

Signaling in a Global Game: Coordination and Policy Traps

George-Marios Angeletos

Massachusetts Institute of Technology and National Bureau of Economic Research

Christian Hellwig

University of California, Los Angeles

Alessandro Pavan

Northwestern University

This paper introduces signaling in a global game so as to examine the informational role of policy in coordination environments such as currency crises and bank runs. While exogenous asymmetric information has been shown to select a unique equilibrium, we show that the endogenous information generated by policy interventions leads to multiple equilibria. The policy maker is thus trapped into a

We are grateful to the editor, Nancy Stokey, and two anonymous referees for suggestions that greatly helped us improve the paper. For comments we thank Daron Acemoglu, Andy Atkeson, Philippe Bacchetta, Gadi Barlevy, Marco Bassetto, Olivier Blanchard, Ricardo Caballero, V. V. Chari, Eddie Dekel, Glenn Ellison, Paul Heidhues, Patrick Kehoe, Robert Lucas Jr., Narayana Kocherlakota, Kiminori Matsuyama, Stephen Morris, Balázs Szentes, Jean Tirole, Muhamet Yildiz, Iván Werning, and seminar participants at Athens University of Economics and Business, Berkeley, Carnegie-Mellon, Harvard, Iowa, Lausanne, London School of Economics, Mannheim, Massachusetts Institute of Technology, Northwestern, Stanford, Stony Brook, Toulouse, University of California at Los Angeles, University of Pennsylvania, the Federal Reserve Banks of Boston, Cleveland, and Minneapolis, the 2002 annual meeting of the Society for Economic Dynamics, the 2002 Universitat Pompeu Fabra workshop on coordination games, the 2003 Centre for Economic Policy Research European Summer Symposium in International Macroeconomics, and the 2003 workshop of the Stanford Institute for Theoretical Economics. We also thank Emily Gallagher for excellent editing assistance. Finally, Angeletos thanks the Federal Reserve Bank of Minneapolis for its hospitality during the latest revision of the paper.

[*Journal of Political Economy*, 2006, vol. 114, no. 3]
© 2006 by The University of Chicago. All rights reserved. 0022-3808/2006/11403-0002\$10.00

position in which self-fulfilling expectations dictate not only the coordination outcome but also the optimal policy. This result does not rely on the freedom to choose out-of-equilibrium beliefs, nor on the policy being a public signal; it may obtain even if the policy is observed with idiosyncratic noise.

I. Introduction

Coordination failures are often invoked as justification for government intervention; they play a prominent role in bank runs, currency attacks, debt crises, investment crashes, adoption of new technologies, and sociopolitical change.

A vast literature models these phenomena as coordination games featuring multiple equilibria (e.g., Diamond and Dybvig 1983; Katz and Shapiro 1986; Obstfeld 1986, 1996; Calvo 1988; Cooper and John 1988; Cole and Kehoe 2000). Multiplicity poses a difficulty for policy analysis since the impact of policy on equilibrium outcomes is indeterminate. Morris and Shin (1998, 2001, 2003) show that a unique equilibrium often survives in these games when agents have heterogeneous information about the underlying payoff structure (the “fundamentals”). They then argue that policy analysis can be guided by the comparative statics of the unique equilibrium with respect to the policy instrument at hand.

However, this argument fails to take into account that policy choices often convey information that is relevant for the coordination game. Policy analysis then cannot be reduced to a simple comparative statics exercise.

Consider, for example, currency crises. A central bank can try to defend the peg against a speculative attack by borrowing reserves from abroad, by raising domestic interest rates, by taxing capital outflows, or by taking other measures that increase the speculators’ cost of attacking. However, such costly policy interventions may signal that the bank is anxious to escape an attack that would otherwise trigger devaluation. Intervention may thus backfire by reducing the speculators’ uncertainty about the success of a coordinated attack.

In this paper we examine the informational effects of policy interventions in coordination environments. The building block of our analysis is a *global coordination game of regime change*. A large number of small agents are choosing whether to attack the status quo, and a policy maker is defending it. The status quo is maintained as long as the aggregate attack is small enough. The policy maker’s type—her willingness or ability to maintain the status quo—is not common knowledge among the agents. Instead, agents observe noisy private signals about it.

This game is the backbone of recent applications of the global games methodology. In Morris and Shin (1998), for example, the status quo is a currency peg, the agents are speculators deciding whether to attack the domestic currency, and the policy maker is a central bank defending it.¹

To this game we add an earlier stage at which the policy maker controls a policy instrument that affects the agents' payoff from attacking. Since policy choices depend on the type of the policy maker and are observed by the agents before they decide whether to attack, this introduces *signaling* to the global coordination game.

If the policy did not convey any information, the equilibrium would be unique, as in Morris and Shin's article. Our main result is that the informational role of policy leads to multiple equilibria.

There is an *inactive-policy equilibrium* in which the agents coordinate on a strategy that is insensitive to the policy, thus ensuring that the policy maker finds it optimal to choose the cost-minimizing policy no matter her type. In addition, there is a continuum of *active-policy equilibria* in which the agents play less aggressively (i.e., attack for a lower range of their private signals) when the policy is raised sufficiently high. The policy maker then finds it optimal to raise the policy only for intermediate types; for sufficiently low types the value of maintaining the status quo is too low to justify the cost of policy intervention, whereas for sufficiently high types the size of attack is small enough even at the cost-minimizing policy. Intervention thus reveals that the policy maker is of intermediate type, which in turn permits the agents to coordinate on multiple reactions. Different equilibria then correspond to different levels of the policy at which the agents switch to a reaction that is favorable to the policy maker (i.e., to a smaller attack).

Our multiplicity result originates in the combination of signaling and coordination. Indeed, letting noise vanish would select a unique equilibrium outcome either if there were a single large agent interacting with the policy maker (no coordination) or if policy choices did not convey information (no signaling). The signaling role of policy is therefore essential for our result, but the freedom to choose out-of-equilibrium beliefs is not. All our equilibria survive the intuitive criterion of Cho and Kreps (1987) and, most important, can be obtained as the limit of equilibria in a perturbation of the game in which beliefs are always pinned down by Bayes' rule.

¹ Other applications include bank runs (Rochet and Vives 2004; Goldstein and Pauzner 2005), debt crises (Corsetti, Guimaraes, and Roubini 2006), investment (Chamley 1999; Dasgupta, forthcoming), liquidity crashes (Morris and Shin 2004), and sociopolitical change (Atkeson 2001; Edmond 2005). See Carlsson and van Damme (1993) for the pioneering contribution on global games and Morris and Shin (2003) for recent developments.

In the benchmark model, the policy is common knowledge and serves as a public signal about the policy maker's type. Moreover, in any of the active-policy equilibria, the equilibrium policy generates common certainty on whether the type of the policy maker is intermediate or extreme. However, our multiplicity result does not rely on either of these properties. It extends to variants of the benchmark model in which the policy is observed with idiosyncratic noise. Although agents receive noisy private signals about the policy, they can still coordinate on different self-fulfilling interpretations of, and reactions to, the same policy signals. As in the benchmark model, different reactions then sustain different strategies for the policy maker, which in turn justify different interpretations of the same policy signals.

On the theoretical side, the paper highlights the importance of endogenous information in global coordination games: while exogenous private information often leads to a unique equilibrium, endogenous information, *even if asymmetric*, sustains multiple equilibria in our framework. Our multiplicity result is thus clearly different from that in global games with exogenous public signals (Morris and Shin 2001; Hellwig 2002). It is also different from that in Angeletos and Werning (2006) or Hellwig, Mukherji, and Tsyvinski (2006), where the public signal—a financial price—is endogenous but its informational content does not vary across equilibria. Here, instead, different equilibrium policies lead to different information structures in the coordination game: multiplicity emerges in the information structure itself.

On the applied side, the paper raises questions about the merits of certain policy proposals that ignore the information content of policy choices. Can a central bank prevent a financial crisis by injecting liquidity, or will such intervention be interpreted as a signal of distress? And do interventions by the International Monetary Fund (IMF) ease debt crises, or might they do more harm than good by revealing that country fundamentals are weak enough to require IMF aid? Our results suggest that policy makers may find themselves trapped in a position in which the effectiveness of such interventions is dictated by self-fulfilling expectations—a form of *policy trap* that contrasts with the view that policy can fashion market behavior.

The rest of the paper is organized as follows. Section II introduces the model. Section III presents the main results. Section IV examines robustness to idiosyncratic noise in the observation of the policy. Section V concludes with a discussion of alternative payoff structures. All proofs are confined to Appendix A.

II. The Model

Actions and payoffs.—There are two possible regimes, the status quo and an alternative. A continuum of agents of measure one, indexed by i and uniformly distributed over $[0, 1]$, is choosing whether to attack (i.e., take an action that favors regime change) or abstain from attacking (i.e., take an action that favors the status quo). Let $r \in [\underline{r}, \bar{r}] \subset (0, 1)$ denote the opportunity cost of attacking and $D \in \{0, 1\}$ the regime outcome; both r and D are controlled by a policy maker.

The payoff for an agent who does not attack is normalized to zero, whereas the payoff from attacking is $1 - r$ in the event that the status quo is abandoned ($D = 1$) and $-r$ otherwise ($D = 0$). The payoff for the policy maker, on the other hand, has two components: the cost of policy intervention and the net benefit from maintaining the status quo. The former is $C(r)$, where C is strictly increasing and Lipschitz continuous, with $C(\underline{r}) = 0$. The latter is $\theta - A$, where $\theta \in \mathbb{R}$ is the type of the policy maker (the “fundamentals”) and $A \in [0, 1]$ is the mass of agents attacking. Hence, the policy maker’s payoff is $U = (1 - D)(\theta - A) - C(r)$.

