsilon > 0$,

$$U_1(x_1^* - \varepsilon) > 0 > U_1(x_1^* + \varepsilon).$$

By the convergence of V_1 to U_1 as $\delta \rightarrow 0$, we can find $\delta_1(\varepsilon) > 0$ such that, for any $\delta < \delta_1(\varepsilon)$,

$$V_1(x_1^* - \varepsilon; \delta) > 0 > V_1(x_1^* + \varepsilon; \delta).$$

From the continuity of $V_1(x_1)$ in x_1 , it follows that there exists a solution $\bar{x}_1(\delta)$ to $V_1(x_1; \delta) = 0$ such that $x_1^* - \varepsilon < \bar{x}_1(\delta) < x_1^* + \varepsilon$. By Step 3, we can then construct an equilibrium $\{\bar{x}_t(\delta)\}_{t=1}^\infty$ for $\Gamma(\delta)$ such that $|\bar{x}_1(\delta) - x_1^*| < \varepsilon$. This proves the result for $T = 1$.

Suppose now the result holds for $T - 1 \geq 1$. Fix an $\varepsilon > 0$ and consider period T .

Assume first $x_T^* > -\infty$. By local strict monotonicity of U_T around x_T^* implied by the fact that $\{x_t^*\}_{t=0}^\infty$ is a strict equilibrium, there exists $\varepsilon_T < \varepsilon$ such that either

$$U_T(x^{*T-1}, x_T^* - \varepsilon_T) > 0 > U_T(x^{*T-1}, x_T^* + \varepsilon_T),$$

or

$$U_T(x^{*T-1}, x_T^* - \varepsilon_T) < 0 < U_T(x^{*T-1}, x_T^* + \varepsilon_T).$$

Without loss of generality, assume the first condition holds; the argument for the second is identical. From the continuity of U_T in all its arguments, the pointwise convergence of V_T to U_T , and the fact that the result holds for $T - 1$, there exists *some* $\delta_T \in (0, \delta(\varepsilon_T, T - 1))$ such that, for any $\delta < \delta_T$, there is a sequence $\bar{x}^{T-1}(\delta)$ such that the following are true:

- (1) $V_1(\bar{x}_1(\delta); \delta) = 0$; and for $\tau \in \{2, \dots, T - 1\}$, either $\bar{x}_\tau(\delta) = -\infty$ and $V_\tau(\bar{x}^\tau(\delta)) < 0$, or $-\infty < \bar{x}_\tau < +\infty$ and $V_\tau(\bar{x}^\tau(\delta); \delta) = 0$;
- (2) for $\tau \in \{1, \dots, T - 1\}$, $|x_\tau^* - \bar{x}_\tau(\delta)| < \varepsilon_T$ or $\max\{x_\tau^*, \bar{x}_\tau(\delta)\} < -1/\varepsilon_T$;
- (3) in period T ,

$$V_T(\bar{x}^{T-1}(\delta), x_T^* - \varepsilon_T; \delta) > 0 > V_T(\bar{x}^{T-1}(\delta), x_T^* + \varepsilon_T; \delta)$$

By the continuity of $V_T(\bar{x}^{T-1}(\delta), x_T; \delta)$ in x_T , there exists $\bar{x}_T(\delta) \in (x_T^* - \varepsilon_T, x_T^* + \varepsilon_T)$ that solves $V_T(\bar{x}^{T-1}(\delta), \bar{x}_T(\delta); \delta) = 0$.

¹²As shown in the previous section, $U_1(x_1)$ is globally strictly decreasing when $\alpha_1 > \beta^2 / (2\pi)$. For the result in the theorem, however, it suffices that $U_1(x_1)$ is locally monotonic at x_1^* , which follows from the assumption that $\{x_t^*\}_{t=0}^\infty$ is a strict equilibrium.

Next, suppose $x_T^* = -\infty$. Recall that, for any $t \geq 2$, $U_t(x^{*t-1}, x_t) = -c < 0$ for all $x_t < \tilde{x}_t$, where $\tilde{x}_t > -\infty$ solves $\theta_t(\tilde{x}_t) = \max_{\tau \leq t-1} \theta_\tau(x_\tau^*)$. Fix now ε and pick some $x'_T < \min\{-1/\varepsilon, \tilde{x}_T\}$. From the continuity of U_T in all its arguments, the pointwise convergence of V_T to U_T , and the fact that the result holds for $T-1$, there exists *some* $\delta_T \in (0, \delta(\varepsilon, T-1))$ such that, for any $\delta < \delta_T$, there is a sequence $\bar{x}^{T-1}(\delta)$ which satisfies conditions (1) and (2) above and in period $t = T$, $V_T(\bar{x}^{T-1}(\delta), x'_T; \delta) < 0$. If in addition there is $x''_T \in (-\infty, x'_T)$ such that $V_T(\bar{x}^{T-1}(\delta), x''_T; \delta) \geq 0$, then by the continuity of V_T in x_T we can find an $\bar{x}_T(\delta) < x'_T < -1/\varepsilon$ such that $V_T(\bar{x}^{T-1}(\delta), \bar{x}_T(\delta); \delta) = 0$. If instead $V_T(\bar{x}^{T-1}(\delta), x_T; \delta) < 0$ for all $x_T \in (-\infty, x'_T)$, then $V_T(\bar{x}^{T-1}(\delta), -\infty) \leq \lim_{x_T \rightarrow -\infty} \sup_{x \in (-\infty, x_T)} V_T(\bar{x}^{T-1}(\delta), x) \leq 0$, in which case we let $\bar{x}_T(\delta) = -\infty$.

