

Angeletos, George-Marios; Hellwig, Christian; Pavan, Alessandro

Working Paper

Dynamic global games of regime change: Learning, multiplicity, and timing of attacks

Discussion Paper, No. 1497

Provided in Cooperation with:

Kellogg School of Management - Center for Mathematical Studies in Economics and Management Science, Northwestern University

Suggested Citation: Angeletos, George-Marios; Hellwig, Christian; Pavan, Alessandro (2006) : Dynamic global games of regime change: Learning, multiplicity, and timing of attacks, Discussion Paper, No. 1497, Northwestern University, Kellogg School of Management, Center for Mathematical Studies in Economics and Management Science, Evanston, IL

This Version is available at:

<https://hdl.handle.net/10419/59643>

Standard-Nutzungsbedingungen:

Die Dokumente auf EconStor dürfen zu eigenen wissenschaftlichen Zwecken und zum Privatgebrauch gespeichert und kopiert werden.

Sie dürfen die Dokumente nicht für öffentliche oder kommerzielle Zwecke vervielfältigen, öffentlich ausstellen, öffentlich zugänglich machen, vertreiben oder anderweitig nutzen.

Sofern die Verfasser die Dokumente unter Open-Content-Lizenzen (insbesondere CC-Lizenzen) zur Verfügung gestellt haben sollten, gelten abweichend von diesen Nutzungsbedingungen die in der dort genannten Lizenz gewährten Nutzungsrechte.

Terms of use:

Documents in EconStor may be saved and copied for your personal and scholarly purposes.

You are not to copy documents for public or commercial purposes, to exhibit the documents publicly, to make them publicly available on the internet, or to distribute or otherwise use the documents in public.

If the documents have been made available under an Open Content Licence (especially Creative Commons Licences), you may exercise further usage rights as specified in the indicated licence.

Dynamic Global Games of Regime Change: Learning, Multiplicity, and Timing of Attacks*

George-Marios Angeletos
MIT and NBER

Christian Hellwig
UCLA

Alessandro Pavan
Northwestern University

August 2006

Abstract

Global games of regime change—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 as currency attacks, bank runs, debt crises, and political change. We extend the static benchmark examined in the literature by allowing agents to take actions in many periods and to learn about the underlying fundamentals over time. We first provide a simple recursive algorithm for the characterization of monotone equilibria. We then show how the interaction of the knowledge that the regime survived past attacks with the arrival of information over time, or with changes in fundamentals, leads to interesting equilibrium properties. First, multiplicity may obtain under the same conditions on exogenous information that guarantee uniqueness in the static benchmark. Second, fundamentals may predict the eventual regime outcome but not the timing or the number of attacks. Finally, equilibrium dynamics can alternate between phases of tranquility—where no attack is possible—and phases of distress—where a large attack can occur—even without changes in fundamentals.

JEL Codes: C7, D7, D8, F3.

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

*Earlier versions of this paper were entitled “Information Dynamics and Equilibrium Multiplicity in Global Games of Regime Change.” We are grateful to the editor, Andrew Postlewaite, and three anonymous referees for suggestions that greatly helped us improve the paper. For useful comments, we also thank Andy Atkeson, Daron Acemoglu, Pierpaolo Battigalli, Alberto Bisin, V.V. Chari, Lars Hansen, Patrick Kehoe, Alessandro Lizzeri, Kiminori Matsuyama, Stephen Morris, Hyun Song Shin, Iván Werning, and seminar participants at Berkeley, Bocconi, Bologna, British Columbia, Chicago, MIT, Northwestern, NYU, Pompeu Fabra, Princeton, UCL, UCLA, UPenn, Yale, the Minneapolis FRB, the 2003 and 2004 SED meetings, the 2005 CEPR-ESSET, the 2005 IDEI conference in tribute to J.J. Laffont, the 2005 NBER Summer Institute, the 2005 Cowles workshop on coordination games, and the 2005 World Congress of the Econometric Society. *Emails:* angelet@mit.edu, chris@econ.ucla.edu, alepavan@northwestern.edu

1 Introduction

Games of regime change are coordination games in which a status quo is abandoned, causing a discrete change in payoffs, once a sufficiently large number of agents takes an action against it. These games have been used to model a variety of crises phenomena: an attack against the status quo is interpreted as speculation against a currency peg, as a run against a bank, or as a revolution against a dictator.

Most applications of these games to crises have been confined to static frameworks: they abstract from the possibility that agents take multiple shots against the status quo and that their beliefs about their ability to induce regime change vary over time.¹ Yet, these two possibilities are important from both an applied and a theoretical perspective. First, crises are intrinsically dynamic phenomena. In the context of currency crises, for example, speculators can attack a currency again and again until they induce devaluation; and their expectations about the ability of the central bank to defend the currency in the present may naturally depend on whether the bank has successfully defended it in the past. Second, learning in a dynamic setting may critically affect the level of strategic uncertainty (i.e., uncertainty about one another’s actions) and thereby the dynamics of coordination and the determinacy of equilibria.

In this paper, we consider a dynamic global game that extends the static benchmark used in the literature so as to incorporate precisely the two possibilities highlighted above.² There is a large number of agents and two possible regimes, the status quo and an alternative. The game continues as long as the status quo is in place. In each period, each agent can either *attack* the status quo (i.e., take an action that favors regime change), or *not attack*. The net payoff from attacking is positive if the status quo is abandoned in that period and negative otherwise. Regime change, in turn, occurs if and only if the fraction of agents attacking exceeds a threshold $\theta \in \mathbb{R}$ that parameterizes the strength of the status quo. θ captures the component of the payoff structure (the “fundamentals”) that is never common knowledge; as time passes, agents receive more and more private information about θ .

We first provide an algorithm for the characterization of monotone equilibria, based on a simple recursive structure. A difficulty with extending global games to dynamic settings is the need to keep track of the evolution of the cross-sectional distribution of beliefs. Our framework overcomes this difficulty by summarizing the private information of an agent about θ at any given period in a one-dimensional sufficient statistic, and capturing the dynamics of the cross-sectional distribution

¹This is particularly true for recent applications that introduce incomplete information. See Morris and Shin (1998) for currency crises; Goldstein and Pauzner (2005) and Rochet and Vives (2004) for bank runs; Morris and Shin (2004) and Corsetti, Guimaraes and Roubini (2005) for debt crises; Atkeson (2000) and Edmond (2005) for riots and political change.

²Global games are incomplete-information games that often admit a unique, iteratively-dominant equilibrium; see Carlsson and Van Damme (1993) and Morris and Shin (2003). The applications cited in footnote 1 are all based on one-shot global games.

of this statistic in a parsimonious way. We then apply this algorithm to examine the effects of learning on the determinacy of equilibria and the dynamics of coordination.

Multiplicity. We find that multiple equilibria can exist in our dynamic game under the same conditions on the precision of exogenous private and public information that would guarantee uniqueness in the static benchmark that is the backbone of most recent applications of global games (Morris and Shin, 1998, 2001, 2003). Multiplicity originates in the interaction between the endogenous learning induced by the knowledge that the regime survived past attacks and the exogenous learning induced by the arrival of new private information over time.

Iterated deletion of dominated strategies ensures that equilibrium play is uniquely determined in the first period: an attack necessarily takes place for every θ , but succeeds in triggering regime change if and only if θ is sufficiently low. In any subsequent period, the knowledge that the status quo is still in place makes it common certainty that it is not too weak, and ensures that no agent ever again finds it dominant to attack. As a result, there always exists an equilibrium in which no attack occurs after the first period. This would actually be the unique equilibrium of the game if agents did not receive any information after the first period.

When instead new private information about θ arrives over time, this has two effects on posterior beliefs about the strength of the status quo and hence on the agents' incentives to attack. On the one hand, it dilutes the upward shift in posterior beliefs induced by the knowledge that the regime survived the first-period attack, which contributes to making further attacks possible. On the other hand, it reduces the dependence of posterior beliefs on the common prior, which in general has an ambiguous effect. When the prior mean is high (i.e., favorable to the status quo), discounting the prior also contributes to making a new attack possible; the opposite is true when the prior mean is low. A high prior mean thus ensures existence of an equilibrium where a second attack occurs once private information becomes sufficiently precise.

More generally, we show that, when the prior mean is sufficiently high, the arrival of private information over time suffices for the existence of arbitrarily many equilibria, which differ in both the number and the timing of attacks.

Dynamics of coordination. The multiplicity discussed above does not mean that “anything goes”: equilibrium outcomes in any given period depend critically on available information and the history of past play.

The learning induced by the knowledge that the status quo survived past attacks introduces a form of strategic substitutability across periods: the more aggressive the agents' strategy in one period, the higher the threshold in θ below which regime change occurs in that period; but then the larger the upward shift in posterior beliefs induced by the knowledge that the regime survived this attack, and hence the lower the incentive to attack in subsequent periods. When an aggressive attack takes place in one period but fails to trigger regime change, then a significant increase in the precision of private information is necessary to offset the endogenous upward shift in posterior beliefs and make a new attack possible in equilibrium. As a result, dynamics take the form

of sequences of periods in which attacks can not occur and agents only accumulate information, followed by periods in which an attack is possible but does not materialize, eventually resulting in a new attack.

Moreover, although it is possible that attacks continue indefinitely as long as new information arrives over time, strategic uncertainty significantly limits the size of attacks. For θ high enough, the status quo may survive forever, independently of which equilibrium is played and despite the fact that it might have been vulnerable to a sufficiently strong attack.

Implications for crises. These results translate to interesting predictions for the dynamics of crises. First, fundamentals may determine eventual outcomes—e.g., whether a currency is devalued—but not the timing and number of attacks. Second, an economy can transit from phases of “tranquility”, where the unique possible outcome is no attack, to phases of “distress”, where a significant change in outcomes can be triggered by a shift in “market sentiments”. Finally, the transition from one phase to another can be caused by a small change in information or, in a later extension, by a small change in fundamentals.

These predictions strike a delicate balance between two alternative views of crises. The first associates crises with multiple self-fulfilling equilibria: large and abrupt changes in outcomes are attributed to shifts in “market sentiments” or “animal spirits” (Obstfeld, 1996). The second associates crises with a discontinuity, or strong non-linearity, in the dependence of the unique equilibrium to exogenous variables: large and abrupt changes in outcomes are attributed to small changes in fundamentals, or in agents’ information (Morris and Shin, 1998, 2001). Our results combine a refined role for multiple self-fulfilling expectations with a certain discontinuity in equilibrium outcomes with respect to information and fundamentals.

Extensions. The benchmark model focuses on the arrival of private information as the only exogenous source of change in beliefs. In an extension we show how the analysis can accommodate public news about the underlying fundamentals. This only reinforces the multiplicity result. Moreover, equilibrium dynamics continue to be characterized by phases of tranquility and phases of distress, but now the transition from one phase to another can be triggered by public news.

The benchmark model also deliberately assumes away the possibility that the critical size of attack that triggers regime change may vary over time. This permits us to isolate the impact of changes in information (beliefs), as opposed to changes in fundamentals (payoffs), on the dynamics of coordination. Nevertheless, introducing shocks to fundamentals is important for applications, as well as for understanding the robustness of our results.

We first examine the case in which the shocks are perfectly observable. Shocks then provide an additional driving force for dynamics: a transition from tranquility to distress may now be triggered by a deterioration in fundamentals. Moreover, a sufficiently bad shock can push the economy into a phase where an attack becomes inevitable—a possibility absent in the benchmark model.

We next consider the case in which the shocks are unobservable (or observed with private noise). The novel effect is that the uncertainty about the shocks “noises up” the learning induced by the

knowledge that the regime survived past attacks: whereas in the benchmark model this knowledge leads to a truncation in the support of posterior beliefs about the strength of the status quo, here posterior beliefs retain full support over the entire real line. Thus, in contrast to the benchmark model, agents with very low signals may now find it dominant to attack in every period; and, other things equal, a unique equilibrium outcome may obtain in any given period when private information in that period is sufficiently precise. Nevertheless, our results are robust as long as the noise in learning is small: any equilibrium of the benchmark model is approximated arbitrarily well by an equilibrium of the game with shocks as the volatility of the shocks vanishes.

Thus, what sustains the multiplicity of equilibria and the structure of dynamics identified in this paper is the combination of exogenous changes in information or fundamentals with the endogenous learning induced by the knowledge that the regime survived past attacks. That in the benchmark model this learning takes the sharp form of a truncation in the support of beliefs simplifies the analysis but is not essential for the results. What is essential is that this learning implies a significant change in common beliefs about the strength of the status quo.

Related literature. This paper contributes to the literature on global games by highlighting the importance of learning for equilibrium determinacy. In this respect, it shares with Angeletos, Hellwig and Pavan (2006)—which considers the signaling effects of policy interventions in a static environment—the idea that natural sources of *endogenous* information may qualify the applicability of global-game uniqueness results, while at the same time reinforcing the more general point that information is important for coordination. In our framework this leads to novel predictions that would have not been possible with either common knowledge or a unique equilibrium.³

The paper also contributes to a small but growing literature on *dynamic* global games. Morris and Shin (1999) consider a dynamic model whose stage game is similar to ours, but where the strength of the status quo follows a random walk and is commonly observed at the end of each period. This reduces the analysis to a sequence of essentially unrelated static games, each with a unique equilibrium. Heidhues and Melissas (2006) and Giannitsarou and Toxvaerd (2003) establish uniqueness results for dynamic global games on the basis of dynamic strategic complementarities. Dasgupta (2006) examines the role of noisy social learning in a two-period investment model with irreversible actions. Levin (2001) considers a global game with overlapping generations of players. Goldstein and Pauzner (2004) and Goldstein (2005) consider models of contagion. Frankel and Pauzner (2000) examine a dynamic coordination game where uniqueness is obtained by combining aggregate shocks with idiosyncratic inertia. Abreu and Brunnermeier (2003) consider a setting in which speculators become gradually and asymmetrically aware of the mispricing of a financial asset. All these papers feature multi-period coordination problems; but none of them features the form

³Information is endogenized also in Angeletos and Werning (2006), Hellwig, Mukherji and Tsyvinski (2006), and Tarashev (2005), where financial prices aggregate and publicize disperse private information, and in Edmond (2005), where a dictator manipulates the distribution of private signals.

of learning that is the center of our analysis.⁴ Our methodological approach is also quite different: instead of forcing uniqueness, we wish to understand how a natural form of learning sustained by repeated play may affect both the determinacy of equilibria and the structure of dynamics.

Finally, this paper shares with Chari and Kehoe (2003) the motivation that information is important for understanding crises: our benchmark model offers a theory where changes in information are the sole source for the dynamics of crises. However, there are two important differences. First, Chari and Kehoe focus on the effect of herding in an environment without strategic complementarities. In contrast, we focus on the impact of learning on the dynamics of coordination. The coordination element is crucial for the prediction that there is a phase of distress during which an attack is possible but does not necessarily take place, as well as for the prediction that attacks occur as sudden and synchronized events. Second, the main learning effect in Chari and Kehoe is the negative information about the fundamentals revealed by the choice by some agents to attack—a form of learning that generates “build-up” or “snow-balling” effects. In contrast, the main learning effect in our benchmark model is the positive information revealed by the failure of an attack to trigger regime change—a form of learning that is crucial for our prediction that phases of distress are eventually followed by phases of tranquility. In Section 5.2 we discuss an extension of our benchmark model in which agents observe noisy signals about the size of past attacks. This extension combines our cycles between phases of distress and tranquility with snow-balling effects similar to those stressed in the herding literature.

The rest of the paper is organized as follows. Section 2 reviews the static benchmark and introduces the dynamic model. Section 3 characterizes the set of monotone equilibria. Section 4 establishes the multiplicity result and examines the properties of equilibrium dynamics. Section 5 considers a few extensions of the benchmark model and examines robustness. Section 6 concludes. Proofs omitted in the main text are in the Appendix.

2 A simple game of regime change

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 zero, whereas the payoff from attacking ($a_i = 1$) is $1 - c > 0$ if the status quo is abandoned ($R = 1$) and $-c < 0$ otherwise ($R = 0$), where $c \in (0, 1)$ parameterizes the relative cost of attacking. An agent hence finds it optimal to attack if and only if he expects regime change with probability at least

⁴Chamley (2003) also considers learning in a dynamic coordination game. However, his model is not a global game; all information is public, and so is learning.

equal to c . The status quo, in turn, is abandoned if and only if the measure of agents attacking, which we denote by A , is no less than a critical value $\theta \in \mathbb{R}$, which parameterizes the strength of the status quo. An agent's incentive to attack thus increases with the aggregate size of attack, implying that agents' actions are strategic complements.⁵

	<i>Regime Change</i> ($A \geq \theta$)	<i>Status Quo</i> ($A < \theta$)
<i>Attack</i> ($a_i = 1$)	$1 - c$	$-c$
<i>Not Attack</i> ($a_i = 0$)	0	0

Table I: Payoffs

Agents have heterogeneous information about the strength of the status quo. Nature first draws θ from a normal distribution $\mathcal{N}(z, 1/\alpha)$, which defines the initial common prior about θ . Each agent then receives a private signal $x_i = \theta + \xi_i$, where $\xi_i \sim \mathcal{N}(0, 1/\beta)$ is noise, i.i.d. across agents and independent of θ . The Normality assumptions allow us to parameterize the information structure parsimoniously with (β, α, z) , that is, the precision of private information and the precision and the mean of the common prior.

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 to roll over its debt and is hence forced into 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. (For references, see footnote 1.)

Equilibrium analysis. Note that the c.d.f. of an agent's posterior about θ is decreasing in his private signal x . Moreover, it is strictly dominant to attack for sufficiently low signals (namely for $x < \underline{x}$, where \underline{x} solves $\Pr(\theta \leq 0 | \underline{x}) = c$) and not to attack for sufficient high signals (namely for $x > \bar{x}$, where \bar{x} solves $\Pr(\theta \leq 1 | \bar{x}) = c$). It is thus natural to look at *monotone* Bayesian Nash equilibria in which the agents' strategy is non-increasing in x .