Timing and information.—The game has three stages. In stage 1, the policy maker learns θ and sets r . In stage 2, agents decide simultaneously whether to attack after observing the policy r and private signals about θ . In stage 3, the policy maker observes the aggregate attack A and decides whether to maintain the status quo.

Since the policy maker finds it sequentially optimal to abandon the status quo when $A > \theta$ and maintain it when $A < \theta$, multiple equilibria exist when it is common knowledge that $\theta \in [0, 1]$. Following Morris and Shin (1998, 2001, 2003), we instead assume that agents have incomplete information about θ . The initial common prior is an improper uniform over \mathbb{R} .² The signal agent i receives is $x_i = \theta + \sigma \xi_i$, where $\sigma > 0$ parameterizes the quality of private information and ξ_i is idiosyncratic noise, independently and identically distributed (i.i.d.) across agents and independent of θ , with absolutely continuous cumulative distribution function (cdf) Ψ and density ψ . We allow the support of the noise to be either $[-1, +1]$ (bounded) or the entire real line (unbounded) and denote with $\Theta(x) \equiv \{\theta : \psi((x - \theta)/\sigma) > 0\}$ the set of types that are in the support of the beliefs of an agent with signal x . Bounded noise has the advantage that the freedom to choose out-of-equilibrium beliefs vanishes in the limit as $\sigma \rightarrow 0$, for then $\Theta(x) \rightarrow \{x\}$. On the other hand,

² Our multiplicity results hold for any smooth prior over any interval $\Theta = [\underline{\theta}, \bar{\theta}]$, with $\underline{\theta} < 0 < 1 < \bar{\theta}$. By assuming an uninformative prior, we bias the analysis *against* multiplicity. See Morris and Shin (2003) for a discussion of the role of priors in global coordination games.

with unbounded noise, $\Theta(x) = \mathbb{R}$, which captures the possibility that the agents may not be able to rule out any type.

Equilibrium.—The defining elements of the model are the signaling game in stage 1 and the coordination game in stage 2; that stage 3 is strategic is not essential. Indeed, as we discuss in Section V, the results extend to more general environments in which the regime outcome is not a choice of the policy maker. Thus we henceforth suppress stage 3 by setting $D = I_{[A > \theta]}$, where $I_{[A > \theta]}$ is the indicator function assuming value one if $A > \theta$ and zero otherwise.³ The policy maker’s payoff then reduces to $U(\theta, r, A) = \max\{0, \theta - A\} - C(r)$. For future reference, we also define $\underline{x} = -\sigma$ and $\bar{x} = 1 + \sigma$ if ξ is bounded, $\underline{x} = -\infty$ and $\bar{x} = +\infty$ if ξ is unbounded, $\tilde{\theta} = 1 - \underline{r} \in (0, 1)$, and $\tilde{r} = C^{-1}(\tilde{\theta}) \in [\underline{r}, \bar{r}]$.⁴

We consider symmetric perfect Bayesian equilibria. Let $r(\theta)$ denote the policy chosen by type θ , $a(x, r)$ the action of an agent given x and r , and $A(\theta, r)$ the aggregate size of attack, given θ and r . Finally, let $\mu(\theta \in S|x, r)$ denote the posterior probability that $\theta \in S$ (for some $S \subseteq \mathbb{R}$) conditional on x and r ; to save on notation, also let $\mu(\theta'|x, r) = \mu(\theta < \theta'|x, r)$.

DEFINITION. An equilibrium consists of a strategy for the policy maker $r: \mathbb{R} \rightarrow [\underline{r}, \bar{r}]$, a (symmetric) strategy for the agents $a: \mathbb{R} \times [\underline{r}, \bar{r}] \rightarrow \{0, 1\}$, and a cdf $\mu: \mathbb{R} \times \mathbb{R} \times [\underline{r}, \bar{r}] \rightarrow [0, 1]$ such that

- a. $r(\theta) \in \arg \max_{r \in [\underline{r}, \bar{r}]} U(\theta, r, A(\theta, r))$,
- b. $a(x, r) \in \arg \max_{a \in \{0, 1\}} a \{ \int_{\Theta(x)} I_{[A(\theta, r) > \theta]} d\mu(\theta|x, r) - r \}$,
- c. $\mu(\theta|x, r)$ is obtained from $r(\cdot)$ using Bayes’ rule for any $r \in r(\Theta(x))$,

where

$$A(\theta, r) \equiv \int_{-\infty}^{+\infty} a(x, r) \frac{1}{\sigma} \psi\left(\frac{x - \theta}{\sigma}\right) dx,$$

and $r(\Theta(x)) \equiv \{r: r = r(\theta) \text{ for some } \theta \in \Theta(x)\}$. For any θ , the equilibrium regime outcome $D(\theta) \equiv I_{[A(\theta, r(\theta)) > \theta]}$.

Conditions *a* and *b* require that the policy choice in stage 1 and the agents’ strategy in stage 2 be sequentially rational, whereas condition *c* requires that beliefs be pinned down by Bayes’ rule along the equilibrium path. We naturally impose that out-of-equilibrium beliefs assign positive measure only to $\theta \in \Theta(x)$. We further require that the equilibrium satisfies the *intuitive criterion* (Cho and Kreps 1987), namely, that it does not rely on out-of-equilibrium beliefs assigning positive measure to types for whom the deviation is dominated in equilibrium.

³ To simplify the exposition, and unless otherwise stated, we assume that the policy maker abandons the status quo, and the agents do not attack, when indifferent.

⁴ Letting $\tilde{r} = C^{-1}(\tilde{\theta})$ presumes $C(\tilde{r}) \geq \tilde{\theta}$. If the latter is not satisfied, the results hold with $\tilde{r} = \bar{r}$.

III. Policy Traps

When the policy is exogenously fixed, say $r = \underline{r} = \bar{r}$, the game reduces to a standard global game with exogenous information, as in Morris and Shin (1998, 2003).

PROPOSITION 1. Suppose $r = \underline{r} = \bar{r}$. The equilibrium is unique and is such that

$$a(x, r) = \begin{cases} 1 & \text{if } x < \tilde{x} \\ 0 & \text{otherwise,} \end{cases} \quad D(\theta) = \begin{cases} 1 & \text{if } \theta < \tilde{\theta} \\ 0 & \text{otherwise,} \end{cases}$$

where

$$\tilde{x} \equiv 1 - \underline{r} + \sigma \Psi^{-1}(1 - \underline{r}), \quad \tilde{\theta} \equiv 1 - \underline{r}. \quad (1)$$

This result follows directly from Morris and Shin (2003); here we merely sketch the proof. Suppose that an agent attacks if and only if $x < x^*$ for some $x^* \in \mathbb{R}$. The aggregate attack is $A(\theta, r) = \Psi((x^* - \theta)/\sigma)$, implying that the status quo is abandoned if and only if $\theta < \theta^*$, where θ^* solves $\theta^* = \Psi((x^* - \theta^*)/\sigma)$. It follows that the posterior probability of regime change for an agent with signal x is $\mu(\theta^*|x) = 1 - \Psi((x - \theta^*)/\sigma)$, and hence x^* must solve $r = 1 - \Psi((x^* - \theta^*)/\sigma)$. Combining the two conditions gives $x^* = \tilde{x}$ and $\theta^* = \tilde{\theta}$. Clearly, this is the sole monotone equilibrium; iterated deletion of strictly dominated strategies then selects this as the unique equilibrium of the game.

Proposition 1 implies that in an economy in which the cost of attacking is higher, the range of θ for which the status quo is abandoned is smaller. This seems to suggest that a policy maker who can control r can use this policy instrument to reduce the size of attack and hence the range of fundamentals for which regime change occurs.

However, this line of reasoning neglects that policy choices convey information about θ . For example, consider the case with bounded noise. Since agents find it dominant to attack when $x < \underline{x}$, for $\theta < \underline{x} - \sigma$, the policy maker expects every agent to attack no matter what r is and hence finds it dominant to set \underline{r} . Similarly, since agents find it dominant not to attack when $x > \bar{x}$, the policy maker necessarily sets \underline{r} also for $\theta > \bar{x} + \sigma$. Any policy intervention thus signals that θ is neither too low nor too high—information that may interfere with the ability of the policy maker to control equilibrium outcomes.

Our main result is that when policy conveys information, this leads to multiple equilibria. Multiplicity is sustained by the agents coordinating on different interpretations of, and different reactions to, the same policy choices, thus inducing different incentives for the policy maker.

If agents expect no policy intervention and follow a strategy $a(x, r)$

that is insensitive to r , the policy maker has no incentive to intervene. If, instead, agents expect the policy maker to intervene for some θ and coordinate on a strategy $a(x, r)$ that is decreasing not only in x but also in r , then the policy maker can decrease the size of the attack by raising r . However, the cost of intervention may exceed the value of maintaining the status quo for low θ ; similarly, the attack faced when setting \underline{r} may be too small to justify the cost of intervention when θ is high. This reasoning suggests the existence of *active-policy equilibria* in which intervention occurs only for intermediate θ , along with *inactive-policy equilibria* in which intervention never occurs.