Finally, note that the set of x_T^* that can be part of a strict equilibrium is finite. This implies that we can always find a $\delta'(\varepsilon) > 0$ such that, for $\delta < \delta'(\varepsilon)$ the above results hold for every strict equilibrium of $\Gamma(0)$. Letting then $\delta(\varepsilon, T) = \min\{\delta(\varepsilon, T-1), \delta'(\varepsilon)\}$ and using Step 3, for any strict equilibrium $\{x_t^*\}_{t=1}^\infty$ of $\Gamma(0)$ we can construct an entire equilibrium $\{\bar{x}_t(\delta)\}_{t=1}^\infty$ for $\Gamma(\delta)$ such that either $|\bar{x}_t(\delta) - x_t^*| < \varepsilon$ or $\max\{x_t^*, \bar{x}_t(\delta)\} < -1/\varepsilon$ for all $t \leq T$, which completes the proof of part (i). ■

Proof of Theorem (4). For any $\theta_{-1}, \alpha, \beta$, and any $\theta > \theta_{-1}$, $U(\theta; \theta_{-1}, z, \alpha, \beta) \rightarrow 1 - c > 0$ as $z \rightarrow -\infty$ and $U(\theta; \theta_{-1}, z, \alpha, \beta) \rightarrow -c < 0$ as $z \rightarrow +\infty$. By the monotonicity of U in z , it follows that there exists $\bar{z}(\theta_{-1}, \alpha, \beta)$ such that $\max_{\theta \in [\theta_{-1}, 1]} U(\theta; \theta_{-1}, z, \alpha, \beta) \geq 0$, if and only if $z \leq \bar{z}(\theta_{-1}, \alpha, \beta)$. Therefore a solution to (??) exists in period t , if and only if $z_t \leq \bar{z}(\theta_{t-1}^*, \alpha_t, \beta_t)$. Since $z_t \leq \bar{z}(\theta_{t-1}^*, \alpha_t, \beta_t)$ with positive probability in any period t , this immediately implies the result. ■

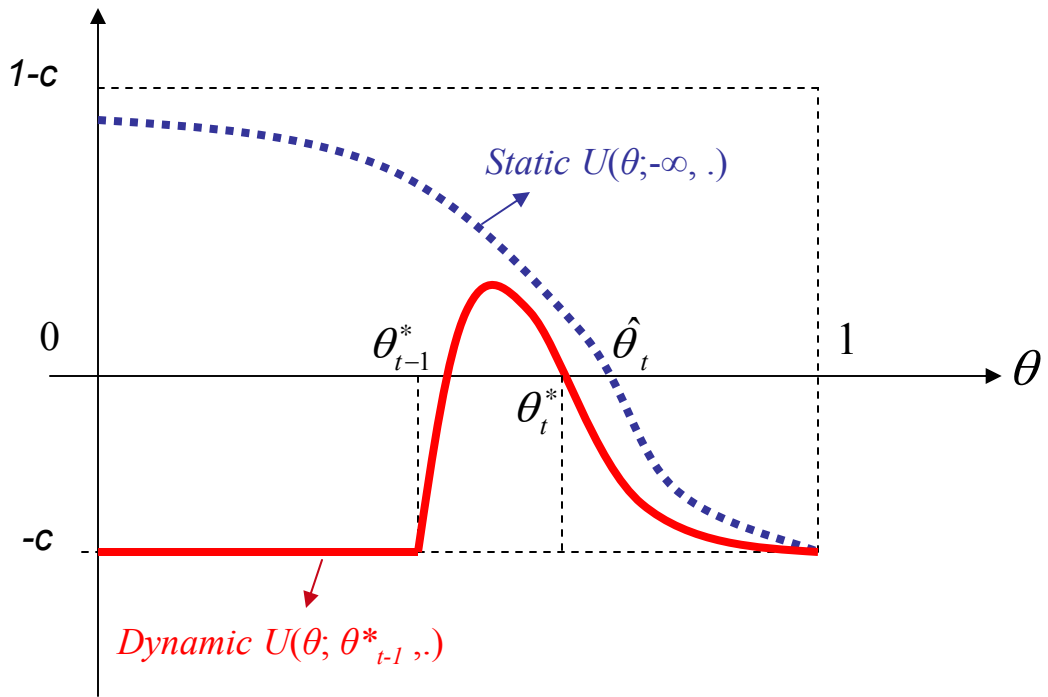


Figure 1: Static and dynamic payoffs for the marginal agent