Indeed, suppose there is a threshold $\hat{x} \in \mathbb{R}$ 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

$$A(\theta) = \Pr(x \leq \hat{x} | \theta) = \Phi(\sqrt{\beta}(\hat{x} - \theta)),$$

⁵The role of coordination is most evident when θ is commonly known by all agents: for $\theta \in (0, 1]$, there exist two pure-strategy equilibria, one in which all agents attack and the status quo is abandoned ($A = 1 \geq \theta$) and another in which no agent attacks and the status quo is maintained ($A = 0 < \theta$).

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 $\hat{\theta} = A(\hat{\theta})$, or equivalently

$$\hat{\theta} = \Phi(\sqrt{\beta}(\hat{x} - \hat{\theta})). \quad (1)$$

By standard Gaussian updating, the posterior about θ conditional on private signal x is Normal with mean $\frac{\beta}{\beta+\alpha}x + \frac{\alpha}{\beta+\alpha}z$ and precision $\beta + \alpha$. It follows that the posterior probability of regime change is simply

$$\Pr(R = 1|x) = \Pr(\theta \leq \hat{\theta}|x) = 1 - \Phi\left(\sqrt{\beta + \alpha}\left(\frac{\beta}{\beta+\alpha}x + \frac{\alpha}{\beta+\alpha}z - \hat{\theta}\right)\right).$$

Since the latter is decreasing in x , an agent finds it optimal to attack if and only if $x \leq \hat{x}$, where \hat{x} solves $\Pr(\theta \leq \hat{\theta}|\hat{x}) = c$, or equivalently

$$1 - \Phi\left(\sqrt{\beta + \alpha}\left(\frac{\beta}{\beta+\alpha}\hat{x} + \frac{\alpha}{\beta+\alpha}z - \hat{\theta}\right)\right) = c. \quad (2)$$

A monotone equilibrium is thus identified by a joint solution $(\hat{x}, \hat{\theta})$ to (1) and (2). Such a solution always exists and is unique for all z if and only if $\beta \geq \alpha^2 / (2\pi)$. Moreover, iterated elimination of strictly dominated strategies implies that, when the monotone equilibrium is unique, there is no other equilibrium.

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

In the limit as $\beta \rightarrow \infty$ for given α , the threshold $\hat{\theta}$ converges to $\theta_\infty \equiv 1 - c$, and the size of attack $A(\theta)$ converges to 1 for all $\theta < \theta_\infty$ and to 0 for all $\theta > \theta_\infty$. Hence, when the noise in private information is small and θ is in the neighborhood of θ_∞ , a small variation in θ can trigger a large variation in the size of attack and in the regime outcome. This kind of discontinuity, or strong non-linearity, in the response of equilibrium outcomes to exogenous variables underlies the view of crises advocated by most global-game applications.⁶

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.

Time is discrete and indexed by $t \in \{1, 2, \dots\}$. The game continues as long as the status quo is in place and is over once the status quo is abandoned. We denote by $R_t = 0$ the event that the status quo is in place at the beginning of period t , by $R_t = 1$ the alternative event, by $a_{it} \in \{0, 1\}$ the action of agent i , and by $A_t \in [0, 1]$ the measure of agents attacking at date t . Conditional on

⁶A related strong non-linearity emerges in the response of equilibrium outcomes to noise in public information; see the discussion of the “publicity multiplier” in Morris and Shin (2003) and that of “non-fundamental volatility” in Angeletos and Werning (2005).

the regime being in place at the beginning of period t ($R_t = 0$), the regime is abandoned in that period ($R_{t+1} = 1$) if and only if $A_t \geq \theta$, where θ again represents the strength of the status quo. Agent i 's flow payoff for period t (conditional on $R_t = 0$) is thus $\pi_{it} = a_{it}(R_{t+1} - c)$, while his payoff from 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 discount factor.

Like in the static model, θ is drawn at the beginning of the game from $\mathcal{N}(z, 1/\alpha)$, which defines the initial common prior, and never becomes common knowledge. Private information, however, evolves over time. 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, hence the public history in period t simply consists of the information that the regime is still in place, whereas the private history of an agent is the sequence of own private signals and own past actions. Finally, we let $\beta_t \equiv \sum_{\tau=1}^t \eta_{\tau}$ and assume that

$$\infty > \beta_t \geq \alpha^2/(2\pi) \quad \forall t \quad \text{and} \quad \lim_{t \rightarrow \infty} \beta_t = \infty.$$

As shown in the next section, β_t parameterizes the precision of private information accumulated up to period t . The assumptions we make here for β_t ensure (i) that the static game defined by the restriction that agents can move only in period t has a unique equilibrium for every t , and (ii) that private information becomes infinitely precise only in the limit.

Remark. While this dynamic 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 strength of the status quo. 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, for example, the role of wealth accumulation or liquidity in currency crises. However, abstracting from these other dimensions allows us to isolate information as the driving force for the dynamics of coordination and crises.

Equilibrium. In what follows, we limit attention to *monotone equilibria*, that is, symmetric Perfect Bayesian equilibria in which the probability an agent attacks in period t , which we denote by $a_t(\tilde{x}^t)$, is non-increasing in his private signals \tilde{x}^t and independent of his own past actions.⁷ Restricting attention to this class of equilibria suffices to establish our results.

3 Equilibrium characterization

Let $a_t : \mathbb{R}^t \rightarrow [0, 1]$ denote the strategy for period t and $a^t = \{a_{\tau}\}_{\tau=1}^t$ the strategy up to period t , with $a^{\infty} = \{a_{\tau}\}_{\tau=1}^{\infty}$ denoting the complete strategy for the dynamic game. Since \tilde{x}^t is i.i.d. across agents conditional on θ , for any given strategy a^{∞} the size of attack and the regime outcome in period t depend only on θ . Thus let $p_t(\theta; a_t)$ denote the probability that the status quo is abandoned

⁷We do *not* restrict the set of available strategies: we look at equilibria in which these properties are satisfied.

in period t when all agents follow the strategy a_t , conditional on the status quo being in place at the beginning of period t and the fundamentals being θ . Finally, let $\Psi_1(\theta|\tilde{x}_1)$ denote the c.d.f. of the posterior beliefs in period 1, while for any $t \geq 2$ let $\Psi_t(\theta|\tilde{x}^t; a^{t-1})$ denote the c.d.f. of the posterior beliefs in period t conditional on the knowledge that the status quo is still in place (i.e., $R_t = 0$) and that agents have played in past periods according to a^{t-1} .

Since neither individual nor aggregate actions are observable, and $R_t = 0$ is always compatible with any strategy profile at any t , no agent can detect out-of-equilibrium play as long as the status quo is in place.⁸ It follows that beliefs are pinned down by Bayes' rule in any relevant history of the game. Furthermore, as long as the status quo is in place, payoffs in one period do not depend on own or other players' actions in any other period, and hence strategies are sequentially rational if and only if the action prescribed for any given period maximizes the payoff for that period. We conclude that the strategy $a^\infty = \{a_t\}_{t=1}^\infty$ is part of an equilibrium if and only if the following hold: at $t = 1$, for all \tilde{x}_1 ,

$$a_1(\tilde{x}_1) \in \arg \max_{a \in [0,1]} \left\{ \left[\int p_1(\theta; a_1) d\Psi_1(\theta|\tilde{x}_1) - c \right] a \right\}; \quad (3)$$

and at any $t \geq 2$, for all \tilde{x}^t ,

$$a_t(\tilde{x}^t) \in \arg \max_{a \in [0,1]} \left\{ \left[\int p_t(\theta; a_t) d\Psi_t(\theta|\tilde{x}^t; a^{t-1}) - c \right] a \right\}. \quad (4)$$

Next, define x_t and β_t recursively by

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

with $x_1 = \tilde{x}_1$ and $\beta_1 = \eta_1$. By standard Gaussian updating, the distribution of θ conditional on $\tilde{x}^t = \{\tilde{x}_\tau\}_{\tau=1}^t$ is Normal with mean $\frac{\beta_t}{\alpha + \beta_t} x_t + \frac{\alpha}{\alpha + \beta_t} z$ and precision $\beta_t + \alpha$. It follows that x_t is a sufficient statistic for \tilde{x}^t with respect to θ , and hence with respect to the event of regime change as well. As we show below (and further discuss in Section 5.5), this ability to summarize private information into a one-dimensional sufficient statistic greatly simplifies the analysis.

Clearly, condition (3) implies that in *any* equilibrium of the dynamic game agents play in the first period exactly as in the static game in which they can attack only at $t = 1$. Hence, by Proposition 1, equilibrium play is uniquely determined in the first period and is characterized in terms of thresholds for x_1 and θ . The following lemma shows that a similar property holds for subsequent periods.⁹

Lemma 1 *Any monotone equilibrium is characterized by a sequence $\{x_t^*, \theta_t^*\}_{t=1}^\infty$, where $x_t^* \in \mathbb{R} \cup \{-\infty\}$, $\theta_t^* \in (0, 1)$, and $\theta_t^* \geq \theta_{t-1}^*$ for all $t \geq 2$, such that:*

- (i) *at any $t \geq 1$, an agent attacks if $x_t < x_t^*$ and does not attack if $x_t > x_t^*$;*
- (ii) *the status quo is in place in period $t \geq 2$ if and only if $\theta > \theta_{t-1}^*$.*

⁸Indeed, the regime always survives any attack for $\theta > 1$ and no realization of the private signal rules out $\theta > 1$.

⁹To simplify the notation, we allow for $x_t^* = -\infty$ and $x_t^* = +\infty$, with which we denote the case where an agent attacks for, respectively, none and every realization of his private information.

Proof. We prove the claim by induction. For $t = 1$, the result follows from Proposition 1. Consider next any $t \geq 2$ and suppose that the result holds for any $\tau \leq t - 1$. Since a_t is non-increasing in \tilde{x}^t , the size of attack $A_t(\theta)$ is non-increasing in θ , implying that either $A_t(\theta) < \theta$ (and therefore $R_{t+1} = 0$) for all $\theta > \theta_{t-1}^*$, in which case $\theta_t^* = \theta_{t-1}^*$, or there exists $\theta_t^* > \theta_{t-1}^*$ such that $A_t(\theta) < \theta$ if and only if $\theta > \theta_t^*$. In the former case, the posterior probability of regime change is 0 for all x_t and hence $x_t^* = -\infty$. In the latter, the posterior probability of regime change is given by

$$\int p_t(\theta; a_t) d\Psi_t(\theta | \tilde{x}^t; a^{t-1}) = \Pr(\theta \leq \theta_t^* | x_t, \theta > \theta_{t-1}^*) = 1 - \frac{\Phi\left(\sqrt{\beta_t + \alpha} \left[\frac{\beta_t x_t + \alpha z}{\beta_t + \alpha} - \theta_t^*\right]\right)}{\Phi\left(\sqrt{\beta_t + \alpha} \left[\frac{\beta_t x_t + \alpha z}{\beta_t + \alpha} - \theta_{t-1}^*\right]\right)}; \quad (5)$$

and since this is continuous and strictly decreasing in x_t , and converges to 1 as $x_t \rightarrow -\infty$ and to 0 as $x_t \rightarrow +\infty$, there exists $x_t^* \in \mathbb{R}$ such that $\Pr(\theta \leq \theta_t^* | x_t, \theta > \theta_{t-1}^*) = c$ for $x_t = x_t^*$, $\Pr(\theta \leq \theta_t^* | x_t, \theta > \theta_{t-1}^*) > c$ for $x_t < x_t^*$ and $\Pr(\theta \leq \theta_t^* | x_t, \theta > \theta_{t-1}^*) < c$ for $x_t > x_t^*$. In either case, $A_t(\theta) < 1$ for all θ and hence $\theta_t^* < 1$, which together with $\theta_t^* \geq \theta_1^* > 0$, implies that $\theta_t^* \in (0, 1)$ for all t , which completes the proof. ■

Clearly, since the status quo can not be in place in one period without also being in place in the previous, the sequence $\{\theta_t^*\}$ is non-decreasing. On the other hand, the sequence $\{x_t^*\}$ is non-monotonic in general: periods where some agents attack ($x_t^* > -\infty$) may indefinitely alternate with periods where nobody attacks ($x_t^* = -\infty$).

As mentioned above, the first period in our dynamic game is similar to the static game; but any subsequent period is very different. In any $t \geq 2$, the fact that the status quo is still in place makes it common certainty that $\theta > \theta_{t-1}^*$.¹⁰ Since $\theta_{t-1}^* \geq \theta_1^* > 0$, this immediately implies that there always exist 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 one and never thereafter. 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 that in the static benchmark. In what follows we thus examine under what conditions there also exist equilibria with further attacks.

Lemma 1 rules out $x_t^* = +\infty$ (situations where everybody attacks). This follows directly from the fact that the status quo always survives for $\theta > 1$ and hence it is dominant not to attack for x_t sufficiently high. We thus look for equilibria in which $x_t^* \in \mathbb{R}$.

The size of attack is then given by $A_t(\theta) = \Pr(x_t \leq x_t^* | \theta) = \Phi(\sqrt{\beta_t}(x_t^* - \theta_t))$, which is continuous and strictly decreasing in θ , while the probability of regime change for an agent with

¹⁰Clearly, the knowledge that the regime is in place in period t is a form of public information. However, this is very different from the type of information conveyed by additive public signals of θ (Morris and Shin, 2001, 2003; Hellwig, 2002). First, the information here is endogenous, as it depends on the particular equilibrium being played; and second, it leads to a first-order stochastic-dominance shift in beliefs. We introduce additive public signals about θ in Section 5.1.

statistic x_t is given by (5), which is continuous and strictly decreasing in x_t if $\theta_t^* > \theta_{t-1}^*$. It follows that, in any equilibrium in which an attack occurs in period t , θ_t^* and x_t^* solve $\theta_t^* = A_t(\theta_t^*)$ and $\Pr(\theta \leq \theta_t^* | x_t^*, \theta > \theta_{t-1}^*) = c$, or equivalently

$$\theta_t^* = \Phi(\sqrt{\beta_t}(x_t^* - \theta_t^*)), \quad (6)$$

$$1 - \frac{\Phi\left(\sqrt{\beta_t + \alpha}\left(\frac{\beta_t}{\beta_t + \alpha}x_t^* + \frac{\alpha}{\beta_t + \alpha}z - \theta_t^*\right)\right)}{\Phi\left(\sqrt{\beta_t + \alpha}\left(\frac{\beta_t}{\beta_t + \alpha}x_t^* + \frac{\alpha}{\beta_t + \alpha}z - \theta_{t-1}^*\right)\right)} = c. \quad (7)$$

Conditions (6) and (7) are the analogs in the dynamic game of conditions (1) and (2) in the static game: (6) states that the equilibrium size of an attack is equal to the critical size that triggers regime change if and only if the fundamentals are θ_t^* , while (7) states that an agent is indifferent between attacking and not attacking if and only if his private information is x_t^* .

An alternative representation of the equilibrium conditions is also useful. Define the functions $u : \mathbb{R} \times [0, 1] \times \overline{\mathbb{R}} \times \mathbb{R}_+^2 \times \mathbb{R} \rightarrow [-c, 1 - c]$, $X : [0, 1] \times \mathbb{R}_+ \rightarrow \overline{\mathbb{R}}$, and $U : [0, 1] \times \overline{\mathbb{R}} \times \mathbb{R}_+^2 \times \mathbb{R} \rightarrow [-c, 1 - c]$ as follows:¹¹

$$u(x, \theta^*, \theta_{-1}^*, \beta, \alpha, z) \equiv \begin{cases} 1 - \frac{\Phi\left(\sqrt{\beta + \alpha}\left(\frac{\beta}{\beta + \alpha}x + \frac{\alpha}{\beta + \alpha}z - \theta^*\right)\right)}{\Phi\left(\sqrt{\beta + \alpha}\left(\frac{\beta}{\beta + \alpha}x + \frac{\alpha}{\beta + \alpha}z - \theta_{-1}^*\right)\right)} - c & \text{if } \theta^* > \theta_{-1}^* \\ -c & \text{if } \theta^* \leq \theta_{-1}^* \end{cases}$$

$$X(\theta^*, \beta) \equiv \theta^* + \frac{1}{\sqrt{\beta}}\Phi^{-1}(\theta^*)$$

$$U(\theta^*, \theta_{-1}^*, \beta, \alpha, z) \equiv \begin{cases} \lim_{x \rightarrow -\infty} u(x, \theta^*, \theta_{-1}^*, \beta, \alpha, z) & \text{if } \theta^* = 0 \\ u(X(\theta^*, \beta), \theta^*, \theta_{-1}^*, \beta, \alpha, z) & \text{if } \theta^* \in (0, 1) \\ \lim_{x \rightarrow +\infty} u(x, \theta^*, \theta_{-1}^*, \beta, \alpha, z) & \text{if } \theta_{-1}^* = 1 \end{cases}$$

These functions have a simple interpretation: $u(x_t, \theta^*, \theta_{-1}^*, \beta_t, \alpha, z)$ is the net payoff from attacking in period t for an agent with statistic x_t when it is known that $\theta > \theta_{-1}^*$ and that regime change will occur if and only if $\theta \leq \theta^*$; $X(\theta^*, \beta_t)$ is the threshold x^* such that, if agents attack in period t if and only if $x_t \leq x^*$, then $A_t(\theta) \geq \theta$ if and only if $\theta \leq \theta^*$; $U(\theta^*, \theta_{-1}^*, \beta_t, \alpha, z)$ is the net payoff from attacking for the “marginal agent” with signal $x^* = X(\theta^*, \beta_t)$ when it is known that $\theta > \theta_{-1}^*$.

Next, solving (6) for x_t^* gives $x_t^* = X(\theta_t^*, \beta_t)$; substituting the latter into (7) gives

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

which represents the indifference condition for the marginal agent in period t , for $t \geq 2$. As for $t = 1$, since the regime has never been challenged in the past, the corresponding indifference condition is $U(\theta_1^*, -\infty, \beta_1, \alpha, z) = 0$; clearly, $U(\theta, -\infty, \beta, \alpha, z)$ coincides with the payoff of the marginal agent in the static benchmark.