PROPOSITION 2. Suppose $\underline{r} < \tilde{r}$. There are multiple equilibria.

a. There is an equilibrium in which $r(\theta) = \underline{r}$ for all θ ,

$$a(x, r) = \begin{cases} 1 & \text{if } x < \tilde{x} \\ 0 & \text{otherwise,} \end{cases} \quad D(\theta) = \begin{cases} 1 & \text{if } \theta < \tilde{\theta} \\ 0 & \text{otherwise,} \end{cases}$$

where \tilde{x} and $\tilde{\theta}$ are as in (1).

b. For any $r^* \in (\underline{r}, \tilde{r}]$, there is an equilibrium in which

$$r(\theta) = \begin{cases} r^* & \text{if } \theta \in [\theta^*, \theta^{**}] \\ \underline{r} & \text{otherwise,} \end{cases}$$

$$a(x, r) = \begin{cases} 1 & \text{if } x < \underline{x} \text{ or } (x, r) < (x^*, r^*) \\ 0 & \text{otherwise,} \end{cases}$$

$$D(\theta) = \begin{cases} 1 & \text{if } \theta < \theta^* \\ 0 & \text{otherwise,} \end{cases}$$

where θ^* , θ^{**} , and x^* are given by

$$\theta^* = C(r^*),$$

$$\theta^{**} = \theta^* + \sigma \left[\Psi^{-1} \left(1 - \frac{\underline{r}}{1 - \underline{r}} \theta^* \right) - \Psi^{-1}(\theta^*) \right],$$

$$x^* = \theta^{**} + \sigma \Psi^{-1}(\theta^*). \tag{2}$$

The construction of an inactive-policy equilibrium (part a above) is illustrated in figure 1 for the case in which ξ is unbounded. Given the agents' strategy, the size of attack $A(\theta, r)$ is decreasing in θ and independent of r . The status quo is thus abandoned if and only if $\theta < \tilde{\theta}$, where $\tilde{\theta}$ solves $\tilde{\theta} = A(\tilde{\theta}, r)$. When $r = \underline{r}$, beliefs are pinned down by Bayes' rule and satisfy $\mu(\theta | x, \underline{r}) > \underline{r}$ if and only if $x < \tilde{x}$. For $r \neq \underline{r}$, we consider out-of-equilibrium beliefs that assign zero measure to types for whom the deviation is dominated in equilibrium (whenever possible). Since the policy maker's equilibrium payoff is zero for $\theta \leq \theta$ and




FIG. 1.—Inactive-policy equilibrium

$\theta - A(\theta, \underline{r}) > 0$ for $\theta > \tilde{\theta}$, any $r > \tilde{r}$ is dominated in equilibrium for all θ , in which case the only restriction is that beliefs have support $\Theta(x)$. A deviation to some $r' \in (\underline{r}, \tilde{r})$, on the other hand, is dominated if and only if $\theta \notin [\theta', \theta'']$, where θ' and θ'' solve $\theta' = C(r') = A(\theta'', \underline{r})$: for $\theta < \theta'$, the cost of r' exceeds the value of maintaining the status quo; for $\theta > \theta''$, the attack faced in equilibrium is smaller than the cost of r' . Since $\tilde{\theta} \in [\theta', \theta'']$, we can construct beliefs that assign zero measure to $\theta \notin [\theta', \theta'']$ whenever $[\theta', \theta''] \cap \Theta(x) \neq \emptyset$ and at the same time satisfy $\mu(\theta|x, r) > r$ if and only if $x < \tilde{x}$, for all r . Given these beliefs, an agent who expects all other agents to follow the proposed strategy, thus triggering regime change if and only if $\theta < \theta$, finds it optimal to do the same. Since the size of the attack $A(\theta, r)$ is independent of r , the optimal policy is then clearly $r(\theta) = \underline{r}$ for all θ .

The construction of an active-policy equilibrium (part *b* above) is illustrated in figure 2. When the agents coordinate on the equilibrium strategy, it never pays to raise the policy at any $r \neq r^*$. Furthermore, r^* is preferred to \underline{r} if and only if the cost of intervention, $C(r^*)$, is less than the opportunity cost of setting \underline{r} , which is equal to the value of maintaining the status quo for types that cannot withstand the attack $A(\theta, \underline{r})$ and to the cost of attack for types that can overcome it. Let θ^* and θ^{**} solve, respectively, $C(r^*) = \theta^*$ and $C(r^*) = A(\theta^{**}, \underline{r})$. As illustrated in figure 2, the interval $[\theta^*, \theta^{**}]$ is the set of types for whom $C(r^*) \leq \min\{\theta, A(\theta, \underline{r})\}$, that is, the types for whom r^* is preferred to \underline{r} . Combining the two indifference conditions for θ^* and θ^{**} with that for the marginal agent, namely with $\underline{r} = \mu(\theta^*|x^*, \underline{r})$, leads to condition (2), while ensuring that the set of types that intervene is nonempty puts an




FIG. 2.—Active-policy equilibrium

upper bound on the cost of intervention and hence on r^* . Finally, because a deviation to $r' > r^*$ is dominated in equilibrium for all types, whereas a deviation to $r' \in (\underline{r}, r^*)$ is not dominated for intermediate types (namely for $\theta \in [\theta', \theta'']$, where θ' and θ'' solve $\theta' = C(r^*) = A(\theta', \underline{r})$ and hence satisfy $0 < \theta' < \theta^* \leq \theta^{**} < \theta''$), there exist out-of-equilibrium beliefs that sustain the proposed strategy for the agents while at the same time attaching zero probability to types for whom a deviation is dominated in equilibrium.

Our result underscores the informational effects of policy interventions in coordination environments. Indeed, multiplicity originates in the combination of signaling and coordination. First, if the policy did not convey any information, as in the case in which r is fixed before the policy maker learns θ , then the equilibrium would be unique. Second, if a single “big” agent were to play against the policy maker, this agent would recognize that he can trigger regime change for all $\theta < 1$. When noise is bounded, he would necessarily attack for $x < 1 - \sigma$ irrespective of r . As $\sigma \rightarrow 0$, the status quo would then be abandoned if and only if $\theta < 1$. In our coordination setting, instead, the regime outcome remains indeterminate for all $\theta \in (0, \theta]$, no matter what σ is.⁵

Note that the exact θ is never revealed in any of the active-policy equilibria of part *b*: the equilibrium policy only makes it common

⁵ By (2) we have that, for any given $r^* \in (\underline{r}, \bar{r})$, while θ^* is independent of σ , θ^{**} decreases with σ and converges to θ^* as $\sigma \rightarrow 0$. Hence, with regard to the policy, as $\sigma \rightarrow 0$, the active-policy equilibria of part *b* become observationally equivalent to one another, as well as to the inactive-policy equilibrium of part *a*; however, they remain very different with regard to the regime outcome.

certainty (i.e., common $p = 1$ belief) whether $\theta \in [\theta^*, \theta^{**}]$ or $\theta \notin [\theta^*, \theta^{**}]$. Common certainty that $\theta \in [\theta^*, \theta^{**}]$ in turn permits *all* agents to coordinate on not attacking. However, as discussed in the next section, this form of common certainty and the property that no agent attacks when the policy maker intervenes are not essential for multiplicity.

Also note that there exist other equilibria that sustain the same policy and regime outcomes as in proposition 2. For example, there is an inactive-policy equilibrium in which the policy maker sets $r = \underline{r}$ for all θ and agents attack if and only if $x < \bar{x}$ when $r = \underline{r}$, and if and only if $x < \bar{x}$ when $r > \underline{r}$. This equilibrium features more “aggressive” reactions to policy deviations than those in part *a* above and is sustained by out-of-equilibrium beliefs that assign probability one to $\theta < 1$ whenever $x < \bar{x}$. Similarly, for any $r^* \in (\underline{r}, \bar{r}]$, there is an active-policy equilibrium in which agents coordinate on attacking for any $x < \bar{x}$ whenever $r \notin \{\underline{r}, r^*\}$, inducing the policy maker to play the same strategy as in part *b* of proposition 2.

The beliefs and strategies we consider in the proof of proposition 2 have two advantages. First, the beliefs satisfy a simple forward-induction argument as in Cho and Kreps’s (1987) intuitive criterion (see App. A). Second, the strategies are the limit of equilibrium strategies in a perturbation of the game in which r is observed with unbounded idiosyncratic noise and there is no room for out-of-equilibrium beliefs (see Sec. IV.B).

Finally, note that proposition 2 does not exhaust the set of possible equilibrium outcomes. For example, when the noise ξ is unbounded, for any $r^* \in (\underline{r}, \bar{r}]$, there also exist equilibria in which the policy is raised to r^* for all $\theta \geq \theta^*$, where $\theta^* = C(r^*)$. These equilibria are sustained by the agents coordinating on attacking if and only if $r < r^*$ no matter what x is and therefore exist only if agents do not find it dominant to refrain from attacking for sufficiently high x , which in turn is possible only if noise is unbounded (Angeletos, Hellwig, and Pavan 2003). See also Section V for a discussion of how alternative payoff structures can lead to different equilibrium outcomes.

IV. Idiosyncratic Policy Observation

The payoff structure of the coordination game played among the agents depends on two variables, θ and r . We have assumed that, while θ is observed with idiosyncratic noise, r is observed publicly. Although this is a reasonable assumption for many applications, from a global games

perspective it is important to consider perturbations that remove common knowledge of the policy.⁶

In this section, we consider two such perturbations. In the first, the policy is observed with small bounded idiosyncratic noise; in the second, the support of the policy signals does not shift with the actual policy choice. In both cases, there is no public information about either r or θ ; in the second, there is also no room for out-of-equilibrium beliefs.