We can thus characterize the set of monotone equilibria as follows.

¹¹With a slight abuse of notation, we let $\Phi(+\infty) = 1$, $\Phi(-\infty) = 0$, $\Phi^{-1}(1) = +\infty$ and $\Phi^{-1}(0) = -\infty$.

Proposition 2 $\{a_t(\cdot)\}_{t=1}^\infty$ 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 $U(\theta_1^*, -\infty, \beta_1, \alpha, z) = 0$ and $x_1^* = X(\theta_1^*, \beta_1)$.
 - (iii) for any $t \geq 2$, either $\theta_t^* = \theta_{t-1}^* > 0$ and $x_t^* = -\infty$, or $\theta_t^* > \theta_{t-1}^*$ is a solution to $U(\theta_t^*, \theta_{t-1}^*, \beta_t, \alpha, z) = 0$ and $x_t^* = X(\theta_t^*, \beta_t)$.
- A monotone equilibrium always exists.

Proposition 2 provides a simple algorithm for constructing the entire set of monotone equilibria: first, start with $t = 1$ and let θ_1^* be the unique solution to $U(\theta_1^*, -\infty, \beta_1, \alpha, z) = 0$; next, proceed to period $t = 2$; if $U(\theta_2^*, \theta_1^*, \beta_2, \alpha, z) = 0$ admits no solution, set $\theta_2^* = \theta_1^*$; if it admits a solution, either let θ_2^* be such a solution, or simply set $\theta_2^* = \theta_1^*$; finally, repeat for all $t \geq 3$ the same step as for $t = 2$. The set of sequences $\{\theta_t^*\}_{t=1}^\infty$ constructed this way, together with the associated sequences $\{x_t^*\}_{t=1}^\infty$, gives the set of monotone equilibria.

This recursive algorithm is based on the property that equilibrium learning takes the simple form of a truncation in the support of beliefs about θ : the knowledge that the regime has survived past attacks simply translates into the knowledge that θ is above a threshold θ_{t-1}^* . In Section 5 we examine how this property may, or may not, extend to richer environments. Note also that 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 $T < \infty$.

Existence of at least one monotone equilibrium follows immediately from the fact that the equation $U(\theta_1^*, -\infty, \beta_1, \alpha, z) = 0$ always admits a solution, and $\theta_t^* = \theta_1^*$ for all t is always an equilibrium. To understand whether there are other monotone equilibria, the next lemma investigates the properties of U and the existence of solutions to condition (8).

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

(ii) Let $\hat{\theta}_t$ be the unique solution to $U(\hat{\theta}_t, -\infty, \beta_t, \alpha, z) = 0$. 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 β_t sufficiently high.

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

(v) If θ_{t-1}^* is the highest solution to (8) for period $t - 1$, there exists $\underline{\beta} > \beta_{t-1}$ such that (8) admits no solution for any period $\tau \geq t$ such that $\beta_\tau < \underline{\beta}$.

To understand why $U(\theta^*, \theta_{-1}^*, \cdot)$ is non-monotonic in θ^* whenever $\theta_{-1}^* \in (0, 1)$, recall that the higher θ^* , the higher the threshold $x^* = X(\theta^*, \beta_t)$ such that, if agents attack in period t if and only if $x_t \leq x^*$, then $A_t(\theta) \geq \theta$ if and only if $\theta \leq \theta^*$. When $\theta^* < \theta_{-1}^*$, the threshold x^* is so low that the size of attack is smaller than θ for all $\theta > \theta_{-1}^*$; but then the marginal agent attaches zero probability




Figure 1: The payoff of the marginal agent.

to regime change, which explains why $U(\theta^*, \theta_{-1}^*, \cdot) = -c$ for $\theta^* < \theta_{-1}^*$. When instead $\theta^* > \theta_{-1}^*$, the threshold x^* is high enough that regime change occurs for a positive measure of $\theta > \theta_{-1}^*$; but then the marginal agent attaches positive probability to regime change, which explains why $U(\theta^*, \theta_{-1}^*, \cdot) > -c$ for $\theta^* > \theta_{-1}^*$. Finally, when $\theta^* \rightarrow 1$, $x^* \rightarrow \infty$ and hence the probability that the marginal agent attaches to the event that $\theta > 1$ converges to 1; but then the probability he attaches to regime change converges to zero, which explains why $U(\theta^*, \theta_{-1}^*, \cdot) \rightarrow -c$ as $\theta^* \rightarrow 1$.

We thus have that U is flat at $-c$ for $\theta^* < \theta_{-1}^*$, it then increases with θ^* and eventually decreases with θ^* and converges again to $-c$ as $\theta^* \rightarrow 1$. This is illustrated by the solid curve in Figure 1. Any intersection of this curve with the horizontal axis corresponds to a solution to (8).¹² The dashed line instead represents the payoff of the marginal agent in the static game in which agents can attack only in period t ; when β_t is sufficiently high, this is monotonic in θ^* . While the monotonicity of the payoff of the marginal agent in the static game ensures uniqueness, the non-monotonicity in the dynamic game leaves open the possibility for multiple equilibria.

Next, to understand why U decreases with z , note that an increase in the prior mean implies a first-order stochastic-dominance change in posterior beliefs about θ : the higher z , the lower the probability of regime change for any given monotone strategy, and hence the lower the net payoff from attacking for the marginal agent.

Similarly, since an increase in θ_{-1}^* also corresponds to an upward shift in posterior beliefs

¹²It can be shown that $U(\theta^*, \theta_{-1}^*, \cdot)$ is single-peaked in θ^* when $\theta_{-1}^* \geq 1/2$. Numerical simulations suggest that this is true even when $\theta_{-1}^* < 1/2$, although we have not been able to prove it. Single-peakedness of U implies that (8) admits at most two solutions (generically none or two). When there are two solutions, as in the case of the solid line in Figure 1, the lowest one corresponds to an unstable equilibrium, while the highest one corresponds to a stable equilibrium. None of these properties, however, are needed for our results. All that matters is that U is non-monotonic in θ^* when $\theta_{-1}^* \in (0, 1)$, with a finite number of stationary points.

about θ , U also decreases with θ_{-1}^* . This implies that, at any $t \geq 2$, the payoff of the marginal agent is always lower than $U(\theta^*, -\infty, \beta_t, \alpha, z)$, that is, than the payoff in the static game where the precision of private information is β_t . This in turn explains why the static-game threshold $\hat{\theta}_t$ (which corresponds to the intersection of the dashed line with the horizontal axis in Figure 1) is an upper bound for any solution to (8).

To understand (iii) and (iv), note that as $\beta_t \rightarrow \infty$ the impact on posterior beliefs of the knowledge that $\theta > \theta_{t-1}^*$ vanishes for any $x_t > \theta_{t-1}^*$. By implication, as $\beta_t \rightarrow \infty$, the difference between $U(\theta^*, \theta_{t-1}^*, \beta_t, \alpha, z)$ and $U(\theta^*, -\infty, \beta_t, \alpha, z)$ also vanishes for any $\theta^* > \theta_{t-1}^*$. Combined with the fact that $U(\theta^*, -\infty, \beta_t, \alpha, z) \rightarrow \theta_\infty - \theta^*$ as $\beta_t \rightarrow \infty$, this implies that, for β_t sufficiently high, (8) necessarily admits at least one solution if $\theta_{t-1}^* < \theta_\infty$, and no solution if $\theta_{t-1}^* > \theta_\infty$, where $\theta_\infty = \lim_{t \rightarrow \infty} \hat{\theta}_t$ is the limit of the equilibrium threshold in the static game for $\beta \rightarrow \infty$.

Finally, to understand (v), suppose that the largest possible attack (that is, the one corresponding to the highest solution to (8)) is played in one period and is unsuccessful. Then the upward shift in posterior beliefs induced by the observation that the status quo survived the attack is such that, if no new information arrives, no further 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.

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 prior mean, as anticipated in the Introduction.

When z is sufficiently low (“aggressive prior”), discounting the prior contributes to less aggressive behavior in the sense that $\hat{\theta}_t$ decreases with β_t and hence $\hat{\theta}_t < \hat{\theta}_1$ for all $t \geq 2$. It follows that an agent who is aware of the fact that the regime survived period one (i.e., that $\theta > \hat{\theta}_1$) would not be willing to attack in any period $t \geq 2$ if he expected all other agents to play as if no attack occurred prior to period t (i.e., as in the equilibrium of the static game where attacking is allowed only in period t). The anticipation that other agents will also take into account the fact that the regime survived past attacks then makes that agent even less willing to attack. Therefore, when z is low, the game has a unique equilibrium, with no attack occurring after the first period.

When, instead, z is sufficiently high (“lenient prior”), discounting the prior contributes to more aggressive behavior in the sense that $\hat{\theta}_t$ increases with β_t . This effect can offset the incentive not to attack induced by the knowledge that the regime survived past attacks, making new attacks eventually possible. Indeed, Lemma 2 implies that, when $\theta_1^* < \theta_\infty$ (which is the case for z high enough), a second attack necessarily becomes possible once β_t is large enough. Such an example is illustrated in Figure 2. The dashed line represents the payoff of the marginal agent in period 1. Its intersection with the horizontal axis defines $\theta_1^* < \theta_\infty$. The payoff of the marginal agent in period 2




Figure 2: Equilibria with multiple attacks.

is represented by the dotted line, and that in period 3 by the solid line. Clearly, β_2 is low enough that no attack is possible in period 2. In contrast, β_3 is high enough that a new attack is possible. Thus, there exist at least three equilibria in this example: one in which $\theta_t^* = \theta_1^*$ for all t , another in which $\theta_2^* = \theta_1^*$ and $\theta_t^* = \theta_3^'$ for all $t \geq 3$, and a third one in which $\theta_2^* = \theta_1^*$ and $\theta_t^* = \theta_3''$ for all $t \geq 3$, where $\theta_3^'$ and θ_3'' correspond to the two intersections of the solid line with the horizontal axis.

In the example of Figure 2, both $\theta_3^'$ and θ_3'' are lower than θ_∞ . By Lemma 2, then, a third attack also becomes possible at some future date. More generally, if z is sufficiently high, any solution to (8) is strictly less than θ_∞ in all periods, which ensures that a new attack eventually becomes possible after any unsuccessful one. Hence, for z sufficiently high, not only there are multiple equilibria, but any arbitrary number of attacks can be sustained in equilibrium.

Theorem 1 *There exist thresholds $\underline{z} \leq \bar{z} \leq \bar{\bar{z}}$ such that:*

(i) *If $z \leq \underline{z}$, there is a **unique** monotone equilibrium and is such that an attack occurs only in period one.*

(ii) *If $z \in (\underline{z}, \bar{z})$, there are at most finitely many monotone equilibria and there exists $\bar{t} < \infty$ such that in any of these equilibria, no attack occurs after period \bar{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. Recall that $\theta_1^* = \hat{\theta}_1$ and, for all $t \geq 2$, $\theta_t^* < \hat{\theta}_t$, where $\hat{\theta}_t = \hat{\theta}(\beta_t, \alpha, z)$ is the unique solution to $U(\hat{\theta}, -\infty, \beta_t, \alpha, z) = 0$. As proved in Lemma A1 in the Appendix, there exist thresholds $\underline{z} \leq \bar{z} \leq \bar{\bar{z}}$ (possibly functions of β_1 and α) 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 $z \geq (\leq) \bar{\bar{z}}$; 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 part (ii) of Lemma 2, (8) admits no solution at any $t \geq 2$. The unique monotone equilibrium is thus $\theta_t^* = \theta_1^*$ for all t .

(ii) Next, consider $z \in (\underline{z}, \bar{z})$, in which case $\hat{\theta}_1 = \theta_1^* > \theta_\infty$, but we can not rule out the possibility that there exists a period $t \geq 2$ such that $\hat{\theta}_t > \hat{\theta}_1$ and $U(\theta^*, \theta_1^*, \beta_t, \alpha, z) = 0$ admits a solution. Nevertheless, since $\theta_{t-1}^* \geq \theta_1^* > \theta_\infty$ for all t , by part (iii) of Lemma 2 and the fact that $\beta_t \rightarrow \infty$ as $t \rightarrow \infty$, there exists $\bar{t} < \infty$ such that (8) admits no solution for $t \geq \bar{t}$. Moreover, since (8) admits at most finitely many solutions for any $t < \bar{t}$, there are at most finitely many monotone equilibria, and in any such equilibrium no attack occurs after period \bar{t} .

(iii) Finally, consider $z > \bar{z}$, in which case $\theta_1^* < \theta_\infty$. Then, by part (iv) of Lemma 2, there exists a $t' < \infty$ such that $U(\theta^*, \theta_1^*, \beta_t, \alpha, z) = 0$ admits a solution for all $t \geq t'$. Hence, for any $t \geq t'$, there is a monotone equilibrium in which $\theta_\tau^* = \theta_1^*$ for $\tau < t$, θ_t^* solves $U(\theta_t^*, \theta_1^*, \beta_t, \alpha, z) = 0$, 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 hence by part (ii) of Lemma 2, $\theta_t^* < \theta_\infty$, for all t . But then by part (iv), a new attack eventually becomes possible after any unsuccessful one. It follows that, for any $t \geq 1$ and any $N \geq 1$, there exist increasing sequences $\{t_2, \dots, t_N\}$ and $\{\theta_2, \dots, \theta_N\}$, with $t_2 \geq t$, such that $U(\theta_2, \theta_1^*, \beta_{t_2}, \alpha, z) = 0$, $U(\theta_3, \theta_2, \beta_{t_3}, \alpha, z) = 0$, and so on. The following is then an equilibrium: $\theta_\tau^* = \theta_1^*$ for $\tau < t_2$, $\theta_\tau^* = \theta_j$ for $\tau \in \{t_j, \dots, t_{j+1} - 1\}$ and $j \in \{2, \dots, N - 1\}$, and $\theta_\tau^* = \theta_N$ for $\tau \geq t_N$. That is, for any $t \geq 1$ and any $N \geq 1$, there exists an equilibrium in which N attacks occur after period t . ■

The existence of infinitely many equilibria in the case $z > \bar{z}$ relies on the assumption that the game continues forever as long as the status quo is in place: if the game ended for exogenous reasons at a finite date, there would exist only finitely many equilibria. Nevertheless, as long as $z > \bar{z}$ and $\beta_t \rightarrow \infty$ as $t \rightarrow \infty$, then, for any M , there exists a finite T such that the game would have at least M equilibria if it ended at date T . Moreover, even when $T = 2$, the game has multiple equilibria if β_2 is sufficiently high and $z > \bar{z}$.

In the remainder of this section, we identify equilibrium properties that seem useful in understanding the dynamics of crises.

Corollary 1 *Suppose $\theta > \theta_\infty$ and $z > \bar{\bar{z}}$. The status quo survives in any monotone equilibrium. Nevertheless, there exists $\underline{t} < \infty$ such that, at any $t \geq \underline{t}$, an attack can occur, yet does not necessarily take place. Furthermore, any arbitrary number of attacks is possible.*

This seems to square well with the common view that economic fundamentals may help predict eventual outcomes (e.g., whether a currency is eventually devalued) but not when a crisis will occur or whether attacks will cease. On the contrary, this view is inconsistent with the common-knowledge version of the model, in which fundamentals fail to predict both the timing of attacks

and the eventual regime outcome whenever they are inside the critical region. It is also inconsistent with unique-equilibrium models like Morris and Shin (1999), in which both the timing of attacks and the ultimate fate of the regime are uniquely pinned down by the fundamentals.

Consider now how the dynamics of attacks depend on the dynamics of information.

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

Along with the property that for $\theta > \theta_\infty$ and $z > \bar{z}$ a new attack eventually becomes possible after any unsuccessful one, the above result implies that dynamics may take the form of cycles in which the economy alternates from phases of tranquillity to phases of distress, eventually resulting into a new attack, without any change in the underlying fundamentals. Once again, this would not have been possible in our framework if θ were common knowledge, or if there were a unique equilibrium.¹³

Also note that the set of equilibrium outcomes in any given period exhibits a discontinuity with respect to the precision of private information in that period: a transition from a phase of tranquillity, where nobody attacking is the unique equilibrium outcome, to a phase of distress, where the size of attack associated with any solution of (8) is bounded away from zero, can be triggered by a small change in β_t . As we will see in Section 5.3, a similar discontinuity emerges with respect to shocks that affect the strength of the status quo: a transition from one phase to another can then be triggered by an arbitrarily small change in fundamentals (that is, in payoffs).

Finally, note that these results raise some interesting possibilities for policy in the context of currency crises. On the one hand, since an increase in c shifts U downwards, a central bank might be able to prevent a transition from a phase of tranquillity to a phase of distress—and thus eliminate the risk of a speculative attack—by raising interest rates, or otherwise increasing the opportunity cost of attacking, up to the point that the (8) admits no solution. On the other hand, the level of policy intervention required to achieve this may increase over time as speculators become more informed about the underlying fundamentals, and may eventually become prohibitively expensive. Thus an interesting possibility is that certain defense policies succeed in postponing but not in escaping a crisis.¹⁴

¹³Broner (2005) considers a model that combines a common-knowledge coordination problem à la Obstfeld (1996) with a negative trend in fundamentals à la Krugman (1979). The first feature delivers multiplicity, while the second ensures that devaluation is eventually inevitable for exogenous reasons. His analysis thus shares with ours the property that it may be easier to predict the eventual outcome than the timing of attacks; but it does not share our predictions about the repeated succession of phases of tranquillity and phases of distress, nor our focus on changes in information, rather than changes in fundamentals, as the source of dynamics.

¹⁴Another possibility is that such defense measures themselves convey valuable information; Angeletos, Hellwig and Pavan (2006) examine such signaling effects in a static global game.

5 Extensions

In this section, we consider a few extensions of the benchmark model. The purpose of these extensions is to show how the analysis can accommodate additional elements which the benchmark model has deliberately abstracted from, but which can be relevant for applications. At the same time, these extensions show robustness to alternative information assumptions and further clarify the driving forces behind our results.

5.1 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, normally distributed with zero mean and precision $\eta_t^z > 0$, serially uncorrelated, and independent of θ and the noise in the agents' private information. These signals may represent, for example, the information generated by news in the media, publication of government statistics, or announcements by policy makers. We also allow for the possibility that the game ends for exogenous reasons at a finite date and denote the horizon of the game with T , where $T \in \{2, 3, \dots\}$ or $T = \infty$.

The common posterior about θ conditional on $\tilde{z}^t \equiv \{\tilde{z}_\tau\}_{\tau=1}^t$ is Normal with mean z_t and precision α_t , where

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

with $(z_0, \alpha_0) = (z, \alpha)$. However, since equilibrium play in past periods now depends on the realizations of past public signals, z_t is not a sufficient statistic *conditional* on the event that the regime is still in place. We thus allow agents to condition their actions on the entire sequence \tilde{z}^t , or equivalently on $z^t \equiv \{z_\tau\}_{\tau=1}^t$. Apart from this modification, the set of monotone equilibria can be constructed following the same algorithm as in the benchmark model.

Proposition 3 *In the game with public signals, $\{a_t(\cdot)\}_{t=0}^T$ is a monotone equilibrium if and only if there exists a sequence of functions $\{x_t^*, \theta_t^*\}_{t=1}^T$, with $x_t^* : \mathbb{R}^t \rightarrow \overline{\mathbb{R}}$ and $\theta_t^* : \mathbb{R}^t \rightarrow (0, 1)$, such that:*

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

$$U(\theta_t^*, \theta_{t-1}^*(z^{t-1}), \beta_t, \alpha_t, z_t) = 0 \tag{9}$$

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

As in the benchmark model without public signals, there always exist equilibria in which attacks cease after any arbitrary period. However, since for any $\theta_{t-1}^*(z^{t-1}), \beta_t, \alpha_t$, (9) admits a solution if and only if $z_t \leq \bar{z}_t$, where $\bar{z}_t = \bar{z}(\theta_{t-1}^*, \beta_t, \alpha_t)$ is always finite, there also exist equilibria in which an attack occurs in period t for sufficiently low realizations of z_t , which proves the following.

Theorem 2 *In the game with public signals, there always exist multiple equilibria.*

This result extends and reinforces Theorem 1: multiplicity now emerges no matter the mean z of the prior, the precisions $\{\beta_t, \alpha_t\}_{t=1}^T$ of private and public information, and the horizon T of the game. This stronger version of multiplicity relies on the combination of two properties: that sufficiently low realizations of z_t make an attack possible in every period; and that the lower dominance region is eliminated in all periods $t \geq 2$ so that no attack also remains possible in every period after the first one.

Consider now how the introduction of public news affects the ability of an “econometrician” to predict the regime outcome and/or the occurrence of an attack in any given period. For any t , any $\theta \in (0, 1)$, and any $\theta_{t-1}^*(z^{t-1}) < \theta$, condition (9) admits a solution higher than θ if and only if z_t is low enough, implying that, conditional on θ , the probability that the status quo is abandoned in any given period is strictly between 0 and 1. It follows that an econometrician who can observe θ but can not observe z^t , necessarily faces uncertainty about the event of regime change. On the other hand, if he also knows z^t , he may be able to predict the regime outcome in a given period for some combinations of θ and z^t , without, however, being able to predict whether an attack will occur or not. For example, take any $t \geq 2$, let $\underline{\theta}_1(z_1)$ and $\bar{\theta}_t(z^t)$ be the lowest and the highest solutions to $U(\theta^*, -\infty, \beta_1, \alpha_1, z_1) = 0$ and $U(\theta^*, \underline{\theta}_1(z_1), \beta_t, \alpha_t, z_t) = 0$, respectively, and assume that $\theta > \bar{\theta}_t(z^t) > \underline{\theta}_1(z_1)$. There is no equilibrium in which the status is abandoned in period t , but there exist both an equilibrium in which an attack occurs and one in which no attack takes place in that period.¹⁵ Therefore, the combination of fundamentals *and* public information may help predict regime outcomes but not the occurrence of attacks, as in the benchmark model.

Also note that the threshold \bar{z}_t , below which (9) admits a solution, decreases with θ_{t-1}^* . Hence, an unsuccessful attack, other things equal, causes a discrete increase in the probability that the game enters a phase during which no attack is possible. In this sense, the prediction of the benchmark model that equilibrium dynamics are characterized by the alternation of phases of tranquility and phases of distress survives the introduction of public news; the novelty is that the transition from one phase to another is now stochastic, as it depends on the realization of z^t .

5.2 Signals about past attacks

In the analysis so far, agents learn from the *outcome* of past attacks but receive no information about the *size* of these attacks. For many applications, however, it seems natural to allow agents to observe noisy private and/or public signals about the size of past attacks.

In the online Supplementary Material we show how this can be done without any sacrifice in

¹⁵Since in any equilibrium necessarily $\theta_{t-1}^*(z^{t-1}) \geq \underline{\theta}_1(z_1)$, from part (i) in Lemma 1, $\bar{\theta}_t(z^t)$ is an upper bound for $\theta_t^*(z^t)$. This implies that, if $\theta > \bar{\theta}_t(z^t) > \underline{\theta}_1(z_1)$, there is no equilibrium in which the regime is abandon in period t . On the other hand, since $U(\theta^*, \underline{\theta}_1(z_1), \beta_t, \alpha_t, z_t) = 0$ admits a solution $\theta^* = \bar{\theta}_t(z^t) > \underline{\theta}_1(z_1)$, there exists an equilibrium in which the second attack occurs exactly in period t .

tractability. The key is to maintain the Normality of the information structure. The algorithm for monotone equilibria then remains the same as in Proposition 3, except for the fact that the sequence $\{\beta_t, \alpha_t\}_{t=1}^{\infty}$ is now part of the equilibrium: the precisions of private and public information in any period $t \geq 2$ depend on whether an attack occurred in the previous period.

As for the structure of equilibrium dynamics, the novelty is that the upward shift in posterior beliefs caused by an unsuccessful attack may now be diluted by the information about θ conveyed by the size of the attack. This in turn may lead to situations where new attacks become possible immediately after unsuccessful ones, even without any exogenous arrival of information. As a result, equilibrium dynamics may now feature snow-balling effects reminiscent of the ones highlighted in herding models of crises (e.g., Chari and Kehoe, 2003).

5.3 Observable shocks: changes in fundamentals as a source of dynamics

In this section we introduce shocks to the sustainability of the regime. In particular, we modify the benchmark model as follows. The regime is abandoned in period t if and only if $A_t \geq h(\theta, \delta\omega_t)$. The variable θ continues to represent the “strength of the status quo”, while ω_t is an exogenous disturbance, independent of θ and i.i.d. over time, with absolutely continuous c.d.f. F and support \mathbb{R} . The scalar $\delta > 0$ parameterizes the volatility of these disturbances. Finally, for simplicity, the function h is assumed to be linear, with $h(\theta, \delta\omega_t) = \theta + \delta\omega_t$. We denote this game with $\Gamma(\delta)$, nesting the baseline model as $\delta = 0$.

We assume here that ω_t is publicly observable, which may be relevant for some applications of interest. In the case of currency attacks, for example, θ may represent the “type” of the central banker, whereas ω_t may capture the role of interest rates, financial prices, and other macroeconomic variables that are readily observable by economic agents and that may affect the willingness or ability of the central banker to defend the peg.¹⁶

As we show below, observable shocks are easy to incorporate in the analysis, because they affect equilibrium dynamics without introducing any noise in the learning about θ . The exercise here is thus useful not only for applications, but also for separating the role of shocks as drivers of equilibrium dynamics from their role as additional sources of noise in learning.

Equilibrium characterization. Since the shocks are observable, the strategy of the agents in period t is contingent on both x^t and $\omega^t \equiv \{\omega_1, \dots, \omega_t\}$. Accordingly, the regime outcome in period t is contingent on both θ and ω^t . The set of monotone equilibria can thus be characterized by a sequence $\{x_t^*(\omega^t), \theta_t^*(\omega^t)\}_{t=1}^{\infty}$ such that an agent attacks in period t if and only if $x_t < x_t^*(\omega^t)$ and the status quo is in place in next period if and only if $\theta > \theta_t^*(\omega^t)$.

Note that sufficiently negative shocks re-introduce the lower dominance region: whenever $\theta_{t-1}^*(\omega^{t-1}) + \delta\omega_t < 0$, it is dominant for agents with sufficiently low x_t to attack. Otherwise,

¹⁶We could have allowed for shocks in the opportunity cost of attacking by letting c depend on ω_t ; we omitted this possibility only for expositional simplicity.

the structure of equilibria, and the algorithm for constructing them, is similar to the one in the benchmark model.

Proposition 4 *In the game with shocks, $\{a_t(\cdot)\}_{t=1}^{\infty}$ is a monotone equilibrium if and only if there exists a sequence $\{x_t^*(\omega^t), \theta_t^*(\omega^t)\}_{t=1}^{\infty}$ such that:*

- (i) *for all t , $a_t(\cdot) = 1$ if $x_t < x_t^*(\omega^t)$ and $a_t(\cdot) = 0$ if $x_t > x_t^*(\omega^t)$.*
- (ii) *for $t = 1$, $U(\theta_1^*(\omega_1) + \delta\omega_1, -\infty; \beta_1, \alpha, z + \delta\omega_1) = 0$ and $x_1^*(\omega_1) + \delta\omega_1 = X(\theta_1^*(\omega_1) + \delta\omega_1; \beta_1)$*
- (iii) *for any $t \geq 2$, either $\theta_t^*(\omega^t) > \theta_{t-1}^*(\omega^{t-1})$,*

$$U(\theta_t^*(\omega^t) + \delta\omega_t, \theta_{t-1}^*(\omega^{t-1}) + \delta\omega_t; \beta_t, \alpha, z + \delta\omega_t) = 0 \quad (10)$$

and $x_t^(\omega^t) + \delta\omega_t = X(\theta_t^*(\omega^t) + \delta\omega_t; \beta_t)$, or $\theta_t^*(\omega^t) = \theta_{t-1}^*(\omega^{t-1}) \geq -\delta\omega_t$ and $x_t^*(\omega^t) = -\infty$.*

This result can be understood as follows. While the critical size of attack that is necessary for regime change is constant in the benchmark model, here it varies over time as a consequence of shocks. However, since shocks are observable, the structure of beliefs remains the same apart from a “change of variables” in the following sense. Let $h_t \equiv \theta + \delta\omega_t$ be the critical size for period t , $x'_t \equiv x_t + \delta\omega_t$, and $z'_t \equiv z + \delta\omega_t$. The distribution of h_t conditional on x^t is Normal with mean $\frac{\beta_t}{\beta_t + \alpha}x'_t + \frac{\alpha}{\beta_t + \alpha}z'_t$ and precision $\beta_t + \alpha$, while the knowledge that $\theta > \theta_{t-1}^*(\omega^{t-1})$ is equivalent to the knowledge that $h_t > \theta_{t-1}^*(\omega^{t-1}) + \delta\omega_t$. It follows that the net payoff from attacking for the marginal agent in period t is $U(\theta_t^*(\omega^t) + \delta\omega_t, \theta_{t-1}^*(\omega^{t-1}) + \delta\omega_t; \beta_t, \alpha, z + \delta\omega_t)$. The result then follows from the same arguments as in the proof of Proposition 2.

Multiplicity and dynamics. Interesting new effects can emerge because of the interaction of information and shocks. The equilibrium dynamics again feature phases of tranquility, where an attack is impossible, and phases of distress, where an attack is possible but does not necessarily take place. However, shocks provide a second channel through which a transition from one phase to another can occur. In particular, a transition from distress to tranquility may now be triggered either by an unsuccessful attack, or by an improvement in fundamentals (a positive ω_t); and a transition from tranquility to distress can be caused either by the arrival of new private information, or by a deterioration in fundamentals. What is more, the economy can now enter a phase where an attack is inevitable—a scenario that was impossible in the benchmark model, but becomes possible here because sufficiently bad shocks re-introduce a lower dominance region.¹⁷

If the benchmark game $\Gamma(0)$ admits multiple equilibria, then the game with shocks $\Gamma(\delta)$ also admits multiple equilibria, no matter δ ; to see this, it suffices to consider realizations of ω_t close enough to zero. Moreover, since the impact of shocks on the conditions that characterize the equilibrium dynamics clearly vanishes as $\delta \rightarrow 0$, the following equilibrium convergence result holds: for any $T > 0$, any $\varepsilon > 0$, and any equilibrium $\{x_t^*, \theta_t^*\}$ of the benchmark game $\Gamma(0)$, there exists

¹⁷Another difference is that not attacking becomes dominant for sufficiently high ω_t no matter x_t , whereas in the benchmark model not attacking is at most iteratively dominant for sufficiently low x_t . This, however, makes little difference in terms of observable dynamics.

a $\hat{\delta} > 0$ such that for all $\delta < \hat{\delta}$ the game $\Gamma(\delta)$ admits an equilibrium $\{x_t^\delta(\omega^t), \theta_t^\delta(\omega^t)\}$ such that the unconditional probability that $|\theta_t^\delta(\omega^t) - \theta_t^*| < \varepsilon$ for all $t \leq T$ is higher than $1 - \varepsilon$.¹⁸

None of these results, however, should be surprising given that shocks do not interfere with the learning process. Indeed, what is important for the results of this section is not the absence of uncertainty about these shocks, but the fact that the shocks do not introduce noise in learning.

To see this, consider the case that ω_t is unobservable in period t but becomes commonly known at the beginning of period $t + 1$. Then agents face additional uncertainty about the regime outcome—indeed, the regime outcome would remain uncertain even if agents had known θ —but the knowledge that the regime has survived past attacks still translates into common certainty that θ is above a certain threshold. That is, the form of learning remains as sharp as in the benchmark model. Not surprisingly then, the equilibrium convergence result described above extends to this case as well.¹⁹

We conclude that with respect to robustness the question of interest is whether equilibrium convergence obtains in situations where shocks also introduce noise in learning. To examine this question, we next turn to the case that ω_t remains unobservable in all periods.

5.4 Unobservable shocks: noisy learning

We now modify the game with shocks examined in the previous section by letting ω_t be unobservable. The unobservability of shocks “noises up” the learning process and ensures that the updating of beliefs caused by the knowledge that the regime is still in place never takes the form of a truncation—agents’ posteriors have full support in \mathbb{R} in all periods. The case of unobservable shocks that we examine in this section is therefore most significant from a theoretical perspective.

Below we first explain how unobservable shocks affect the algorithm for the construction of equilibria. We then show how the equilibria in the benchmark model can be approximated arbitrarily well by equilibria of the perturbed game as $\delta \rightarrow 0$. It follows that the key qualitative properties of the equilibrium dynamics identified in the benchmark model—the multiplicity and the succession of phases of tranquility and distress—extend to the case with unobservable shocks provided that the volatility of these shocks is small enough and that we reinterpret a phase of tranquility as one where at most an (arbitrarily) small attack is possible.

Equilibrium characterization. Because ω_t affects the regime outcome and is unobserved, the ability to characterize the set of monotone equilibria in terms of a sequence of truncation points for θ is lost. Nevertheless, as long as private information can be summarized by a sufficient statistic $x_t \in \mathbb{R}$, we can still characterize monotone equilibria as sequences of thresholds $\{x_t^*\}_{t=1}^\infty$ such that an agent attacks in period t if and only if $x_t \leq x_t^*$, where $x_t^* \in \overline{\mathbb{R}}$.

To see this, consider an arbitrary monotone strategy, indexed by $\{\bar{x}_t\}_{t=1}^\infty$, such that an agent attacks in period t if and only if $x_t < \bar{x}_t$. Given this strategy, the size of the attack in period t

¹⁸This result is proved in the online Supplementary Material.

¹⁹This case is also examined in the online Supplementary Material.

is $A_t(\theta) = \Phi(\sqrt{\beta_t}(\bar{x}_t - \theta))$, and hence the status quo is abandoned in that period if and only if $\omega_t \leq \bar{\omega}_t^\delta(\theta; \bar{x}_t)$, where

$$\bar{\omega}_t^\delta(\theta; \bar{x}_t) \equiv \frac{1}{\delta} \left[\Phi\left(\sqrt{\beta_t}(\bar{x}_t - \theta)\right) - \theta \right].$$

It follows that the probability of regime change in period t conditional on θ is

$$p_t^\delta(\theta; \bar{x}_t) = F(\bar{\omega}_t^\delta(\theta; \bar{x}_t)).$$

Next, consider the learning induced when the strategy $\{\bar{x}_t\}_{t=1}^\infty$ is played. For any $t \geq 2$, let $\psi_t^\delta(\theta; \bar{x}^{t-1})$ denote the density of the *common* posterior about θ , when in previous periods agents followed monotone strategies with thresholds $\bar{x}^{t-1} = \{\bar{x}_1, \dots, \bar{x}_{t-1}\}$. By Bayes' rule,

$$\psi_t^\delta(\theta; \bar{x}^{t-1}) = \frac{[1 - p_{t-1}^\delta(\theta; \bar{x}_{t-1})] \psi_{t-1}^\delta(\theta; \bar{x}^{t-2})}{\int_{-\infty}^{+\infty} [1 - p_{t-1}^\delta(\theta'; \bar{x}_{t-1})] \psi_{t-1}^\delta(\theta'; \bar{x}^{t-2}) d\theta'} = \frac{\prod_{s=1}^{t-1} [1 - p_s^\delta(\theta; \bar{x}_s)] \psi_1^\delta(\theta)}{\int_{-\infty}^{+\infty} \prod_{s=1}^{t-1} [1 - p_s^\delta(\theta'; \bar{x}_s)] \psi_1^\delta(\theta') d\theta'}$$

where $\psi_1^\delta(\theta) = \sqrt{\alpha}\phi(\sqrt{\alpha}(\theta - z))$ is the density of the initial prior. When $\delta = 0$, the above reduces to a truncated Normal distribution, with truncation point $\bar{\theta}_{t-1}(\bar{x}^{t-1}) \equiv \min\{\theta : \theta \geq \Phi(\sqrt{\beta_\tau}(\bar{x}_\tau - \theta)) \forall \tau \leq t-1\}$. When instead $\delta > 0$, learning is “smoother” in the sense that $\psi_t^\delta(\theta; \bar{x}^{t-1})$ is strictly positive and continuous over the entire real line.