The key difference from standard global games is the endogeneity of the information structure: whereas the informational content of the private signals about θ is exogenously given, the informational content of the private signals about r is determined in equilibrium. As we show next, this has important implications for the determinacy of equilibria.

A. *Bounded Policy Noise*

Consider the following modification of the benchmark model. The policy choice r is no longer publicly observed. Instead, each agent receives a private signal $z_i = r + \eta \zeta_i$ about the policy, where $\eta > 0$ parameterizes the precision of the policy signal and ζ_i is noise, i.i.d. across agents, independent of θ and of ξ_j for any j , and distributed over $[-1, 1]$ with absolutely continuous cdf Φ and density ϕ bounded away from zero.

This extension requires a modification of the equilibrium definition: agents now condition their strategy on (x, z) and form beliefs not only about θ but also about r . We denote with $\mu(\{\theta \in S, r \in S'\} | x, z)$ the posterior probability that an agent with signals x and z assigns to the joint event that $\theta \in S$ and $r \in S'$ (where $S \subseteq \mathbb{R}$ and $S' \subseteq [\underline{r}, \bar{r}]$), and with $R(z) \equiv [z - \eta, z + \eta] \cap [\underline{r}, \bar{r}]$ the set of policy choices that are consistent with the observation of a policy signal z . An agent who observes (x, z) such that $R(z) \cap r(\Theta(x)) \neq \emptyset$ necessarily believes that no deviation has occurred, in which case his beliefs are pinned down by Bayes' rule. On the contrary, the observation of (x, z) such that $R(z) \cap r(\Theta(x)) = \emptyset$ signals that a deviation occurred, in which case the agent has to form beliefs not only about θ but also about which policy was chosen. As in the previous section, we naturally impose that out-of-equilibrium beliefs assign zero measure to pairs (θ, r) that are inconsistent with the signals (x, z) , that is, to pairs such that either $\theta \notin \Theta(x)$ or $r \notin R(z)$. In the spirit of the intuitive criterion, we further require that the equilibrium does not rely on out-of-equilibrium beliefs assigning positive measure to pairs (θ, r) such that choice r is dominated in equilibrium for type θ .

Since there is *no* public information about either θ or r , one may

⁶ Another possibility is that r is observed with aggregate noise. This case may be relevant for some applications but is less interesting from a theoretical perspective, for it maintains common knowledge of the policy signal.

expect a unique equilibrium. However, the fact that equilibrium policies can be nonmonotone and discontinuous—as in the benchmark model—can generate common certainty about θ and r despite the absence of public information.

PROPOSITION 3. Consider the game with bounded policy noise as described above. There exists $\bar{\eta} > 0$ such that, for any $\eta < \bar{\eta}$, there exist multiple equilibria.

- a. There is an equilibrium in which $r(\theta) = \underline{r}$ for all θ ,

$$a(x, z) = \begin{cases} 1 & \text{if } x < \tilde{x} \\ 0 & \text{otherwise,} \end{cases} \quad D(\theta) = \begin{cases} 1 & \text{if } \theta < \tilde{\theta} \\ 0 & \text{otherwise,} \end{cases}$$

where \tilde{x} and $\tilde{\theta}$ are as in (1).

- b. There exists $\underline{r}_\eta \in (\underline{r}, \tilde{r})$ such that, for any $r^* \in (\underline{r}_\eta, \tilde{r}]$, there is an equilibrium in which

$$r(\theta) = \begin{cases} r^* & \text{if } \theta \in [\theta^*, \theta^{**}] \\ \underline{r} & \text{otherwise,} \end{cases}$$

$$a(x, z) = \begin{cases} 1 & \text{if } x < \underline{x} \text{ or } (x, z) < (x^*, z^*) \\ 0 & \text{otherwise,} \end{cases}$$

$$D(\theta) = \begin{cases} 1 & \text{if } \theta < \theta^* \\ 0 & \text{otherwise,} \end{cases}$$

where θ^* and θ^{**} are as in (2) and $z^* = r^* - \eta$. Moreover, $\underline{r}_\eta \rightarrow \underline{r}$ as $\eta \rightarrow 0$.

In the inactive-policy equilibrium, agents attack if and only if $x < \tilde{x}$, no matter their (idiosyncratic) policy signal z . Given this strategy, the aggregate size of attack and the regime outcome are insensitive to r , in which case it is clearly optimal for the policy maker to set $r(\theta) = \underline{r}$ for all θ .

For the agents, then, receiving a signal $z \in [\underline{r} - \eta, \underline{r} + \eta]$ is consistent with equilibrium play by the policy maker and is completely uninformative about θ . Sequential rationality of the agents' equilibrium strategy then follows from exactly the same arguments as in the benchmark model. Receiving a signal $z > \underline{r} + \eta$, on the other hand, is informative of a deviation. Sequential rationality would then be trivial if out-of-equilibrium beliefs were free. For example, it would suffice that an agent believes with probability one that $\underline{r} = \max R(z)$ and $\theta < \tilde{\theta}$ whenever $x < \tilde{x}$ and that $\underline{r} = \max R(z)$ and $\theta > \tilde{\theta}$ whenever $x > \tilde{x}$. However, for some (x, z) , these beliefs assign measure one to pairs (θ, r) such that r is dominated in equilibrium for θ , despite the fact that there exist pairs

(θ', r') such that r' is not dominated for θ' .⁷ If the equilibrium were to rely on this property, it would fail the intuitive criterion test. The beliefs we construct in Appendix A guarantee that this is not the case.

In an active-policy equilibrium, on the other hand, agents coordinate on a strategy such that, for any θ , $A(\theta, r) = A(\theta, \underline{x})$ for $r \in [\underline{x}, r^* - 2\eta]$, $A(\theta, r)$ is strictly decreasing in r for $r \in (r^* - 2\eta, r^*)$, and $A(\theta, r) = A(\theta, r^*)$ for any $r \geq r^*$. It follows that, as in the benchmark model, \underline{x} is preferred to any $r \in (\underline{x}, r^* - 2\eta]$, and r^* is preferred to any $r > r^*$. But, in contrast to the benchmark model, it no longer suffices to check whether $\theta \in [\theta^*, \theta^{**}]$ prefers r^* to \underline{x} (and the converse for $\theta \notin [\theta^*, \theta^{**}]$). It is also needed that any $\theta \in [\theta^*, \theta^{**}]$ does not find it profitable to deviate from r^* to any $r \in (r^* - 2\eta, r^*)$, in which case only a fraction of the agents would detect the deviation and switch to more “aggressive” behavior (i.e., from attacking for $x < \underline{x}$ to attacking for $x < x^*$). In Appendix A we show that, while the marginal cost saving of reducing r below r^* is independent of η , the marginal increase in the size of the attack goes to infinity as $\eta \rightarrow 0$. It follows that η small enough suffices for such a deviation to be unprofitable and hence for r^* to be optimal for $\theta \in [\theta^*, \theta^{**}]$. The same bound on η also guarantees that any $\theta \notin [\theta^*, \theta^{**}]$ prefers \underline{x} to any $r \in [r^* - 2\eta, r^*]$.

Consider next the agents. Although the policy signals are private, in equilibrium the observation of $z \in [r^* - \eta, r^* + \eta]$ generates common certainty that $r = r^*$ and hence that $\theta \in [\theta^*, \theta^{**}]$. Similarly, $z \in [\underline{x} - \eta, \underline{x} + \eta]$ generates common certainty that $r = \underline{x}$ and $\theta \notin [\theta^*, \theta^{**}]$. Sequential rationality then follows from the same arguments as in the benchmark model. On the other hand, an agent who detects a deviation needs to form beliefs about whether other agents also detected the deviation in order to best respond. If such an agent believed that *all* other agents detected the deviation, then strategic uncertainty about the dependence of other agents’ actions on z would vanish, and sequential rationality would follow from essentially the same arguments as in the benchmark model. However, this would again leave open the possibility that the equilibrium fails the intuitive criterion. In Appendix A we consider beliefs that guarantee that this is never the case.

B. Unbounded Policy Noise

In the active-policy equilibria of the previous section, although agents observe both r and θ with idiosyncratic noise, the equilibrium policy creates common certainty on whether $r = \underline{x}$ or $r = r^*$ and, correspondingly, on whether θ is extreme or intermediate. This is not a consequence

⁷ For example, this is the case when $z \in (\bar{r}, \bar{r} + \eta)$ and $\Theta(x) \ni \bar{\theta}$, since $r' = z - \eta$ is not dominated for $\theta' = \bar{\theta}$.

of bounded noise alone: if the equilibrium policy had no discontinuities, no policy choice would ever lead to common certainty about either r or θ . Nevertheless, one may argue that such common certainty is necessary for the multiplicity result. Moreover, with bounded noise, agents can still detect deviations. As in standard signaling games, one may thus argue that multiplicity relies on the freedom to choose out-of-equilibrium beliefs.

To show that neither of the above concerns is necessarily valid, we consider the following variant of the benchmark model. The private signal about the policy is $z_i = w_i \underline{r} + (1 - w_i)r + \eta \zeta_i$, where w_i is a binary variable assuming value one with probability $\rho \in [0, 1)$ and zero otherwise, whereas ζ_i is distributed exponentially over $[0, +\infty)$ and $\eta > 0$. The realizations of w_i and ζ_i are i.i.d. across agents, independent of each other, as well as of θ and of ξ_j , for any $j \in [0, 1]$. We also assume that C is linear and that ψ , the probability density function of the noise ξ , is log concave and strictly positive over \mathbb{R} .