Finally, consider payoffs. For any $t \geq 1$, $x \in \mathbb{R}$, and $\bar{x}^t \in \bar{\mathbb{R}}^t$, let $v_t^\delta(x; \bar{x}^t)$ 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 or equal to \bar{x}_τ . This is given by

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

where $\psi_t^\delta(\theta|x; \bar{x}^{t-1})$ denotes the density of the *private* posterior in period t . (The latter are computed applying Bayes' rule to the common posteriors described above.) Note that $v_t^\delta(x; \bar{x}^t)$ 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 $\bar{x}^t \in \bar{\mathbb{R}}^t$, let

$$V_t^\delta(\bar{x}^t) \equiv \begin{cases} \lim_{x \rightarrow +\infty} v_t^\delta(x; \bar{x}^t) & \text{if } \bar{x}_t = +\infty \\ v_t^\delta(\bar{x}_t; \bar{x}^t) & \text{if } \bar{x}_t \in \mathbb{R} \\ \lim_{x \rightarrow -\infty} v_t^\delta(x; \bar{x}^t) & \text{if } \bar{x}_t = -\infty \end{cases} \quad (11)$$

V_t^δ 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 threshold \bar{x}_t .

In Lemma A2 in the Appendix, we prove that, for any $\delta > 0$, $V_t^\delta(\bar{x}^t)$ is continuous in \bar{x}^t for any $\bar{x}^t \in \bar{\mathbb{R}}^{t-1} \times \mathbb{R}$, which we use to establish the existence, and complete the characterization, of monotone equilibria.

Proposition 5 For any $\delta > 0$, $\{a_t(\cdot)\}_{t=1}^{\infty}$ is a monotone equilibrium for $\Gamma(\delta)$ if and only if there exists a sequence $\{x_t^*\}_{t=1}^{\infty}$ 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 $t = 1$, $x_1^* \in \mathbb{R}$ and $V_1^\delta(x_1^*) = 0$;
- (iii) for any $t \geq 2$, either $x_t^* = -\infty$ and $V_t^\delta(x^{*t}) \leq 0$, or $x_t^* \in \mathbb{R}$ and $V_t^\delta(x^{*t}) = 0$.

A monotone equilibrium exists for any $\delta > 0$.

The equilibrium algorithm provided above clearly applies also to $\delta = 0$ and is similar to the one in Proposition 2: start with $t = 1$ and let x_1^* be the unique solution to $V_1^\delta(x_1^*) = 0$; proceed to period $t = 2$ and either let $x_2^* = -\infty$ if $V_2^\delta(x_1^*, -\infty) \leq 0$, or let x_2^* be the solution to $V_2^\delta(x_1^*, x_2^*) = 0$; repeat for any $t \geq 3$. The difference is that here at each step t we need to keep track of the entire sequence of past thresholds x^{*t-1} , while in the algorithm of Proposition 2 the impact of x^{*t-1} on period- t beliefs was summarized by θ_{t-1}^* .

Multiplicity and dynamics. As $\delta \rightarrow 0$, the dependence of the regime outcome on the shock ω_t vanishes. By implication, the posteriors in any period $t \geq 2$ converge pointwise to truncated Normals as in the benchmark model. The pointwise convergence of p_t^δ and ψ_t^δ in turn implies pointwise convergence of the payoff of the marginal agent: for $t = 1$ and any \bar{x}_1 , $V_1^\delta(\bar{x}_1) \rightarrow V_1^0(\bar{x}_1) \equiv U(\bar{\theta}_1(\bar{x}_1), -\infty, \beta_1, \alpha, z)$; similarly, for any $t \geq 2$, \bar{x}^{t-1} and $\bar{x}_t > -\infty$,

$$V_t^\delta(\bar{x}^t) \rightarrow V_t^0(\bar{x}^t) \equiv U(\bar{\theta}_t(\bar{x}_t), \bar{\theta}_{t-1}(\bar{x}^{t-1}), \beta_t, \alpha, z).$$

Pointwise convergence of payoffs, however, can fail for $t \geq 2$ at $\bar{x}_t = -\infty$. To see why, note that, in the presence of shocks, an agent with sufficiently low x_t may attach probability higher than c to regime change in period $t \geq 2$ even if he expects no other agent to attack in that period. When this is the case, a positive measure of agents may attack in every period in the perturbed game, unlike the benchmark model.

Nevertheless, the pointwise convergence of $V_t^\delta(\bar{x}^t)$ for any $\bar{x}_t > -\infty$ ensures that this dominance region vanishes as $\delta \rightarrow 0$. It also ensures that whenever $V_t^0(\bar{x}^t)$ has an intersection with the horizontal axis, $V_t^\delta(\bar{x}^t)$ also has a nearby intersection for $\delta > 0$ small enough. These properties together imply that any equilibrium in the benchmark game can be approximated arbitrarily well by an equilibrium in the perturbed game, except for knife-edge cases where V_t^0 (or equivalently U) is tangent to the horizontal axis instead of intersecting it.

Theorem 3 For any $\varepsilon > 0$ and any $T < \infty$, there exists $\delta(\varepsilon, T) > 0$ such that the following is true for all $\delta < \delta(\varepsilon, T)$:

For any equilibrium $\{x_t^*\}_{t=1}^{\infty}$ of $\Gamma(0)$ such that $x_t^* \notin \arg \max_x V_t^0(x^{*t-1}, x)$ for all $t \in \{2, \dots, T\}$, there exists an equilibrium $\{x_t^\delta\}_{t=1}^{\infty}$ of $\Gamma(\delta)$ such that, for all $t \leq T$, either $|x_t^* - x_t^\delta| < \varepsilon$, or $x_t^* = -\infty$ and $x_t^\delta < -1/\varepsilon$.

The result is illustrated in Figures 3 and 4 for an example where $T = 2$ and where $\Gamma(0)$ admits multiple equilibria. The solid line in Figure 3 represents the p.d.f. of the common posterior in




Figure 3: Common posteriors with and without shocks.




Figure 4: Payoff of marginal agent with and without shocks.

period 2 generated by equilibrium play in period 1 in the game without shocks ($\delta = 0$). This is simply the initial prior truncated at $\theta_1^* = \bar{\theta}_1(x_1^*)$, where x_1^* is the unique solution to $V_1^0(x_1^*) = 0$ (or equivalently where θ_1^* is the unique solution to $U(\theta_1^*, -\infty, \beta_1, \alpha, z) = 0$). The other two lines represent the equilibrium common posteriors $\psi_2^\delta(\theta; x_1^\delta)$ for the game with shocks ($\delta > 0$), where x_1^δ is the unique solution to $V_1^\delta(x_1^\delta) = 0$; the dotted line corresponds to a relatively high δ and the dashed one to a low δ . Since the support of ω_t is the entire real line, the probability of regime change is less than 1 for any θ and therefore ψ_2^δ assigns positive probability to all θ . However, as δ becomes smaller, x_1^δ converges to x_1^* and the probability of regime change converges to 1 for $\theta < \theta_1^*$ and to 0 for $\theta > \theta_1^*$. By implication, the smooth common posterior of the perturbed game in period 2 converges to the truncated one of the benchmark model.

In Figure 4, the solid line represents the payoff $V_2^0(x_1^*, x_2)$ of the marginal agent in period 2 for $\delta = 0$, whereas the other two lines represent $V_2^\delta(x_1^\delta, x_2)$ for $\delta > 0$.²⁰ Note that, for x_2 small enough, V_2^0 is negative but V_2^δ is positive, which implies that nobody attacking in period 2 is part of an equilibrium in the benchmark model but not in the game with shocks.²¹ Moreover, when δ is high (dotted line), V_2^δ is monotonic in x_2 and therefore has a single intersection with the horizontal line, in which case the equilibrium would be unique if the game ended in period 2. When, instead, δ is sufficiently small (dashed line), V_2^δ is non-monotonic and has three intersections, which correspond to three different equilibria for the two-period game with shocks. Finally, the middle and the highest intersections approximate the two intersections of the solid line, while the lowest intersection is arbitrarily small, thus approximating $x_2^* = -\infty$. Along with the fact that x_1^δ converges to x_1^* , this implies that any equilibrium of the two-period game without shocks can be approximated by an equilibrium of the perturbed game.

At the same time, it is important to recognize that uniqueness is ensured in the alternative limit as $\beta_2 \rightarrow \infty$ for given $\delta > 0$. To see this, note that, for any given $\delta > 0$, the common posterior in period 2 has a strictly positive density over the entire real line, and hence over a connected set of θ that includes both dominance regions ($\theta < 0$ and $\theta > 1$). This in turn ensures that standard global-game uniqueness results apply: uniqueness necessarily obtains in the limit as the noise in private information vanishes (see Proposition 2.2 in Morris and Shin, 2003).

However, away from this limit, the impact of private information here is quite different from that in the static Gaussian benchmark (Section 2.1). There, private information always contributes toward uniqueness.²² Here, instead, it can have a non-monotonic effect on the determinacy of equilibria: for δ small enough, in period 2 uniqueness obtains for β_2 either close to β_1 or close to ∞ , while multiplicity obtains for intermediate β_2 .

This non-monotonic effect of private information in our game highlights the interaction of private information with equilibrium learning. On the one hand, more precise private information increases strategic uncertainty as in the static game; on the other, it dilutes the upward shift in posterior beliefs caused by the knowledge that the regime survived past attacks. Whereas the first effect contributes to uniqueness, the second can contribute to multiplicity in a similar fashion as in the benchmark game without shocks.

In conclusion, what sustains multiplicity in the dynamic game is the property that the knowledge that the regime survived attacks in the past provides relevant common information about

²⁰In order to illustrate V_2^δ over its entire domain, the figure depicts $V_2^\delta(x_1^\delta, x_2)$ against $f(x_2)$ rather than x_2 , where f is a strictly increasing function that maps \mathbb{R} onto a bounded interval (e.g., $f = \Phi$).

²¹In the example discussed here, an agent finds it dominant to attack in period 2 for sufficiently low x_2 . However, this need not be the case if ω_t had a bounded support and δ were small enough.

²²This is true in two senses. First, a higher β makes it more likely that the economy satisfies $\beta \geq \alpha^2/(2\pi)$ in which case the equilibrium is unique. Second, whenever $\beta < \alpha^2/(2\pi)$ the range of z for which (1) and (2) admit multiple solutions shrinks with β , and the distance between the largest and the smallest solutions for any given z also diminishes with β .

the strength of the status quo in the present. That this knowledge resulted in posterior beliefs that assign zero measure to sufficiently low θ in the benchmark model is not essential. What is important is that the effect of this information on posterior beliefs is not diluted too much either by a significant change in fundamentals (sufficiently high δ) or by a significant increase in strategic uncertainty (sufficiently high β).

5.5 Private information about shocks: long- versus short-lived agents

In the environments examined above, agents had no *private* information about the innovations in fundamentals. Relaxing this assumption may compromise tractability by removing the ability to summarize the history of private information with a one-dimensional sufficient statistic.

To see this, consider the following variation of the game with shocks. Let h_t denote again the critical size of attack that triggers regime change in period t and assume that $\{h_t\}$ are jointly Normal with non-zero correlation across time. To simplify, think of h_t following a Gaussian random walk: $h_1 = \theta \sim N(z, 1/\alpha)$ and $h_t = h_{t-1} + \delta\omega_t$ for $t \geq 2$, with $\omega_t \sim N(0, 1)$ being i.i.d. across time and independent of θ .²³ Next, let the private signals agents receive in period t be $\tilde{x}_t = h_t + \xi_t$, where ξ_t is i.i.d. across agents and time, and independent of h_s for any s .

Consider first $t = 1$. Equilibrium play is the same as in the static benchmark: there exist thresholds x_1^* and h_1^* such that an agent attacks if and only if $\tilde{x}_1 \leq x_1^*$ and the status quo is abandoned if and only if $h_1 \leq h_1^*$. Consider next $t = 2$. The posterior beliefs about h_2 given the private signals \tilde{x}_1 and \tilde{x}_2 alone are Normal with mean $x_2 \equiv \lambda_0 + \lambda_1\tilde{x}_1 + \lambda_2\tilde{x}_2$ and variance σ_2^2 , for some coefficients $(\lambda_0, \lambda_1, \lambda_2, \sigma_2)$. This may suggest that x_2 can be used as a sufficient statistic for $(\tilde{x}_1, \tilde{x}_2)$ with respect to h_2 . However, the posterior beliefs about h_2 conditional also on the event that $h_1 > h_1^*$ are not invariant in $(\tilde{x}_1, \tilde{x}_2)$ for given x_2 ; the problem is that x_2 is *not* a sufficient statistic for $(\tilde{x}_1, \tilde{x}_2)$ with respect to h_1 . Thus, private information cannot be summarized in x_2 and equilibrium play in period 2 is characterized by a function $Y_2 : \mathbb{R}^2 \rightarrow \mathbb{R}$ such that an agent attacks if and only if $Y_2(\tilde{x}_1, \tilde{x}_2) \leq 0$ (and a corresponding function $Q_2 : \mathbb{R}^2 \rightarrow \mathbb{R}$ such that regime change occurs if and only if $Q_2(h_1, h_2) \leq 0$). Similarly, equilibrium play in any period $t \geq 2$ is characterized by a function $Y_t : \mathbb{R}^t \rightarrow \mathbb{R}$ such that $a_t(\tilde{x}^t) = 1$ if and only if $Y_t(\tilde{x}^t) \leq 0$.

Contrast this with the formalization in the previous section. There, in each period, we had to solve an equation where the unknown was a real variable $\bar{x}_t \in \mathbb{R}$. Here, instead, we need to solve each period a *functional* equation where the unknown is a function Y_t with domain \mathbb{R}^t —a function whose dimensionality explodes with t . Clearly, this is far less tractable, if at all feasible.

Moreover, it is not clear if this alternative formalization brings any substantial gain from a theoretical perspective. Both formalizations ensure that the critical size of attack h_t (and hence the payoff structure) may change over time, that agents have asymmetric information about h_t

²³Note that this is the same as $h_t = \theta + \delta\tilde{\omega}_t$, where $\tilde{\omega}_t \equiv \omega_1 + \dots + \omega_t$; that is, the same as in the game with shocks in Sections 5.3 and 5.4, but with the shocks correlated across time. Such correlation was not allowed in Sections 5.3 and 5.4, but this was only for simplicity.

in each period, and that the common posterior about h_t is continuous over a connected set that includes both $h_t < 0$ and $h_t > 1$ (and hence that dominance regions are possible for both actions). In these respects, they both seem appropriate extensions of global games to a dynamic setting.²⁴

Nevertheless, this second formalization may be more appropriate for certain applications. One way then to restore tractability is to assume that agents are short-lived. In particular, consider the game described above, in which h_t follows a Gaussian random walk, with the following modification. As long as the status quo is in place, a new cohort of agents replaces the old one in each period. Each cohort is of measure 1 and lives exactly one period. Agents who are born in period t receive private signals $x_t = h_t + \xi_t$, where ξ_t is Normal noise with precision β_t , i.i.d. across agents and independent of h_s for any $s \leq t$.

Given that h_t is correlated across time, the knowledge that the regime survived past attacks is informative about the strength of the regime in the present, as in the case with long-lived agents. But unlike that case, agents playing in period t have no private information other than x_t . Together with the fact that h_t alone pins down the cross-sectional distribution of x_t , this property ensures that monotone equilibria can again be characterized by sequences $\{x_t^*, h_t^*\}_{t=1}^\infty$ such that in period t an agent attacks if and only if $x_t < x_t^*$ and the status quo is abandoned if and only if $h_t < h_t^*$.

The characterization of equilibria then parallels that in the previous section. To see this, let $\Psi_t^\delta(h_t; \bar{x}^{t-1})$ denote the c.d.f. of the common posterior in period t about h_t , when agents in earlier cohorts attacked in periods $\tau \leq t-1$ if and only if $x_\tau < \bar{x}_\tau$. When earlier cohorts followed such strategies, the status quo survived period τ if and only if $h_\tau > \bar{\theta}_\tau(\bar{x}_\tau)$, where $\bar{\theta}_\tau(\bar{x}_\tau)$ is the solution to $\Phi(\sqrt{\beta_\tau}(\bar{x}_\tau - h_\tau)) = h_\tau$. Therefore, for any $t \geq 2$, $\Psi_t^\delta(h_t; \bar{x}^{t-1})$ is recursively defined by

$$\Psi_t^\delta(h_t; \bar{x}^{t-1}) = \frac{\int_{\bar{\theta}_{t-1}(\bar{x}_{t-1})}^{+\infty} \Phi\left(\frac{h_t - h_{t-1}}{\delta}\right) d\Psi_{t-1}^\delta(h_{t-1}; \bar{x}^{t-2})}{1 - \Psi_{t-1}^\delta(\bar{\theta}_{t-1}(\bar{x}_{t-1}); \bar{x}^{t-2})} \quad (12)$$

with $\Psi_1^\delta(h_1) = \Phi(\sqrt{\alpha}(h_1 - z))$. Next, let $\Psi_t^\delta(h_t|x; \bar{x}^{t-1})$ denote the c.d.f. of private posteriors about h_t ; this is obtained by applying Bayes' rule to (12). Then, the expected net payoff from attacking in period t for an agent with signal x is $v_t^\delta(x; \bar{x}_1) = \Psi_1^\delta(\bar{\theta}_1(\bar{x}_1)|x) - c$ for $t = 1$ and

$$v_t^\delta(x; \bar{x}^t) = \Psi_t^\delta(\bar{\theta}_t(\bar{x}_t)|x; \bar{x}^{t-1}) - c$$

for $t \geq 2$. Finally, let V_t^δ denote the payoff of the marginal agent, as defined in condition (11)) but using the function v_t^δ above. With V_t^δ defined this way, the equilibrium algorithm of Proposition 5 applies to the environment examined here as well. What is more, because beliefs—and hence payoffs—again converge to their counterparts in the benchmark game as $\delta \rightarrow 0$, Theorem 3 also applies. (See the Supplementary Material for details.)

The game with short-lived agents thus permits one to examine environments where agents have private information about the innovations in fundamentals while maintaining tractability.

²⁴Note that global-game results do not require that agents have private information about *all* payoff-relevant variables, nor that uncertainty vanishes in the limit for *all* payoff-relevant variables.

6 Conclusion

This paper examined how learning influences the dynamics of coordination in a global game of regime change. Our results struck a delicate balance between the earlier common-knowledge and the more recent global-games literature: the dynamics featured both a refined role for multiplicity and a certain discontinuity of outcomes with respect to changes in information or payoffs. They also led to novel predictions, such as the possibility that fundamentals predict eventual outcomes but not the timing and number of attacks, or that dynamics alternate between phases of tranquility, during which agents accumulate information and no attacks are possible, and phases of distress, during which attacks may occur but do not necessarily take place.

From a methodological perspective, our results offer two lessons with regard to the recent debate about uniqueness versus multiplicity in coordination environments. First, that equilibrium learning can be a natural source of multiplicity in a dynamic setting, despite the heterogeneity of beliefs. Second, and most importantly, that this debate may dilute what, at least in our view, is the central contribution of the global-games approach: the understanding of how the structure of beliefs can lead to interesting and novel predictions about equilibrium behavior well beyond equilibrium determinacy.

From an applied perspective, on the other hand, the predictions we derived may help understand the dynamics of currency attacks, financial crashes, political change, and other crises phenomena. With this in mind, in Section 5 we sought to give some guidance on how the analysis can be extended to accommodate certain features that were absent in the benchmark model but may be important for applications. The scope of these extensions, however, was limited to changes in information or in fundamentals—we remained silent about other dynamic effects (such as those introduced by irreversible actions or liquidity constraints), as well as about the role of large players (such as that of a “Soros” or a policy maker). Extending the analysis in these directions seems a promising line for future research.

Appendix: proofs omitted in the main text

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

$$U^{st}(\hat{\theta}; \beta, \alpha, z) = 0, \quad (13)$$

where

$$U^{st}(\theta; \beta, \alpha, z) \equiv 1 - \Phi\left(\frac{\sqrt{\beta}}{\sqrt{\beta+\alpha}}\left[\Phi^{-1}(\theta) + \frac{\alpha}{\sqrt{\beta}}(z - \theta)\right]\right) - c. \quad (14)$$

Note that $U^{st}(\theta; \cdot)$ is continuous and differentiable 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 (13) therefore always exists. Next, note that

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

Since $\min_{\theta \in (0,1)} [1/\phi(\Phi^{-1}(\theta))] = \sqrt{2\pi}$, the condition $\beta \geq \alpha^2/(2\pi)$ 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 strictly dominated strategies, see Morris and Shin (2001, 2003). ■

Proof of Proposition 2. Necessity follows from the arguments in the main text. For sufficiency, take any sequence $\{x_t^*, \theta_t^*\}_{t=1}^\infty$ that satisfies conditions (ii) and (iii); let $\theta_0^* = -\infty$; suppose all other agents follow strategies as in (i), in which case $R_t = 0$ if and only if $\theta > \theta_{t-1}^*$, for all $t \geq 1$; and consider the best response for an individual agent. If $\theta_t^* = \theta_{t-1}^*$, in which case $t \geq 2$, $\theta_{t-1}^* > 0$ and $x_t^* = -\infty$, then $\Pr(R_{t+1} = 1 | x_t, R_t = 0) = \Pr(\theta \leq \theta_t^* | x_t, \theta > \theta_{t-1}^*) = 0$ for all x_t and therefore not attacking is indeed optimal. If instead $\theta_t^* > \theta_{t-1}^*$, in which case $U(\theta_t^*, \theta_{t-1}^*, \beta_t, \alpha, z) = 0$ and $x_t^* = X(\theta_t^*, \beta_t)$, then, by the monotonicity of the private posterior in x_t and the definitions of $X(\cdot)$ and $U(\cdot)$, $\Pr(\theta \leq \theta_t^* | x_t, \theta > \theta_{t-1}^*) - c \geq (\leq) U(\theta_t^*, \theta_{t-1}^*, \beta_t, \alpha, z)$ if and only if $x_t \leq (\geq) X(\theta_t^*, \beta_t)$ and therefore it is indeed optimal to attack for $x_t < x_t^*$ and not to attack for $x_t > x_t^*$. ■

Proof of Lemma 2. Combining the definitions of u , X and U , we have that

$$U(\theta^*, \theta_{-1}^*, \beta, \alpha, z) = \begin{cases} 1 - c & \text{if } \theta^* = 0 > \theta_{-1}^* \\ 1 - \frac{\Phi\left(\frac{\sqrt{\beta}}{\sqrt{\beta+\alpha}}\left[\Phi^{-1}(\theta^*) + \frac{\alpha}{\sqrt{\beta}}(z - \theta^*)\right]\right)}{\Phi\left(\frac{\sqrt{\beta}}{\sqrt{\beta+\alpha}}\left[\Phi^{-1}(\theta^*) + \frac{\alpha}{\sqrt{\beta}}(z - \theta^*)\right] + \sqrt{\beta+\alpha}(\theta^* - \theta_{-1}^*)\right)} - c & \text{if } \max\{0, \theta_{-1}^*\} < \theta^* < 1 \\ -c & \text{if } \theta^* \leq \theta_{-1}^* \text{ or } \theta^* = 1 > \theta_{-1}^* \end{cases}$$

Part (i) follows by inspecting U .

For (ii), note that $U(\theta^*, \theta_{t-1}^*, \beta_t, \alpha, z) < U(\theta^*, -\infty, \beta_t, \alpha, z)$ for all θ^* (since $\theta_{t-1}^* > -\infty$) and that $U(\theta^*, -\infty, \beta_t, \alpha, z) = U^{st}(\theta^*, \beta_t, \alpha, z)$ is strictly decreasing in θ^* (since $\beta_t \geq \alpha^2/(2\pi)$). It follows that $U(\theta^*, \theta_{t-1}^*, \beta_t, \alpha, z) < 0$ for all $\theta^* \geq \hat{\theta}_t$, which gives the result.

For (iii), take any $\theta_{t-1}^* > \theta_\infty$. Note that for all $\theta^* \in [\theta_{t-1}^*, 1]$, $\lim_{\beta \rightarrow \infty} U(\theta^*, -\infty, \beta, \alpha, z) = \theta_\infty - \theta^* < 0$. Since $U(\theta^*, -\infty, \beta, \alpha, z)$ is continuous in θ^* and $[\theta_{t-1}^*, 1]$ is compact, it follows that

there exists $\bar{\beta}$ such that, for any $\beta > \bar{\beta}$, $U(\theta^*, -\infty, \beta, \alpha, z) < 0$ for all $\theta^* \in [\theta_{t-1}^*, 1]$. Moreover, for all β , $U(\theta^*, \theta_{t-1}^*, \beta, \alpha, z) = -c < 0$ for $\theta \leq \theta_{t-1}^*$ and $U(\theta^*, \theta_{t-1}^*, \beta, \alpha, z) < U(\theta^*, -\infty, \beta, \alpha, z)$ for $\theta^* > \theta_{t-1}^*$. It follows that, for any $\beta > \bar{\beta}$, $U(\theta^*, \theta_{t-1}^*, \beta, \alpha, z) < 0$ for all θ^* and therefore (8) admits no solution.

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

Finally, consider (v). Fix $t \geq 2$, θ_{t-2}^* , β_{t-1} , α , and z (where we use the convention $\theta_0^* = -\infty$) and suppose that θ_{t-1}^* is the highest solution to (8) for period $t-1$, which means that $U(\theta^*, \theta_{t-2}^*, \beta_{t-1}, \alpha, z) < 0$ for all $\theta^* > \theta_{t-1}^*$. This, together with the properties that $U(\theta^*, \theta_{-1}, \beta, \alpha, z)$ is non-increasing in θ_{-1} , continuous in θ^* , and equal to $-c$ for $\theta^* \leq \theta_{-1}$, implies that there exists $\Delta > 0$ such that $U(\theta^*, \theta_{t-1}^*, \beta_{t-1}, \alpha, z) < -\Delta$ for all $\theta^* \in [0, 1]$. Furthermore, by continuity of U in (θ^*, β) , there exists $\bar{\beta} > \beta_{t-1}$ such that $U(\theta^*, \cdot, \beta, \cdot)$ is uniformly continuous over $[0, 1] \times [\beta_{t-1}, \bar{\beta}]$. This also implies that there exists $\underline{\beta} \in (\beta_{t-1}, \bar{\beta})$ such that $U(\theta^*, \theta_{t-1}^*, \beta, \alpha, z) < 0$ for all $\beta \in [\beta_{t-1}, \underline{\beta}]$ and all $\theta^* \in [0, 1]$, which proves that condition (8) admits no solution in any period $\tau > t$ for which $\beta_\tau < \underline{\beta}$. ■

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. For any $\beta \geq \beta_1$ ($\geq \alpha^2/(2\pi)$), let $\hat{\theta} = \hat{\theta}(\beta, \alpha, z)$ be the unique solution to the equation $U(\hat{\theta}, -\infty, \beta, \alpha, z) = 0$ (i.e., the static equilibrium threshold) and

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

The threshold $\tilde{z}(\beta, \alpha)$ is defined by $U(\theta_\infty, -\infty, \beta, \alpha, \tilde{z}(\beta, \alpha)) = 0$ and is such that $\hat{\theta} \geq (\leq) \theta_\infty$ if and only if $z \leq (\geq) \tilde{z}$. The threshold $\hat{z}(\beta, \alpha)$, on the other hand, is defined so that $\partial \hat{\theta} / \partial \beta \geq (\leq) 0$ if and only if $z \geq (\leq) \hat{z}$. To simplify notation, we henceforth suppress the dependence of $\hat{\theta}$ on (α, z) and of \tilde{z} and \hat{z} on α .

First, consider $c = 1/2$, in which case $\hat{z}(\beta) = \tilde{z}(\beta) = 1/2$ for all β . When $z < 1/2$, $\hat{\theta}(\beta) > \theta_\infty$ and $\partial \hat{\theta} / \partial \beta < 0$ for all $\beta \geq \beta_1$ and therefore $\hat{\theta}_1 \geq \hat{\theta}_t > \theta_\infty$ for all t . When instead $z = 1/2$, $\hat{\theta}(\beta) = \theta_\infty$ for any $\beta \geq \beta_1$, and therefore $\hat{\theta}_1 = \hat{\theta}_t = \theta_\infty$ for all t . Finally, when $z > 1/2$, for any $\beta \geq \beta_1$, $\hat{\theta}(\beta) < \theta_\infty$ and $\partial \hat{\theta} / \partial \beta > 0$, and hence $\hat{\theta}_1 \leq \hat{\theta}_t < \theta_\infty$ 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 $\tilde{z}(\beta)$ and $\hat{z}(\beta)$ are both decreasing in β , satisfy $\hat{z}(\beta) > \tilde{z}(\beta) > \theta_\infty$ for all β , and converge to θ_∞ as $\beta \rightarrow \infty$. When $z \leq \theta_\infty$, then clearly $z < \tilde{z}(\beta) < \hat{z}(\beta)$

for all β and therefore $\hat{\theta}(\beta)$ is always higher than θ_∞ and decreasing in β , which implies that $\hat{\theta}_1 \geq \hat{\theta}_t > \theta_\infty$ for all t . When $z \in (\theta_\infty, \tilde{z}(\beta_1))$, there are $\beta'' > \beta' > \beta_1$ such that $\tilde{z}(\beta') = \hat{z}(\beta'') = z$. For $\beta \in [\beta_1, \beta']$, $\hat{\theta}(\beta)$ is higher than θ_∞ and decreases with β . As soon as $\beta \in (\beta', \beta'')$, $\hat{\theta}(\beta)$ becomes lower than θ_∞ and continues to decrease with β . Once $\beta \geq \beta''$, $\hat{\theta}(\beta)$ starts increasing with β , but never exceeds θ_∞ . Hence, $\hat{\theta}_1 > \theta_\infty$ and $\hat{\theta}_1 \geq \hat{\theta}_t$ for all t . When $z = \tilde{z}(\beta_1)$, $\hat{\theta}_1 = \theta_\infty \geq \hat{\theta}_t$ for all t . Finally, when $z > \tilde{z}(\beta_1)$, $\hat{\theta}(\beta) < \theta_\infty$ for all β , and therefore $\hat{\theta}_t < \theta_\infty$ for all t . We conclude that the result holds for $c < 1/2$ with $\underline{z} = \bar{z} = \bar{\bar{z}} = \tilde{z}(\beta_1)$.

Finally, consider $c > 1/2$, in which case $\hat{z}(\beta)$ and $\tilde{z}(\beta)$ are both increasing in β , satisfy $\hat{z}(\beta) < \tilde{z}(\beta) < \theta_\infty$, and converge to θ_∞ as $\beta \rightarrow \infty$. When $z \leq \hat{z}(\beta_1)$, then clearly $z < \hat{z}(\beta) < \tilde{z}(\beta)$ for all $\beta > \beta_1$ and therefore $\hat{\theta}(\beta)$ is always higher than θ_∞ and decreasing in β , which implies that $\hat{\theta}_1 \geq \hat{\theta}_t > \theta_\infty$ for all t . When $z \in (\hat{z}(\beta_1), \tilde{z}(\beta_1))$, there is $\beta' > \beta_1$ such $\hat{z}(\beta') = z$. For $\beta \in (\beta_1, \beta')$, $\hat{\theta}(\beta)$ is higher than θ_∞ and increasing in β , whereas for $\beta > \beta'$, $\hat{\theta}(\beta)$ decreases with β , converging to θ_∞ from above. It follows that $\max_{t \geq 1} \hat{\theta}_t \geq \hat{\theta}_1 > \theta_\infty$. When $z = \tilde{z}(\beta_1)$, $\max_{t \geq 1} \hat{\theta}_t \geq \hat{\theta}_1 = \theta_\infty$. When $z \in (\tilde{z}(\beta_1), \theta_\infty)$, there are $\beta'' > \beta' > \beta_1$ such that $\tilde{z}(\beta') = \hat{z}(\beta'') = z$. For $\beta \in (\beta_1, \beta')$, $\hat{\theta}(\beta)$ is lower than θ_∞ and increasing in β . For $\beta \in (\beta', \beta'')$, $\hat{\theta}(\beta)$ is higher than θ_∞ and increases with β . And for $\beta > \beta''$, $\hat{\theta}(\beta)$ decreases with β 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 $z > \tilde{z}(\beta) > \hat{z}(\beta)$ for all β and therefore $\hat{\theta}(\beta)$ is always lower than θ_∞ , increases with β , and asymptotes θ_∞ from below. Hence, $\hat{\theta}_1 \leq \hat{\theta}_t < \theta_\infty$ for all t . We conclude that the result holds for $c > 1/2$ with $\underline{z} = \hat{z}(\beta_1)$, $\bar{z} = \tilde{z}(\beta_1)$, and $\bar{\bar{z}} = \theta_\infty$. ■

Proof of Proposition 3. Apart from a notational adjustment – namely the dependence of U in period t on (α_t, z_t) and of (x_t^*, θ_t^*) on z^t – the proof follows exactly the same steps as in the model with only private information, and is thus omitted for brevity. ■

Proof of Theorem 2. Consider first $t = 1$. For any (β_1, α_1, z_1) , $U(\theta^*, -\infty, \beta_1, \alpha_1, z_1)$ is continuous in $\theta^* \in [0, 1]$ with $U(0, -\infty, \cdot) = 1 - c$ and $U(1, -\infty, \cdot) = -c$. Hence a solution $\theta_1^*(z_1)$ to $U(\theta_1^*, -\infty, \beta_1, \alpha_1, z_1) = 0$ always exists.²⁵ Next, consider any $t \geq 2$ and note that, for any $(\theta_{t-1}^*, \beta_t, \alpha_t)$ and any $\theta^* \in (\theta_{t-1}^*, 1)$, $U(\theta^*, \theta_{t-1}^*, \beta_t, \alpha_t, z_t)$ is strictly decreasing in z_t and $U(\theta^*, \cdot, z_t) \rightarrow 1 - c > 0$ as $z_t \rightarrow -\infty$, implying that necessarily $\max_{\theta^* \in [\theta_{t-1}^*, 1]} U(\theta^*, \cdot, z_t) > 0$ for z_t sufficiently low. Furthermore, since $U(\theta^*, \cdot, z_t)$ is continuous in $\theta^* \in [\theta_{t-1}^*, 1]$ for any z_t , and since $U(\theta^*, \theta_{t-1}^*, \cdot) \rightarrow -c$ monotonically for any $\theta^* \in [\theta_{t-1}^*, 1]$ as $z_t \rightarrow +\infty$, from standard Monotone Convergence Theorems, the function $U(\theta^*, \theta_{t-1}^*, \cdot, z_t)$ converges uniformly to $-c$ as $z_t \rightarrow +\infty$, implying that $\max_{\theta^* \in [\theta_{t-1}^*, 1]} U(\theta^*, \theta_{t-1}^*, \beta_t, \alpha_t, z_t) < 0$ for z_t sufficiently high. The strict monotonicity of U in z_t then guarantees that there exists a finite $\bar{z}(\theta_{t-1}^*, \beta_t, \alpha_t)$ such that $\max_{\theta^* \in [\theta_{t-1}^*, 1]} U(\theta^*, \theta_{t-1}^*, \beta_t, \alpha_t, z_t) \geq (\leq) 0$ if and only if $z \leq (\geq) \bar{z}(\theta_{t-1}^*, \beta_t, \alpha_t)$, which also implies that (9) admits a solution $\theta_t^*(z^t) > \theta_{t-1}^*(z^{t-1})$ if and only if $z_t \leq \bar{z}(\theta_{t-1}^*, \beta_t, \alpha_t)$. The following

²⁵Note that the function $\theta_1^*(\cdot)$ is unique if and only if $\beta_1 \geq \alpha_1^2/2\pi$. Hence for $\beta_1 < \alpha_1^2/2\pi$, the game trivially admits multiple equilibria even if $T = 1$.

is then an equilibrium: for $t = 1$, $\theta_1^*(z_1)$ is any solution to $U(\theta_1^*, -\infty, \beta_1, \alpha_1, z_1) = 0$; for all $t \in \{2, \dots, T\}$, $\theta_t^*(z^t) = \max(\{\theta_{t-1}^*(z^{t-1})\} \cup \{\theta^* : U(\theta^*, \theta_{t-1}^*(z^{t-1}), \beta_t, \alpha_t, z_t) = 0\})$. Note that, in this equilibrium, at any $t \geq 2$, $\theta_t^*(z^t) > \theta_{t-1}^*(z^{t-1})$ for all $z_t \leq \bar{z}(\theta_{t-1}^*, \beta_t, \alpha_t)$. Since $\theta_t^*(z^t) = \theta_1^*(z_1)$ for all z^t and all t is also an equilibrium, we conclude that the game admits multiple equilibria for any $\{\beta_t, \alpha_t\}_{t=1}^T$ and any $T \geq 2$. ■

Proof of Proposition 4. Parts (i) and (ii) are immediate. For part (iii), note that in each period $t \geq 2$ there are two possible cases: either an attack takes place ($x_t^*(\omega^t) > -\infty$), or not ($x_t^*(\omega^t) = -\infty$).