The combination of these assumptions permits us to address the concerns raised above while keeping the analysis relatively tractable. In particular, that the noises ξ and ζ are unbounded along with the fact that \underline{r} is dominant for $\theta < 0$ ensures that any (x, z) is consistent with equilibrium, which leaves no room for out-of-equilibrium beliefs. That ζ is exponential, on the other hand, simplifies the construction of posterior beliefs. Finally, the binary variable w is not necessary, but it delivers the appealing property that the support of the policy signal z is the same for any policy choice r .

Since any (x, z) is consistent with any (θ, r) , no policy ever generates certainty—either common or private—about the fundamentals. Hence, in contrast to both the benchmark model and the bounded-noise case, agents can no longer coordinate on *everybody* not attacking whenever the policy maker raises the policy sufficiently high. Indeed, an agent necessarily attacks when x is low enough no matter what z is, and therefore the aggregate attack is bounded away from zero no matter what r is. Nevertheless, agents can still coordinate on multiple self-fulfilling interpretations of, and reactions to, the same idiosyncratic policy signals.

PROPOSITION 4. Consider the game with unbounded policy noise as described above. There exist $(\tilde{\eta}, \tilde{\rho}) > 0$ such that, for any $(\eta, \rho) < (\tilde{\eta}, \tilde{\rho})$, there exist multiple equilibria.

- a. There always exists an equilibrium in which $r(\theta) = \underline{r}$ for all θ ,

$$a(x, z) = \begin{cases} 1 & \text{if } x < \tilde{x} \\ 0 & \text{otherwise,} \end{cases} \quad D(\theta) = \begin{cases} 1 & \text{if } \theta < \tilde{\theta} \\ 0 & \text{otherwise,} \end{cases}$$

where \tilde{x} and $\tilde{\theta}$ are as in (1).

- b. For any $r^* \in (\underline{r}, \bar{r}]$ and any $\epsilon > 0$, there exist $(\hat{\eta}, \hat{\rho}) > 0$ such that the following is true: for any $(\eta, \rho) < (\hat{\eta}, \hat{\rho})$, there is an equilibrium in which

$$r(\theta) = \begin{cases} r^* & \text{if } \theta \in [\theta', \theta''] \\ \underline{r} & \text{otherwise,} \end{cases}$$

$$a(x, z) = \begin{cases} 1 & \text{if } x < \hat{x} \text{ or } (x, z) < (x', z^*) \\ 0 & \text{otherwise,} \end{cases}$$

$$D(\theta) = \begin{cases} 1 & \text{if } \theta < \theta' \\ 0 & \text{otherwise,} \end{cases}$$

where $\theta', \theta'', x', \hat{x}$, and z^* satisfy $0 < \theta' \leq \theta'', |\theta' - \theta^*| < \epsilon$, $|\theta'' - \theta^{**}| < \epsilon$, $|x' - x^*| < \epsilon$, $\hat{x} < -1/\epsilon$, and $z^* = r^*$, with θ^*, θ^{**} , and x^* as in (2).

In an inactive-policy equilibrium, agents expect the policy maker never to intervene and hence interpret variation in z as pure noise. They then condition their behavior only on x , thus making the aggregate attack independent of r . It is then immediate that an inactive-policy equilibrium always exists.⁸

In an active-policy equilibrium, on the other hand, the policy maker raises the policy for intermediate θ . By implication, the private signals z are now informative: an agent who observes a high z attaches high probability not only to θ being intermediate but also to other agents observing high policy signals. If he then expects them to play more “leniently” for sufficiently high z (i.e., to use a lower cutoff for x , as in the equilibrium strategy prescribed above), he finds it optimal to do the same. When in turn all agents follow such a strategy, the policy maker finds it optimal to intervene only for intermediate θ , which in turn justifies the agents’ interpretation of high z as a signal of intermediate types.

This circularity is suggestive of multiplicity, but, to complete the argument, one needs to close the cycle; the payoff and signal structures assumed here make this task possible. In particular, the exponential noise structure ensures that, for any given x , the likelihood ratio of the event that $r = \underline{r}$ relative to the event that $r = r^*$ takes only two values as a function of z : a high one for $z < z^*$ ($= r^*$) and a low one for $z \geq z^*$. Combined with the fact that all θ who abandon the regime set \underline{r} , this implies that, for any given x , the posterior probability of regime change, $\mu(\theta < \theta' | x, z)$, also takes only two values as a function of z and switches

⁸ Indeed, the existence of an inactive-policy equilibrium trivially extends to any noise structure that ensures nonmoving support for the policy signal.

from the higher one to the lower one as z crosses z^* . This in turn implies that, for any given x , the strategy of the agents is a step function of z , with a downward discontinuity at $z = z^*$. The aggregate attack $A(\theta, r)$ is then continuous in r but has a kink at $r = r^*$. As the noise in z vanishes, this kink becomes arbitrarily steep. It follows that (η, ρ) small enough suffices for the marginal cost of reducing r below r^* (i.e., the cost of a higher A) to outweigh the associated marginal benefit (i.e., the benefit of a lower C). This in turn allows one to sustain a discontinuous equilibrium policy as in the benchmark model.

Note that this policy discontinuity is different from the state-space discontinuities that are known to break uniqueness arguments in standard global games. First, the discontinuity here is an equilibrium property, not an assumption on the state space. Second, the state that is relevant for dominance is merely θ , not (θ, r) . Indeed, our results do not rely on r having a direct payoff effect for the agents: policy could have simply been “money burning.” Finally, even if the coordination game played among the agents for a fixed strategy of the policy maker had a unique equilibrium, multiplicity could still emerge in the fixed point between the strategy of the policy maker and the interpretations and reactions of the agents.

V. Alternative Payoff Structures

The model we use in this paper is highly stylized. However, we expect the multiplicity result to extend more generally to environments that share the main two ingredients of our model, signaling and coordination.

Consider, for example, the following extension of the benchmark model. The status quo is maintained ($D = 0$) if and only if $R(\theta, A, r) \geq 0$, where R is continuous in (θ, A, r) , strictly increasing in θ , strictly decreasing in A , and nondecreasing in r , with $R(0, 0, r) = 0 = R(1, 1, r)$ for any r . When $D = 0$, the policy maker’s payoff is given by $U(\theta, A, r)$, where U is continuous in (θ, A, r) , nondecreasing in θ , nonincreasing in A , and strictly decreasing in r . When instead $D = 1$, her payoff is zero if $r = \underline{r}$ and $W(\theta, A, r) < 0$ if $r > \underline{r}$.⁹ The agent’s net payoff from attacking is $D - g(r)$, where g is nondecreasing in r , with $0 < g(\underline{r}) \leq g(\bar{r}) < 1$. The benchmark model is nested with $R(\theta, A, r) = \theta - A$, $U(\theta, A, r) = \max\{0, \theta - A\} - C(r)$, and $g(r) = r$. Here we assume that ψ is log concave and let $\tilde{\theta} \in (0, 1)$ and $\tilde{r} \in (\underline{r}, \bar{r})$ solve $R(\tilde{\theta}, 1 - g(\underline{r}), \underline{r}) = 0$ and $U(\tilde{\theta}, 0, \tilde{r}) = 0$.

This extension captures aspects that are absent in the benchmark

⁹ Note that (i) r does not affect the region of θ for which the regime outcome depends on A , and (ii) \underline{r} is optimal when the status quo is abandoned.

model but may be important for applications. The cost of policy intervention can depend on the policy maker's type, the regime outcome can be beyond her direct control, and the strength of the status quo can differ from the value she attaches to it. Furthermore, the strength of the status quo can be affected directly by the policy instrument, as in the case in which a central bank borrows reserves from abroad in anticipation of a speculative currency attack. Finally, policy intervention can be pure money burning—a costly action for the policy maker that has no effect on agents' payoffs—highlighting that what is most important for the multiplicity result is the signaling role of the policy, not its direct payoff effects.

The existence of an inactive-policy equilibrium as in part *a* of proposition 2 is straightforward; the following proposition generalizes part *b*.

PROPOSITION 5. Consider the extension described above. For any $r^* \in (\underline{r}, \tilde{r})$, there exists a nonempty set Θ^{**} , a function $X: [\underline{r}, \tilde{r}] \rightarrow \mathbb{R} \cup \{\pm\infty\}$, and an equilibrium in which

$$r(\theta) = \begin{cases} r^* & \text{if } \theta \in \Theta^{**} \\ \underline{r} & \text{otherwise,} \end{cases}$$

$$a(x, r) = \begin{cases} 1 & \text{if } x < X(r) \\ 0 & \text{otherwise,} \end{cases}$$

$$D(\theta) = \begin{cases} 1 & \text{if } \theta < \theta^* \\ 0 & \text{otherwise,} \end{cases}$$

where $\theta^* = \min \Theta^{**} \in [0, \tilde{\theta})$.

Unlike the active-policy equilibria of proposition 2, the set of fundamentals for which the policy maker intervenes may now be the union of multiple disjoint intervals. For example, if $R(\theta, A, r) = \theta - A$, $U(\theta, A, r) = \theta - A - C(r, \theta)$, and $C(r, \theta)$ is decreasing in θ for $r > \underline{r}$, the equation $C(r^*, \theta) = A(\theta, \underline{r})$ can admit multiple solutions corresponding to multiple indifference points at which the policy maker switches between r^* and \underline{r} . Furthermore, there are cases in which it is the lowest types in the critical region that intervene. For example, if $R(\theta, A, r) = \theta - A$ and $U(\theta, A, r) = V - C(r)$, where $V > 0$ is a constant, then $\Theta^{**} = [0, \theta^{**})$ for some $\theta^{**} > 0$.