If $x_t^*(\omega^t) > -\infty$, it must be that $\theta_t^*(\omega^t) > \theta_{t-1}^*(\omega^{t-1})$, for otherwise the posterior probability of regime change would be zero for any x_t and attacking would never be optimal; moreover, it must be that the thresholds $\theta_t^*(\omega^t)$ and $x_t^*(\omega^t)$ solve $A_t(\theta_t^*, \omega^t) = \theta_t^* + \delta\omega_t$ and $\Pr(\theta < \theta_t^* | x_t^*, \omega_t, \theta \geq \theta_{t-1}^*) = c$, or equivalently

$$1 - \frac{\theta_t^*(\omega^t) + \delta\omega_t = \Phi\left(\sqrt{\beta_x}(x_t^*(\omega^t) - \theta_t^*(\omega^t))\right) \Phi\left(\sqrt{\beta_t + \alpha}\left(\frac{\beta}{\beta_t + \alpha}x_t^*(\omega^t) + \frac{\alpha}{\beta_t + \alpha}z - \theta_t^*(\omega^t)\right)\right)}{\Phi\left(\sqrt{\beta_t + \alpha}\left(\frac{\beta}{\beta_t + \alpha}x_t^*(\omega^t) + \frac{\alpha}{\beta_t + \alpha}z - \theta_{t-1}^*(\omega^{t-1})\right)\right)} = c.$$

Using the definitions of X and U from the benchmark model, the above two conditions are equivalent to $x_t^*(\omega^t) = X(\theta_t^*(\omega^t) + \delta\omega_t; \beta_t) - \delta\omega_t$ and (10). And conversely, if (10) admits a solution, then there exists an equilibrium with an attack in period t . This establishes the first half of part (iii).

If, on the other hand, $x_t^*(\omega^t) = -\infty$, it must be that $\theta_{t-1}^*(\omega^{t-1}) + \delta\omega_t \geq 0$, for otherwise it would be dominant for some agents to attack. And conversely, $\theta_{t-1}^*(\omega^{t-1}) + \delta\omega_t \geq 0$ ensures that there is an equilibrium in which no attack takes place in period t . This establishes the second half of part (iii). ■

Lemma A2 For $t = 1$, $V_1^\delta(\bar{x}_1)$ is continuous in \bar{x}_1 for any $\bar{x}_1 \in \mathbb{R}$; and for $t \geq 2$, $V_t^\delta(\bar{x}^t)$ is continuous in \bar{x}^t for any $\bar{x}^t \in \overline{\mathbb{R}}^{t-1} \times \mathbb{R}$.²⁶

Proof. Consider first $\delta = 0$, in which case $V_1^0(\bar{x}_1) \equiv U(\bar{\theta}_1(\bar{x}_1), -\infty, \beta_1, \alpha, z)$ and $V_t^0(\bar{x}^t) \equiv U(\bar{\theta}_t(\bar{x}_t), \bar{\theta}_{t-1}(\bar{x}^{t-1}), \beta_t, \alpha, z)$ for $t \geq 2$. Note that, for all t , $\bar{\theta}_t(\bar{x}^t) \equiv \min\{\theta : \theta \geq \Phi(\sqrt{\beta_\tau}(\bar{x}_\tau - \theta)) \forall \tau \leq t\}$ is continuous in $\bar{x}^t \in \overline{\mathbb{R}}^t$ and takes values in $[0, 1]$. Furthermore, $U(\theta, -\infty, \beta, \alpha, z)$ is

²⁶Continuity can be extended in $\overline{\mathbb{R}}^t$ as follows. For any function $f : A \rightarrow \overline{\mathbb{R}}$, where $A \subseteq \overline{\mathbb{R}}^t$ and $t \geq 1$, we say that f is continuous over A if and only if, for any $x^t \in A$ and any $\varepsilon > 0$, there exists $\eta > 0$ such that, for any $\tilde{x}^t \in A$ such that for all $\tau \leq t$: (a) $|\tilde{x}_\tau - x_\tau| < \eta$ if $x_\tau \in \mathbb{R}$; (b) $\tilde{x}_\tau < -1/\eta$ if $x_\tau = -\infty$; (c) $\tilde{x}_\tau > 1/\eta$, if $x_\tau = +\infty$, the following is true: (a') if $f(x^t) \in \mathbb{R}$, then $|f(\tilde{x}^t) - f(x^t)| < \varepsilon$; (b') if $f(x^t) = -\infty$, then $f(\tilde{x}^t) < -1/\varepsilon$; (c') if $f(x^t) = +\infty$, then $f(\tilde{x}^t) > 1/\varepsilon$.

Note that, if $f : A \rightarrow \overline{\mathbb{R}}$, $g : B \rightarrow \overline{\mathbb{R}}$, and $q : C \rightarrow \overline{\mathbb{R}}$ are continuous, respectively, in A , B and C , where $A \subseteq \overline{\mathbb{R}}^t$, $B \subseteq \overline{\mathbb{R}}^k$, and $f(A) \times g(B) \subseteq C \subseteq \overline{\mathbb{R}}^2$, then the function $w : A \times B \rightarrow \overline{\mathbb{R}}$ defined by $w(x^t, x^k) = q(f(x^t), g(x^k))$ is continuous in $A \times B$.

continuous in $\theta \in [0, 1]$ and $U(\theta, \theta_{-1}, \beta, \alpha, z)$ is continuous in $(\theta, \theta_{-1}) \in [0, 1]^2$. It follows that, for all t , $V_t^0(\bar{x}^t)$ is continuous in $\bar{\mathbb{R}}^t$.

Consider next $\delta > 0$. For all $t \geq 1$, the function $p_t^\delta(\theta; \bar{x}_t) = F([\Phi(\sqrt{\beta_t}(\bar{x}_t - \theta)) - \theta]/\delta)$ is continuous in $\theta \in \mathbb{R}$ and $\bar{x}_t \in \bar{\mathbb{R}}$; increasing in \bar{x}_t , and decreasing in θ ; it is bounded in $[0, 1]$; and it satisfies $\lim_{\theta \rightarrow -\infty} p_t^\delta(\theta; \bar{x}_t) = 0$ and $\lim_{\theta \rightarrow +\infty} p_t^\delta(\theta; \bar{x}_t) = 1$ for any $\bar{x}_t \in \bar{\mathbb{R}}$. The p.d.f. of the private posteriors for $t = 1$, $\psi_1^\delta(\theta|x) = \phi\left(\sqrt{\beta_1 + \alpha}\left(\theta - \frac{\beta_1 x + \alpha z}{\beta_1 + \alpha}\right)\right)$, is clearly continuous in $\theta \in \mathbb{R}$ and $x \in \mathbb{R}$; and similarly for the c.d.f. Ψ_1^δ . It follows that $v_1(x; \bar{x}_1) = \int_{-\infty}^{+\infty} p_1^\delta(\theta; \bar{x}_1) d\Psi_1(\theta|x) - c$ is continuous in $(x, \bar{x}_1) \in \mathbb{R} \times \bar{\mathbb{R}}$. For any $t \geq 2$, from Bayes' rule,

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

which is also continuous in $\theta \in \mathbb{R}$ and $(x, \bar{x}^{t-1}) \in \mathbb{R} \times \bar{\mathbb{R}}^{t-1}$; and similarly for Ψ_t^δ .²⁷ It follows that $v_t(x; \bar{x}^{t-1}, \bar{x}_t) = \int_{-\infty}^{+\infty} p_t^\delta(\theta; \bar{x}_t) d\Psi_t^\delta(\theta|x; \bar{x}^{t-1}) - c$ is continuous in $(x, \bar{x}^{t-1}, \bar{x}_t) \in \mathbb{R} \times \bar{\mathbb{R}}^t$. Moreover, for all t , since $p_t^\delta(\theta; \bar{x}_t)$ is bounded in $[0, 1]$, $v_t^\delta(x; \bar{x}^t)$ is bounded in $[-c, 1 - c]$. In addition, since the distribution of x given θ satisfies the MLRP and $p_t^\delta(\theta; \bar{x}_t)$ is decreasing in θ , by standard representation theorems (Milgrom, 1981) we have that $v_t^\delta(x; \bar{x}^t)$ is decreasing in $x \in \mathbb{R}$. It follows that $\lim_{x \rightarrow -\infty} v_t^\delta(x; \bar{x}^t)$ and $\lim_{x \rightarrow +\infty} v_t^\delta(x; \bar{x}^t)$ exist for any $\bar{x}^t \in \bar{\mathbb{R}}^t$ and therefore $V_t^\delta(\bar{x}^t)$ is well-defined for $\bar{x}_t = \pm\infty$. Finally, since $v_t^\delta(x; \bar{x}^{t-1}, \bar{x}_t)$ is continuous in $(x, \bar{x}^{t-1}, \bar{x}_t) \in \mathbb{R} \times \bar{\mathbb{R}}^{t-1} \times \mathbb{R}$, it is immediate that $V_t^\delta(\bar{x}^{t-1}, \bar{x}_t) = v_t^\delta(\bar{x}_t; \bar{x}^{t-1}, \bar{x}_t)$ is continuous in $(\bar{x}^{t-1}, \bar{x}_t) \in \bar{\mathbb{R}}^{t-1} \times \mathbb{R}$. ■

Proof of Proposition 5. Sufficiency. Consider a sequence $\{x_t^*\}_{t=1}^\infty$ that satisfies conditions (ii) and (iii) in the proposition. The monotonicity of $v_t^\delta(x; \bar{x}^t)$ with respect to x (see proof of Lemma A2 above) guarantees that, for any $x \in \mathbb{R}$, $v_t^\delta(x; x^{*t}) \geq (\leq) V_t^\delta(x^{*t})$ if and only if $x \leq (\geq) x_t^*$. It follows that the strategies defined by (i) – (iii) constitute a monotone equilibrium.

Necessity. Conversely, suppose that $\{a_t(\cdot)\}_{t=1}^\infty$ is a monotone equilibrium. Since in any such equilibrium the measure of agents attacking in every period is decreasing in θ , the probability of regime change is also decreasing in θ . Then, by standard representation theorems (Milgrom, 1981), the expected payoff from attacking is decreasing in x_t , implying that agents must follow cut-off strategies. For $\{x_t^*\}_{t=1}^\infty$ to be equilibrium cutoffs, it must be that, for all t , $V_t^\delta(x^{*t}) = 0$ if $-\infty < x_t^* < +\infty$, $V_t^\delta(x^{*t}) \leq 0$ if $x_t^* = -\infty$, and $V_t^\delta(x^{*t}) \geq 0$ if $x_t^* = +\infty$.

We next show that, in any equilibrium, $x_t^* < +\infty$ for all $t \geq 1$ and $x_1^* > -\infty$. Indeed, if $x_t^* = +\infty$, in which case $p_t^\delta(\theta; x_t^*) = F((1 - \theta)/\delta)$, then, for any $t \geq 2$, $(x, \bar{x}^{t-1}) \in \mathbb{R} \times \bar{\mathbb{R}}^{t-1}$, and

²⁷To see this, note that the function q defined by $q(p_1, \dots, p_{t-1}, \phi) = \prod_{s=1}^{t-1} [1 - p_s] \phi$ is continuous in $[0, 1]^{t-1} \times \mathbb{R}$, each p_s is continuous in $\bar{x}_s \in \bar{\mathbb{R}}$, and ϕ in $x \in \mathbb{R}$.

$\theta' \in \mathbb{R}$,

$$\begin{aligned} v_t^\delta(x; \bar{x}^{t-1}, +\infty) &= \int_{-\infty}^{+\infty} F\left(\frac{1}{\delta}(1-\theta)\right) \psi_t^\delta(\theta|x; \bar{x}^{t-1}) d\theta - c \\ &= \int_{-\infty}^{\theta'} F\left(\frac{1}{\delta}(1-\theta)\right) \psi_t^\delta(\theta|x; \bar{x}^{t-1}) d\theta + \int_{\theta'}^{+\infty} F\left(\frac{1}{\delta}(1-\theta)\right) \psi_t^\delta(\theta|x; \bar{x}^{t-1}) d\theta - c \\ &\leq \Psi_t^\delta(\theta'|x; \bar{x}^{t-1}) + F\left(\frac{1}{\delta}(1-\theta')\right) \left[1 - \Psi_t^\delta(\theta'|x; \bar{x}^{t-1})\right] - c, \end{aligned}$$

where $\Psi_t^\delta(\theta'|x; \bar{x}^{t-1}) = \int_{-\infty}^{\theta'} \psi_t^\delta(\theta|x; \bar{x}^{t-1}) d\theta$. Furthermore, since the knowledge that the status quo survived past attacks causes a first-order-stochastic-dominance change in posterior beliefs,²⁸ $\Psi_t^\delta(\theta'|x; \bar{x}^{t-1}) \leq \Phi\left(\sqrt{\beta_t + \alpha}\left(\theta' - \frac{\beta_t x + \alpha z}{\beta_t + \alpha}\right)\right)$. Along with $\lim_{x \rightarrow +\infty} \Phi\left(\sqrt{\beta_t + \alpha}\left(\theta' - \frac{\beta_t x + \alpha z}{\beta_t + \alpha}\right)\right) = 0$, this implies that $\lim_{x \rightarrow +\infty} v_t^\delta(x; \bar{x}^{t-1}, +\infty) \leq F((1-\theta')/\delta) - c$. Since the latter is true for any $\theta' \in \mathbb{R}$, it is also true for $\theta' \rightarrow +\infty$, in which case $F((1-\theta')/\delta) \rightarrow 0$. Together with the fact that v_t^δ is bounded from below by $-c$, this implies that $V^\delta(\bar{x}^{t-1}, +\infty) = \lim_{x \rightarrow +\infty} v_t^\delta(x; \bar{x}^{t-1}, +\infty) = -c < 0$ and hence $x_t^* = +\infty$ can not be part of any equilibrium. A similar argument rules out $x_1^* = +\infty$. Finally, suppose $x_1^* = -\infty$. Then, for any $x \in \mathbb{R}$ and any $\theta' \in \mathbb{R}$,

$$v_1^\delta(x; -\infty) = \int_{-\infty}^{+\infty} F\left(\frac{1}{\delta}(-\theta)\right) \psi_1^\delta(\theta|x) d\theta - c \geq \Psi_1^\delta(\theta'|x) F\left(\frac{1}{\delta}(-\theta')\right) - c,$$

where $\Psi_1^\delta(\theta'|x) = \int_{-\infty}^{\theta'} \psi_1^\delta(\theta|x) d\theta$, and therefore $\lim_{x \rightarrow -\infty} v_1^\delta(x; -\infty) \geq F((-\theta')/\delta) - c$. Since this is true also for $\theta' \rightarrow -\infty$, and since v_1^δ is bounded from above by $1 - c$, we have that $V_1^\delta(-\infty) = \lim_{x \rightarrow -\infty} v_1^\delta(x; -\infty) = 1 - c > 0$, implying that $x_1^* = -\infty$ can not be part of an equilibrium.

We conclude that (i) – (iii) necessarily hold in any monotone equilibrium.

Existence. For any $\delta > 0$, the monotonicity of $v_1^\delta(x; \bar{x}_1)$ in \bar{x}_1 along with its continuity in x for any \bar{x}_1 and the fact that $\lim_{x \rightarrow -\infty} v_1^\delta(x, -\infty) > 0 > \lim_{x \rightarrow +\infty} v_1^\delta(x, +\infty)$, implies that there exist $x', x'' \in \mathbb{R}$ such that $V_1^\delta(x') \geq v_1^\delta(x', -\infty) > 0 > v_1^\delta(x'', +\infty) \geq V_1^\delta(x'')$. The continuity of $V_1^\delta(\bar{x}_1)$ in \bar{x}_1 then ensures existence of a solution $x_1^* \in (x', x'')$ to $V_1^\delta(x_1^*) = 0$.