Another possibility, not addressed by the above extension, is that the strength and the value of the status quo are negatively correlated with each other: countries with the weakest fundamentals might be those that suffer the most from a collapse of their currency or their banking system. We consider an example in online Appendix B (Sec. B). We find again multiple equilibria in which policy intervention occurs for the lowest types in the critical region.

We conclude that, although the equilibria of proposition 2 are clearly

sensitive to the specific payoff assumptions of the benchmark model, the multiplicity result may be relevant for a variety of applications that feature a combination of signaling and coordination.

VI. Concluding Remarks

This paper endogenized policy in a global coordination game. We found that the possibility that policy choices convey information leads to multiple equilibria in which both the optimal policy and the coordination outcome are dictated by self-fulfilling expectations. Different equilibria are sustained by the agents coordinating on different interpretations of, and different reactions to, the same policy choices.

On the theoretical side, our results underscore the importance of endogenous information structures in global games. On the applied side, they warn that policy proposals that ignore the signaling aspect of policy interventions can be dangerously misleading. Further exploring the interaction between policy making and market expectations remains an important direction for future research.

Appendix A

Proofs

Proof of Proposition 2

We first prove the existence of a pooling equilibrium in which $r(\theta) = \underline{r}$ for all θ ; we then construct a continuum of semiseparating equilibria in which $r(\theta) > \underline{r}$ for an intermediate interval of θ . In both cases, we consider beliefs that ensure that the equilibria pass the intuitive criterion.

Part *a*. When the agents play the equilibrium strategy, the aggregate size of attack does not depend on r . Since C is strictly increasing, any θ clearly finds it optimal to set \underline{r} , which establishes the optimality of the strategy of the policy maker.

Next, consider the agents. An agent who expects all other agents to follow the equilibrium strategy expects regime change to occur if and only if $\theta < \tilde{\theta}$, no matter what r is. He thus also finds it optimal to follow the equilibrium strategy if and only if his beliefs satisfy

$$\mu(\tilde{\theta}|x, r) > r \quad \text{if and only if } x < \tilde{x} \text{ for all } r. \quad (\text{A1})$$

Since in equilibrium $r(\theta) = \underline{r}$ for all θ , the observation of \underline{r} is uninformative about θ . Hence, when $r = \underline{r}$, beliefs are pinned down by Bayes' rule and are given by $\mu(\theta|x, \underline{r}) = 1 - \Psi((x - \theta)/\sigma)$; that they satisfy (A1) then follows directly from (1). When, instead, $r > \underline{r}$, a deviation is detected whatever x is. Then let $\bar{\Theta}(r)$ denote the set of types for whom the deviation to $r > \underline{r}$ is dominated in equilibrium and consider any beliefs that satisfy (A1) and the following two

conditions:

$$\mu(\theta \in \Theta(x)|x, r) = 1 \tag{A2}$$

and

$$\mu(\theta \in \bar{\Theta}(r)|x, r) = 0 \text{ whenever } \Theta(x) \not\subseteq \bar{\Theta}(r). \tag{A3}$$

Condition (A2) is the natural restriction that beliefs assign positive measure only to types that could have led to signal x , whereas (A3) requires that beliefs assign zero measure to types for whom the deviation is dominated in equilibrium. The latter suffices—although it is not necessary—for the equilibrium to satisfy the intuitive criterion.

To see that the set of beliefs that satisfy (A1)–(A3) for any $r > \underline{r}$ is nonempty, note that the most favorable reaction to any deviation is that an agent attacks only when $x < \underline{x}$ (i.e., unless it is dominant for him to do otherwise), and therefore

$$\bar{\Theta}(r) = \left\{ \theta : U^*(\theta) > U\left(\theta, r, \Psi\left(\frac{x-\theta}{\sigma}\right)\right) \right\},$$

where

$$U^*(\theta) \equiv U\left(\theta, \underline{r}, \Psi\left(\frac{\tilde{x}-\theta}{\sigma}\right)\right) = \max\left\{0, \theta - \Psi\left(\frac{\tilde{x}-\theta}{\sigma}\right)\right\}$$

is the policy maker’s equilibrium payoff. Since $C(\tilde{r}) = \tilde{\theta} = \Psi((\tilde{x} - \tilde{\theta})/\sigma)$, we have that $\bar{\Theta}(r) = \mathbb{R}$ for $r > \tilde{r}$ and $\bar{\Theta}(r) = \mathbb{R} \setminus \{\tilde{\theta}\}$ for $r = \tilde{r}$.¹⁰ Given that $\tilde{\theta} \in \Theta(\tilde{x})$, the set of beliefs that satisfy (A1)–(A3) for $r \geq \tilde{r}$ is clearly nonempty. On the other hand, a deviation to some $r' \in (\underline{r}, \tilde{r})$ is dominated in equilibrium if and only if $\theta \notin [\theta', \theta'']$, where θ' and θ'' solve $\theta' = C(r') = A(\theta', \underline{r})$, with $A(\theta, \underline{r}) = \Psi((\tilde{x} - \theta)/\sigma)$. (This is illustrated in fig. 1: the cost of raising the policy to r' exceeds the value of maintaining the status quo for $\theta < \theta'$, whereas it exceeds the cost of the attack faced in equilibrium for $\theta > \theta''$.) Since $\theta' < \tilde{\theta} < \theta''$ and $\tilde{\theta} \in \Theta(\tilde{x})$, the set of beliefs that assign zero measure to $\theta \notin [\theta', \theta'']$ whenever $[\theta', \theta''] \cap \Theta(x) \neq \emptyset$ and at the same time satisfy (A1) and (A2) is also nonempty. QED

Part *b*. Take any $r^* \in (\underline{r}, \tilde{r}]$, let $(\theta^*, \theta^{**}, x^*)$ be as in (2), and note that $0 < \theta^* \leq \tilde{\theta}$ and $\theta^* \leq \theta^{**}$, with the equalities holding only for $r^* = \tilde{r}$ (in which case also $x^* = \tilde{x}$). We first show optimality for the policy maker; we next construct beliefs that sustain the strategy for the agents.

When agents coordinate on the equilibrium strategy, the aggregate size of the attack is $A(\theta, r) = A(\theta, \underline{r}) = \Psi((x^* - \theta)/\sigma)$ for any $r < r^*$ and $A(\theta, r) = A(\theta, r^*) = \Psi((\underline{x} - \theta)/\sigma)$ for any $r \geq r^*$. Clearly, the policy maker prefers \underline{r} to any $r \in (\underline{r}, r^*)$ and r^* to any $r > r^*$.¹¹ Furthermore, \underline{r} is dominant for any $\theta \leq 0$. For $\theta > 0$, on the other hand, $A(\theta, r^*) = 0$, and hence the payoff from setting r^* is

¹⁰ If (x, r) is such that $\tilde{\theta} \in \Theta(x)$ and $r = \tilde{r}$, let $\mu(\{\theta = \tilde{\theta}\}|x, r) = 1$ and assume that type $\tilde{\theta}$ (who is indifferent between maintaining the status quo and abandoning it) is expected to abandon it with probability \tilde{r} .

¹¹ We use the convention $\Psi(\pm\infty) \equiv \lim_{\xi \rightarrow \pm\infty} \Psi(\xi)$; hence, when noise is unbounded, in which case $\underline{x} = -\infty$, $A(\theta, r) = 0$ for all θ and $r \geq r^*$.

$\theta - C(r^*)$, whereas the payoff from setting \underline{r} is $\max\{0, \theta - A(\theta, \underline{r})\}$. It follows that r^* is optimal if and only if $C(r^*) \leq \theta$ and $C(r^*) \leq A(\theta, \underline{r})$. (That is, as illustrated in fig. 2, the cost of intervention should not exceed the value of maintaining the status quo, nor should it exceed the opportunity cost of facing the attack associated with setting \underline{r} .) Since $A(\theta, \underline{r})$ is decreasing in θ , we conclude that the strategy of the policy maker is optimal if and only if θ^* and θ^{**} solve

$$\theta^* = C(r^*) = A(\theta^{**}, \underline{r}). \quad (\text{A4})$$

Consider now the agents. An agent who expects all other agents to follow the equilibrium strategy expects the status quo to be abandoned if and only if $\theta < 0$ when $r \geq r^*$, and if and only if $\theta < \hat{\theta}$ when $r < r^*$, where $\hat{\theta}$ solves $\hat{\theta} = \Psi((x^* - \hat{\theta})/\sigma)$. (Note that $\theta^* \leq \hat{\theta} \leq \theta^{**} < -\infty$, with the equalities again holding only for $r^* = \tilde{r}$.) He thus also finds it optimal to follow the equilibrium strategy if and only if his beliefs satisfy

$$\mu(\hat{\theta}|x, r) > r \quad \text{if and only if } x < x^* \text{ for } r < r^* \quad (\text{A5})$$

and

$$\mu(0|x, r) > r \quad \text{if and only if } x < \underline{x} \text{ for } r \geq r^*. \quad (\text{A6})$$

Beliefs are pinned down by Bayes' rule whenever either $r = \underline{r}^{12}$ or $r = r^*$ and $\Theta(x) \cap [\theta^*, \theta^{**}] \neq \emptyset$. In the first case ($r = \underline{r}$), the posterior belief of regime change is given by

$$\mu(\hat{\theta}|x, \underline{r}) = \mu(\theta^*|x, \underline{r}) = \frac{1 - \Psi\left(\frac{x - \theta^*}{\sigma}\right)}{1 - \Psi\left(\frac{x - \theta^*}{\sigma}\right) + \Psi\left(\frac{x - \theta^{**}}{\sigma}\right)}$$

and is decreasing in x . For (A5) to be satisfied, it is thus necessary and sufficient that x^* solves $\mu(\theta^*|x^*, \underline{r}) = \underline{r}$. This in turn holds along with (A4) if and only if (2) holds. In the second case ($r = r^*$ and $\Theta(x) \cap [\theta^*, \theta^{**}] \neq \emptyset$), necessarily $x > \underline{x}$ and from Bayes' rule $\mu(0|x, r^*) = 0$, which clearly satisfies (A6).

Consider next out-of-equilibrium observations. As in part *a*, we require out-of-equilibrium beliefs to satisfy (A2) and (A3). When $r \in (\underline{r}, r^*)$, the set of types for whom r is dominated in equilibrium is $\bar{\Theta}(r) = (-\infty, \theta') \cup (\theta'', +\infty)$, where θ' and θ'' solve $\theta' = C(r) = \Psi((x^* - \theta')/\sigma)$. Note that $0 < \theta' < \theta^* \leq \hat{\theta} \leq \theta^{**} < \theta'' < \infty$ and $\hat{\theta} \in \Theta(x^*)$. Hence, the set of beliefs that satisfy (A5) along with (A2) and (A3) is nonempty. Finally, for any (x, r) such that either $r > r^*$ or $r = r^*$ and $\Theta(x) \cap [\theta^*, \theta^{**}] = \emptyset$, in which case $\Theta(x) \subseteq \bar{\Theta}(r)$, take any beliefs that satisfy (A2) and such that $\mu(0|x, r) = 0$ if $x \geq \underline{x}$ and $\mu(0|x, r) = 1$ otherwise, once again satisfying (A6).

Finally, ensuring that the set of θ who intervene is nonempty puts an upper bound on the policy: $\theta^* \leq \theta^{**}$ if and only if $\theta^* \leq 1 - \underline{r}$ ($= \hat{\theta}$) or, equivalently, $r^* \leq \tilde{r}$ ($= C^{-1}(\hat{\theta})$). QED

¹² That $r = \underline{r}$ is consistent with equilibrium for any x is immediate when noise is unbounded; when it is bounded, it follows from the fact that, by (2), $|\theta^{**} - \theta^*| < 2\sigma$.

Proof of Proposition 3

Since C is Lipschitz continuous, there exists $K < \infty$ such that $|C(r) - C(r')| < K|r - r'|$ for any $(r, r') \in [\underline{x}, \bar{r}]^2$. We prove the result for $\bar{\eta} = \min\{(\bar{r} - \underline{x})/4; (1 - \underline{x})\phi/K\}$, where $\phi = \inf_{\xi \in [-1, 1]} \phi(\xi) > 0$.

Part *a*. Optimality of the strategy of the policy maker is immediate, given the agents' strategies. Thus consider the agents. When $z \leq \underline{x} + \eta$, beliefs are pinned down by Bayes' rule, in which case sequential rationality follows directly from the same argument as in the proof of part *a* of proposition 2.

When instead $z > \underline{x} + \eta$, we consider out-of-equilibrium beliefs that ensure that the equilibrium satisfies the intuitive criterion. Let

$$P(x, z) = \{(\theta, r) \in \Theta(x) \times R(z) : \theta \notin \bar{\Theta}(r)\}$$

denote the set of type-policy pairs (θ, r) that are compatible with the observation of (x, z) and such that r is *not* dominated in equilibrium for θ . We sustain the equilibrium strategy with out-of-equilibrium beliefs that satisfy the following properties:

$$\mu(\{(\theta, r) \in \Theta(x) \times R(z) | x, z\}) = 1 \tag{A7}$$

and

$$\mu(\{(\theta, r) \in P(x, z)\}) = 1 \text{ whenever } P(x, z) \neq \emptyset. \tag{A8}$$

Condition (A7) requires that beliefs assign zero measure to (θ, r) that are inconsistent with the agent's signals, whereas condition (A8) requires that beliefs assign zero measure to type-policy pairs for whom a deviation is dominated in equilibrium whenever possible (which is sufficient, although not necessary, for the intuitive criterion to be satisfied).

Take first (x, z) such that $\theta \notin \Theta(x)$, which is possible only when ξ is bounded. Pick any $(t, \rho) \in P(x, z)$ if $P(x, z) \neq \emptyset$ and any $(t, \rho) \in \Theta(x) \times R(z)$ otherwise, and let $\mu(\{\theta = t, r = \rho | x, z\}) = 1$. Since $\theta \notin \Theta(x)$, either $x < \theta - \sigma < \tilde{x}$, in which case $t \leq \max \Theta(x) < \tilde{\theta}$ and hence $\mu(\{\theta < \tilde{\theta} | x, z\}) = 1$, or $x > \tilde{\theta} + \sigma > \tilde{x}$, in which case $\tilde{\theta} < \min \Theta(x) \leq t$ and hence $\mu(\{\theta < \tilde{\theta} | x, z\}) = 0$.

Next, take (x, z) such that $\theta \in \Theta(x)$ and $z \neq \tilde{r} + \eta$. Note that $P(x, z) = \emptyset$ for $z > \tilde{r} + \eta$, whereas $P(x, z) \neq \emptyset$ for $z < \tilde{r} + \eta$; indeed, $R(z) \cap (\underline{x} + 2\eta, \tilde{r}) \neq \emptyset$ by the assumption that $\eta < \bar{\eta}$, and any $\rho \in R(z) \cap (\underline{x} + 2\eta, \tilde{r})$ is never dominated in equilibrium for all $\theta \in [\theta'(\rho), \theta''(\rho)]$, where $\theta'(\rho)$ and $\theta''(\rho)$ solve $\theta' = C(\rho) = \Psi((\tilde{x} - \theta')/\sigma)$ and satisfy $0 < \theta'(\rho) < \tilde{\theta} < \theta''(\rho) < \infty$.¹³ Pick any $(t, \rho) \in \Theta(x) \times R(z)$ such that $\rho \in (\underline{x} + 2\eta, \tilde{r})$, $t < \tilde{\theta}$ if $x < \tilde{x}$, $t > \tilde{\theta}$ if $x \geq \tilde{x}$, and $t \notin \bar{\Theta}(\rho)$ if $P(x, z) \neq \emptyset$. Then let $\mu(\{\theta = t, r = \rho | x, z\}) = 1$.

Finally, consider (x, z) such that $\theta \in \Theta(x)$ and $z = \tilde{r} + \eta$, in which case $P(x, z) = \{(\tilde{\theta}, \tilde{r})\}$. Then let $\mu(\{\theta = \tilde{\theta}, r = \tilde{r} | x, z\}) = 1$ and, as in the proof of proposition 2, assume that the agent expects θ to abandon the regime with probability \tilde{r} .

¹³ Note that for all $\theta \in [\theta', \theta'']$, $x > \underline{x}$. Hence, an agent with signals (x, z) with $z > \underline{x} + \eta$ who believes that $r' > \underline{x} + 2\eta$ and who expects all other agents not to attack also finds it optimal not to attack. It follows that the minimum size of attack for a type $\theta \in [\theta', \theta'']$ who sets the policy at $r' \in (\underline{x} + 2\eta, \tilde{r})$ is $A(\theta, r') = 0$, which implies that r' is not dominated in equilibrium.

The beliefs specified above satisfy (A7) and (A8) and guarantee that an agent who expects all other agents to attack if and only if $x < \tilde{x}$ for any z finds it optimal to follow the same strategy.

Part *b*. Let $r_\eta \equiv \max\{\underline{r} + 4\eta, C^{-1}(K\eta/\phi)\}$; r_η is increasing in η , $r_\eta \rightarrow \underline{r}$ as $\eta \rightarrow 0$, and $r_\eta < \tilde{r}$ for any $\eta < \tilde{\eta}$. Step 1 establishes optimality for the policy maker whereas step 2 does so for the agents.

Step 1. For $\theta < 0$, \underline{r} is dominant. For $\theta \geq 0$,

$$A(\theta, r) = \Pr(x < x^*, z < z^* | \theta, r) \\ = \begin{cases} 0 & \text{for } r \geq r^* \\ \Psi\left(\frac{x^* - \theta}{\sigma}\right) \Phi\left(\frac{z^* - r}{\eta}\right) & \text{for } r \in [z^* - \eta, r^*) \\ \Psi\left(\frac{x^* - \theta}{\sigma}\right) & \text{for } r \leq z^* - \eta. \end{cases}$$

Clearly, \underline{r} is preferred to any $r \in (\underline{r}, z^* - \eta)$, and r^* is preferred to any $r > r^*$ by all θ .