Next, consider $t \geq 2$. For any given \bar{x}^{t-1} , a similar argument as above ensures the existence of $x'' \in \mathbb{R}$ such that $V_t^\delta(\bar{x}^{t-1}, x'') \leq v_t^\delta(x'', \bar{x}^{t-1}, +\infty) < 0$. Moreover, either there also exists $x' \in \mathbb{R}$ such that $V_t^\delta(\bar{x}^{t-1}, x') \geq 0$, or $V_t^\delta(\bar{x}^{t-1}, \bar{x}_t) < 0$ for all $\bar{x}_t \in \mathbb{R}$. In the former case, the continuity of $V_t^\delta(\bar{x}^{t-1}, \bar{x}_t)$ in \bar{x}_t ensures the existence of $\bar{x}_t \in (x', x'')$ such that $V_t^\delta(\bar{x}^{t-1}, \bar{x}_t) = 0$. In the latter case, $v_t^\delta(x; \bar{x}^{t-1}, -\infty) \leq v_t^\delta(x; \bar{x}^{t-1}, x) = V_t^\delta(\bar{x}^{t-1}, x) < 0$ for any $x \in \mathbb{R}$ and therefore at $\bar{x}_t = -\infty$, $V_t^\delta(\bar{x}^{t-1}, -\infty) \equiv \lim_{x \rightarrow -\infty} v_t^\delta(x; \bar{x}^{t-1}, -\infty) \leq 0$. We conclude that there exists a sequence $\{x_t^*\}_{t=0}^\infty$ that satisfies conditions (ii) and (iii) in the proposition. ■

Proof of Theorem 3. We prove the result in four steps. Step 1 uses the structure of beliefs and payoffs to establish that V_t^δ converges pointwise to V_t^0 as $\delta \rightarrow 0$. Steps 2 and 3 then use this pointwise convergence of payoffs to prove the result by induction: Step 2 proves that the result

²⁸This can be seen by noting that the ratio of the densities $\psi_t^\delta(\theta|x; \bar{x}^{t-1})/\sqrt{\beta_t + \alpha}\phi\left(\sqrt{\beta_t + \alpha}\left(\theta - \frac{\beta_t x + \alpha z}{\beta_t + \alpha}\right)\right)$ is increasing in θ .

holds for $T = 1$, while Step 3 proves that if the result holds for $T' = T - 1$, then it holds also for $T' = T$.

We start by establishing pointwise convergence of V_t^δ as $\delta \rightarrow 0$.

Step 1. First, note that, for any $t \geq 1$ any $\bar{x}_t \in \overline{\mathbb{R}}$ and any $\theta \neq \bar{\theta}_t(\bar{x}_t)$,

$$\lim_{\delta \rightarrow 0} p_t^\delta(\theta; \bar{x}_t) = p_t^0(\theta; \bar{x}_t) \equiv \begin{cases} 1 & \text{if } \theta \leq \bar{\theta}_t(\bar{x}_t), \\ 0 & \text{if } \theta > \bar{\theta}_t(\bar{x}_t). \end{cases} \quad (15)$$

This implies that, for any $t \geq 2$, any $\bar{x}^{t-1} \in \overline{\mathbb{R}}^{t-1}$ and any θ ,

$$\lim_{\delta \rightarrow 0} \psi_t^\delta(\theta; \bar{x}^{t-1}) = \psi_t^0(\theta; \bar{x}^{t-1}) \equiv \begin{cases} 0 & \text{if } \theta \leq \bar{\theta}_{t-1}(\bar{x}^{t-1}) \\ \frac{\sqrt{\alpha}\phi(\sqrt{\alpha}(\theta-z))}{1-\Phi(\sqrt{\alpha}(\bar{\theta}_{t-1}(\bar{x}^{t-1})-z))} & \text{otherwise} \end{cases} \quad (16)$$

and hence $\lim_{\delta \rightarrow 0} \psi_t^\delta(\theta|x; \bar{x}^{t-1}) = \psi_t^0(\theta|x; \bar{x}^{t-1})$ for any $x \in \mathbb{R}$. From (15), it follows that, at $t = 1$, for any $\bar{x}_1 \in \mathbb{R}$

$$\begin{aligned} \lim_{\delta \rightarrow 0} V_1^\delta(\bar{x}_1) &= \lim_{\delta \rightarrow 0} \int_{-\infty}^{+\infty} p_1^\delta(\theta; \bar{x}_1) d\Psi_1^\delta(\theta|\bar{x}_1) - c \\ &= \Psi_1^0(\bar{\theta}_1(\bar{x}_1) | \bar{x}_1) - c \\ &= U(\bar{\theta}_1(\bar{x}_1); -\infty, \beta_1, \alpha, z) \equiv V_1^0(\bar{x}_1). \end{aligned}$$

Similarly, by (15) and (16), for any $t \geq 2$, any $(\bar{x}^{t-1}, \bar{x}_t) \in \overline{\mathbb{R}}^{t-1} \times \mathbb{R}$ and any θ , we have that

$$\begin{aligned} \lim_{\delta \rightarrow 0} V_t^\delta(\bar{x}^{t-1}, \bar{x}_t) &= \lim_{\delta \rightarrow 0} \int_{-\infty}^{+\infty} p_t^\delta(\theta; \bar{x}_t) d\Psi_t^\delta(\theta|\bar{x}_t; \bar{x}^{t-1}) - c \\ &= \Psi_t^0(\bar{\theta}_t(\bar{x}_t) | \bar{x}_t; \bar{x}^{t-1}) - c \\ &= U(\bar{\theta}_t(\bar{x}_t); \bar{\theta}_{t-1}(\bar{x}^{t-1}), \beta_t, \alpha, z) \equiv V_t^0(\bar{x}^{t-1}, \bar{x}_t). \end{aligned}$$

We next prove the result by induction.

Step 2. Consider first $T = 1$ and fix an arbitrary $\varepsilon > 0$. From the strict monotonicity of $V_1^0(\bar{x}_1)$,²⁹

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

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

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

From the continuity of $V_1^\delta(\bar{x}_1)$ in \bar{x}_1 for any $\delta > 0$, it follows that there exists a solution x_1^δ to $V_1^\delta(x_1) = 0$ such that $x_1^* - \varepsilon < x_1^\delta < x_1^* + \varepsilon$. Following the same steps as in the proof of existence in Proposition 5, we can then construct an equilibrium $\{x_t^\delta\}_{t=1}^\infty$ for $\Gamma(\delta)$ such that $|\bar{x}_1(\delta) - x_1^*| < \varepsilon$. This proves the result for $T = 1$.

²⁹This follows from the monotonicity of $U(\theta; -\infty, \beta_1, \alpha, z)$ in θ – which in turn is implied by $\beta_1 \geq \alpha^2/\sqrt{2\pi}$ – and the monotonicity of $\bar{\theta}_1(\bar{x}_1)$ in \bar{x}_1 .

Step 3. Consider next an arbitrary $T \geq 2$, fix $\varepsilon > 0$, and suppose the result holds for $T - 1$. We seek to prove that the result holds for T . In doing so, we distinguish two cases: Step 3.a below considers the case that $x_T^* > -\infty$, whereas Step 3.b considers the case that $x_T^* = -\infty$.

Step 3.a. Take first any equilibrium of $\Gamma(0)$ such that $x_T^* > -\infty$. By the (local) strict monotonicity of V_T^0 around x_T^* implied by the assumption that $x_T^* \notin \arg \max_x V_T^0(x^{*T-1}, x)$, there exists $\varepsilon_T < \varepsilon$ such that either

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

or $V_T^0(x^{*T-1}, x_T^* - \varepsilon_T) < 0 < V_T^0(x^{*T-1}, x_T^* + \varepsilon_T)$. Without loss of generality, assume the first case—the argument for the other case is identical. From the continuity of $V_T^0(x^{T-1}, x_T)$ in $x^{T-1} \in \overline{\mathbb{R}}^{T-1}$ and the fact that the result holds for $T - 1$, there exists *some* $\varepsilon'_T \in (0, \varepsilon_T)$ such that, for *any* $\delta < \delta(\varepsilon'_T, T - 1)$, there is a sequence $x^{\delta, T-1}$ satisfying the following three conditions:³⁰

[C1] for all $t \leq T - 1$, either $x_t^\delta = -\infty$ and $V_t^\delta(x^{\delta, t}) \leq 0$, or $x_t^\delta \in \mathbb{R}$ and $V_t^\delta(x^{\delta, t}) = 0$;

[C2] for all $t \leq T - 1$, $|x_t^* - x_t^\delta| < \varepsilon'_T < \varepsilon$ if $x_t^* \in \mathbb{R}$, and $x_t^\delta < -1/\varepsilon'_T < -1/\varepsilon$ if $x_t^* = -\infty$;

[C3] in period T ,

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

Next, by the convergence of $V_T^\delta(x^{T-1}, x_T)$ to $V_T^0(x^{T-1}, x_T)$ for any $(x^{T-1}, x_T) \in \overline{\mathbb{R}}^{T-1} \times \mathbb{R}$, there exists $\delta_T \in (0, \delta(\varepsilon'_T, T - 1))$ such that, for any $\delta < \delta_T$, there is $x^{\delta, T-1}$ that satisfies [C1]-[C2] and such that:

[C3'] in period T ,

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

But then, by the continuity of $V_T^\delta(x^{T-1}, x_T)$ in x_T , for the same $x^{\delta, T-1}$, there exists an $x_T^\delta \in \mathbb{R}$, with $|x_T^* - x_T^\delta| < \varepsilon_T < \varepsilon$, that solves $V_T^\delta(x^{\delta, T-1}, x_T^\delta) = 0$.

Step 3.b. Next, take any equilibrium of $\Gamma(0)$ such that $x_T^* = -\infty$. Recall that, for any $t \geq 2$, $V_t^0(x^{*t-1}, x_t) = -c < 0$ for all $x_t < \tilde{x}_t$, where $\tilde{x}_t > -\infty$ solves $\bar{\theta}_t(\tilde{x}_t) = \bar{\theta}_{t-1}(x^{*t-1}) \equiv \max_{\tau \leq t-1} \bar{\theta}_\tau(x_\tau^*)$. Pick some $x'_T \in (-\infty, \min\{-1/\varepsilon, \tilde{x}_T\})$. From the continuity of $V_T^0(x^{T-1}, x_T)$ in x^{T-1} and the fact that the result holds for $T - 1$, there exists *some* $\varepsilon' \in (0, \varepsilon)$ such that, for any $\delta < \delta(\varepsilon', T - 1)$, there is a sequence $x^{\delta, T-1}$ which satisfies conditions [C1]-[C2] above (replacing ε'_T with ε') and such that:

[C4] $V_T^0(x^{\delta, T-1}, x'_T) < 0$.

By the pointwise convergence of V_T^δ to V_T^0 , there also exists a $\delta_T \in (0, \delta(\varepsilon', T - 1))$ such that, for any $\delta < \delta_T$, there is $x^{\delta, T-1}$ that satisfies [C1]-[C2] and such that:

[C4'] $V_T^\delta(x^{\delta, T-1}, x'_T) < 0$.

If, for the same $x^{\delta, T-1}$, there exists an $x''_T \in (-\infty, x'_T)$ such that $V_T^\delta(x^{\delta, T-1}, x''_T) \geq 0$, then, by the

³⁰Continuity of V^0 implies existence of ε'_T such that [C3] holds for *any* $x^{\delta, T-1}$ that satisfies [C2]; that the result holds for $T - 1$ then ensures that, for any $\delta < \delta(\varepsilon'_T, T - 1)$, there exists $x^{\delta, T-1}$ that satisfies both [C1] and [C2].

continuity of $V_T^\delta(x^{\delta, T-1}, x_T)$ in $x_T \in \mathbb{R}$, there is also an $x_T^\delta \in \mathbb{R}$, with $x_T'' < x_T^\delta < x_T' < -1/\varepsilon$, such that $V_T^\delta(x^{\delta, T-1}, x_T^\delta) = 0$. If instead $V_T^\delta(x^{\delta, T-1}, x_T) < 0$ for all $x_T \in (-\infty, x_T')$, then $x_T^\delta = -\infty$ satisfies $V_T^\delta(x^{\delta, T-1}, -\infty) \leq 0$.³¹

Finally, recall that (8) admits at most two solutions in every t and therefore the set of x^{T*} that can be part of an equilibrium of $\Gamma(0)$ is finite. Hence, there is $\delta(\varepsilon, T) \in (0, \delta(\varepsilon, T-1))$ such that, for any $\delta < \delta(\varepsilon, T)$ and every equilibrium $\{x_t^*\}_{t=1}^\infty$ of $\Gamma(0)$ for which $x_t^* \notin \arg \max_x V_t^0(x^{*t-1}, x)$ for all $t \leq T$, there exists $x^{\delta, T}$ such that, for all $t \leq T$: if $x_t^* \in \mathbb{R}$, then $|x_t^\delta - x_t^*| < \varepsilon$ and $V_t^\delta(x^{\delta, t}) = 0$; and if $x_t^* = -\infty$, then $x_t^\delta < -1/\varepsilon$ and $V_t^\delta(x^{\delta, t}) \leq 0$. From the same arguments as for the proof of existence in Proposition 5, we conclude that $x^{\delta, T}$ is part of an equilibrium $\{x_t^\delta\}_{t=1}^\infty$ for $\Gamma(\delta)$, which completes the proof. ■

References

- [1] Abreu, Dilip, and Markus Brunnermeier (2003): “Bubbles and Crashes,” *Econometrica*, 71, 173-204.
- [2] Angeletos, George-Marios, Christian Hellwig, and Alessandro Pavan (2006): “Signaling in a Global Game: Coordination and Policy Traps,” *Journal of Political Economy*, 114-3, 452-484.
- [3] Angeletos, George-Marios, and Iván Werning (2006): “Crises and Prices: Information Aggregation, Multiplicity and Volatility,” *American Economic Review*, forthcoming.
- [4] Atkeson, Andrew (2001): “Discussion on Morris and Shin,” *NBER Macroeconomics Annual 2000*.
- [5] Broner, Fernando (2005): “Discrete Devaluations and Multiple Equilibria in a First Generation Model of Currency Crises,” working paper, University Pompeu Fabra.
- [6] Carlsson, Hans, and Eric van Damme (1993): “Global Games and Equilibrium Selection,” *Econometrica* 61, 5, 989-1018.
- [7] Chamley, Christophe (1999): “Coordinating Regime Switches,” *Quarterly Journal of Economics* 114, 3, 869-905.
- [8] Chamley, Christophe (2003): “Dynamic Speculative Attacks,” *American Economic Review* 93, 3, 603-621.

³¹This follows from the same argument used in the proof of Proposition 5.

- [9] Chari, V.V., and Patrick Kehoe (2003): “Hot Money,” *Journal of Political Economy* 111, 1262-1292.
- [10] Corsetti, Giancarlo, Bernardo Guimaraes and Nouriel Roubini, (2005): “International Lending of Last Resort and Moral Hazard: A Model of IMF’s Catalytic Finance,” *Journal of Monetary Economics*, forthcoming.
- [11] Dasgupta, Amil (2006): “Coordination and Delay in Global Games” *Journal of Economic Theory*, forthcoming.
- [12] Edmond, Chris (2005): “Information Manipulation, Coordination and Regime Change,” working paper, NYU Stern School of Business.
- [13] Frankel, David, and Ady Pauzner (2000): “Resolving Indeterminacy in Dynamic Settings: The Role of Shocks,” *Quarterly Journal of Economics* 115, 283-304.
- [14] Giannitsarou, Chryssi, and Flavio Toxvaerd (2003): “Recursive Global Games,” working paper, Universidade Nova de Lisboa and Hebrew University of Jerusalem.
- [15] Goldstein, Itay (2005): “Strategic Complementarities and the Twin Crises,” *Economic Journal*, 115, 368-390.
- [16] Goldstein, Itay, and Ady Pauzner (2004): “Contagion of Self-Fulfilling Financial Crises due to Diversification of Investment Portfolios,” *Journal of Economic Theory*, 119(1), 151-183.
- [17] Goldstein, Itay, and Ady Pauzner (2005): “Demand Deposit Contracts and the Probability of Bank Runs,” *Journal of Finance*, 60(3), 1293-1328.
- [18] Heidhues, Paul, and Nicolas Melissas (2006): “Equilibria in a Dynamic Global Game: The Role of Cohort Effects”, *Economic Theory*, 28, 531-557.
- [19] Hellwig, Christian (2002): “Public Information, Private Information and the Multiplicity of Equilibria in Coordination Games,” *Journal of Economic Theory*, 107, 191-222.
- [20] Hellwig, Christian, Arijit Mukerji, and Aleh Tsyvinski (2006): “Self-fulfilling Currency Crises: the Role of Interest Rates,” *American Economic Review*, forthcoming.
- [21] Krugman, Paul (1979): “A model of balance-of-payments crises,” *Journal of Money, Credit, and Banking* 11, 311-25.
- [22] Levin, Jonathan (2001): “A Note on Global Games with Overlapping Generations,” working paper, Stanford University.
- [23] Milgrom, Paul (1981): “Good News and Bad News: Representation Theorems and Applications,” *Bell Journal of Economics* 12, 380-391.

- [24] Morris, Stephen, and Hyun Song Shin (1998): “Unique Equilibrium in a Model of Self-Fulfilling Currency Attacks,” *American Economic Review*, 88, 3, 587-597.
- [25] Morris, Stephen and Hyun Song Shin (1999): “A Theory of the Onset of Currency Attacks,” in P.-R. Agénor, M. Miller D. Vines, and A. Weber, eds., *Asian Financial Crisis: Causes, Contagion and Consequences*, Cambridge, UK: Cambridge University Press.
- [26] Morris, Stephen, and Hyun Song Shin (2001): “Rethinking Multiple Equilibria in Macroeconomics,” *NBER Macroeconomics Annual 2000*.
- [27] Morris, Stephen, and Hyun Song Shin (2003): “Global Games—Theory and Applications,” in M. Dewatripont, L. Hansen, and S. Turnovsky, eds., *Advances in Economics and Econometrics (8th World Congress of the Econometric Society)*, Cambridge, UK: Cambridge University Press.
- [28] Morris, Stephen and Hyun Song Shin (2004): “Coordination Risk and the Price of Debt,” *European Economic Review* 48, 133-153.
- [29] Obstfeld, Maurice (1996): “Models of Currency Crises with Self-Fulfilling Features,” *European Economic Review* 40, 3-5, 1037-47.
- [30] Rochet, Jean-Charles and Xavier Vives (2004): “Coordination Failures and the Lender of Last Resort: Was Bagehot Right After All?,” *Journal of the European Economic Association*, 6(2), 1116-1147.
- [31] Tarashev, Nikola (2005): “Currency Crises and the Informational Role of Interest Rates,” discussion paper, Bank of International Settlements.