Consider first $\theta \in [\theta^*, \theta^{**}]$. The policy maker prefers r^* to $r \in [z^* - \eta, r^*)$ if and only if

$$C(r^*) - C(r) \leq \min\{\theta, A(\theta, r)\}. \quad (\text{A9})$$

Since

$$\begin{aligned} \min_{\theta \in [\theta^*, \theta^{**}]} (\min\{\theta, A(\theta, r)\}) &= \min\{\theta^*, A(\theta^{**}, r)\} \\ &= A(\theta^{**}, r) = \Psi\left(\frac{x^* - \theta^{**}}{\sigma}\right) \Phi\left(\frac{z^* - r}{\eta}\right) \\ &= C(r^*) \Phi\left(\frac{z^* - r}{\eta}\right), \end{aligned}$$

(A9) holds for all $\theta \in [\theta^*, \theta^{**}]$ if and only if $C(r^*) - C(r) \leq C(r^*) \Phi((z^* - r)/\eta)$. By Lipschitz continuity of C and absolute continuity of Φ , $C(r^*) - C(r) \leq K(r^* - r)$ and

$$\Phi\left(\frac{z^* - r}{\eta}\right) = \int_{-1}^{(z^* - r)/\eta} \phi(\xi) d\xi \geq \frac{1}{\eta} \phi(r^* - r),$$

whereas $r^* > r_\eta$ implies $K < C(r^*)\phi/\eta$. It follows that, for all $r \in [z^* - \eta, r^*)$,

$$C(r^*) - C(r) - C(r^*) \Phi\left(\frac{z^* - r}{\eta}\right) \leq \left[K - \frac{1}{\eta} \phi C(r^*)\right](r^* - r) < 0, \quad (\text{A10})$$

which in turn suffices for (A9). Furthermore, since

$$\theta - C(r^*) \geq \max\left\{\theta - \Psi\left(\frac{x^* - \theta}{\sigma}\right), 0\right\}$$

for all $\theta \in [\theta^*, \theta^{**}]$, r^* is also preferred to \underline{r} , and hence it is optimal.

Consider next $\theta \in [0, \theta^*]$. Since

$$A(\theta, r) = \Psi\left(\frac{x^* - \theta}{\sigma}\right) \Phi\left(\frac{z^* - r}{\eta}\right) > \theta \Phi\left(\frac{z^* - r}{\eta}\right)$$

and $\theta < C(r^*)$,

$$\begin{aligned} U(\theta, r, A(\theta, r)) &= \theta - \min\{A(\theta, r), \theta\} - C(r) \\ &\leq \theta \left[1 - \Phi\left(\frac{z^* - r}{\eta}\right)\right] - C(r) < C(r^*) \left[1 - \Phi\left(\frac{z^* - r}{\eta}\right)\right] - C(r), \end{aligned}$$

for all $r \in [z^* - \eta, r^*]$. By (A10) then, $U(\theta, r, A(\theta, r)) < 0 = U(\theta, \underline{r}, A(\theta, \underline{r}))$, which proves that \underline{r} is optimal.

Consider finally $\theta > \theta^{**}$. Since

$$A(\theta, r) = \Psi\left(\frac{x^* - \theta}{\sigma}\right) \Phi\left(\frac{z^* - r}{\eta}\right) < \Psi\left(\frac{x^* - \theta}{\sigma}\right) < C(r^*) < \theta,$$

$$\begin{aligned} U(\theta, r, A(\theta, r)) &= \theta - A(\theta, r) - C(r) = \theta - \Phi\left(\frac{z^* - r}{\eta}\right) \Psi\left(\frac{x^* - \theta}{\sigma}\right) - C(r) \\ &= \theta - \Psi\left(\frac{x^* - \theta}{\sigma}\right) + \left[1 - \Phi\left(\frac{z^* - r}{\eta}\right)\right] \Psi\left(\frac{x^* - \theta}{\sigma}\right) - C(r) \\ &< \theta - \Psi\left(\frac{x^* - \theta}{\sigma}\right) + \left[1 - \Phi\left(\frac{z^* - r}{\eta}\right)\right] C(r^*) - C(r). \end{aligned}$$

By (A10) then,

$$U(\theta, r, A(\theta, r)) < \theta - \Psi\left(\frac{x^* - \theta}{\sigma}\right) = U(\theta, \underline{r}, A(\theta, \underline{r})),$$

and hence \underline{r} is optimal.

Step 2. Consider next the agents. When $z \in [\underline{r} - \eta, \underline{r} + \eta]$, beliefs are always pinned down by Bayes' rule, since (2) ensures $\Theta(x) \not\subseteq [\theta^*, \theta^{**}]$ for any x . It follows that $\mu(\{r = \underline{r}\} | x, z) = 1$ and

$$\begin{aligned} \mu(\{\theta < \hat{\theta}, r = \underline{r}\} | x, z) &= \mu(\{\theta < \theta^*, r = \underline{r}\} | x, z) \\ &= \frac{1 - \Psi\left(\frac{x - \theta^*}{\sigma}\right)}{1 - \Psi\left(\frac{x - \theta^*}{\sigma}\right) + \Psi\left(\frac{x - \theta^{**}}{\sigma}\right)}, \end{aligned}$$

where $\hat{\theta} \in (\theta^*, \theta^{**})$ solves $\Psi((x^* - \hat{\theta})/\sigma) = \hat{\theta}$ and x^* is as in (2). Therefore, when $z \in [\underline{r} - \eta, \underline{r} + \eta]$, it is optimal for agents to attack if and only if $x \leq x^*$. Beliefs are pinned down by Bayes' rule also for (x, z) such that $z \in [r^* - \eta, r^* + \eta]$ and $\Theta(x) \cap [\theta^*, \theta^{**}] \neq \emptyset$, in which case $\mu(\{\theta \in [\theta^*, \theta^{**}], r = r^*\} | x, z) = 1$, implying that it is optimal not to attack.

Next, take any (x, z) such that $z \in (\underline{r} + \eta, r^* - \eta)$. For any $r' \in (\underline{r} + 2\eta, r^*)$, let θ' and θ'' solve $\theta' = C(r') = \Psi((x^* - \theta)/\sigma)$ and note that

$$0 < \theta' \leq \min \bar{\Theta}(r') \leq \theta^* < \tilde{\theta} < \theta^{**} < \max \bar{\Theta}(r') = \theta'' < \infty,$$

with $\min \bar{\Theta}(r') = \theta'$ for $r' \leq r^* - 2\eta$. It follows that if $\{t, \rho\} \in P(x, z)$ for some t and some $\rho > r^* - 2\eta$, then also $\{t, \rho'\} \in P(x, z)$ for some $\rho' < r^* - 2\eta$. Hence, we can always find $(t, \rho) \in \Theta(x) \times R(z)$ such that $\rho \in [z - \eta, r^* - 2\eta]$, $t < \theta^*$ if $x < x^*$, $t > \theta^*$ if $x > x^*$, and $(t, \rho) \in P(x, z)$ if $P(x, z) \neq \emptyset$. Then let

$$\mu(\{\theta = t, r = \rho\} | x, z) = 1.$$

These beliefs satisfy (A7) and (A8) and guarantee optimality.

For any (x, z) such that $z > r^* + \eta$, in which case necessarily $P(x, z) = \emptyset$, pick any $(t, \rho) \in \Theta(x) \times R(z)$ such that $t < 0$ if and only if $x < \bar{x}$ and let $\mu(\{\theta = t, r = \rho\} | x, z) = 1$. Again, μ satisfies (A7) and (A8) and guarantees optimality.

Finally, take any (x, z) such that $z \in [r^* - \eta, r^* + \eta]$ and $\Theta(x) \cap [\theta^*, \theta^{**}] = \emptyset$, which can be the case only when ξ is bounded. Pick any $(t, \rho) \in \Theta(x) \times R(z)$ such that $t < 0$ if and only if $x < \bar{x}$ and let $\mu(\{\theta = t, r = \rho\} | x, z) = 1$. These beliefs suffice for optimality and satisfy (A7) but not (A8) because $\rho > r^*$ is dominated in equilibrium for all $\theta \in \Theta(x)$. However, this is not a problem since in this case imposing (A8) would only make the agents (weakly) more aggressive, thus making the deviation even less profitable for the policy maker. QED

Proof of Proposition 4

The existence of at least two equilibria follows immediately from parts *a* and *b* by choosing some $r^* \in (\underline{r}, \bar{r})$ and letting $(\bar{\eta}, \bar{\rho}) = (\hat{\eta}, \hat{\rho})$, where $(\hat{\eta}, \hat{\rho})$ are the bounds from part *b*.

Part *a*. When $r(\theta) = \underline{r}$ for all θ, z conveys no information about θ , and hence

$$\Pr [D = 1 | x, z] = \Pr [\theta \leq \tilde{\theta} | x, z] = 1 - \Psi\left(\frac{x - \tilde{\theta}}{\sigma}\right)$$

and $\mathbb{E}[r | x, z] = \underline{r}$ for any z . An agent thus finds it optimal to attack if and only if $x < \tilde{x}$. The size of the attack is then given by $A(\theta, r) = \Psi((\tilde{x} - \theta)/\sigma)$ and is independent of r , implying that the policy maker finds it optimal to set $r(\theta) = \underline{r}$ for all θ and to abandon the regime if and only if $\theta < \theta$.

Part *b*. Fix an $r \in (\underline{r}, r^*)$. We prove the result with two lemmas. Lemma 1 is the heart of the argument: it considers the behavior of the policy maker and the agents and identifies conditions for the thresholds $(x', \hat{x}, \theta', \theta'')$ that guarantee that the proposed strategies are part of an equilibrium. Lemma 2 then proves that a solution to this system of conditions exists if (η, ρ) are small enough and that the thresholds $(x', \hat{x}, \theta', \theta'')$ are arbitrarily close to the corresponding thresholds of the benchmark model as $(\eta, \rho) \rightarrow (0, 0)$.

LEMMA 1. Take any $r^* \in (\underline{r}, \bar{r})$. Suppose that there exist $(x', \hat{x}, \theta', \theta'')$ that satisfy the following conditions:

$$1 - \Psi\left(\frac{x' - \theta'}{\sigma}\right) = \underline{r} - (r - \rho r^*) \left[\Psi\left(\frac{x' - \theta'}{\sigma}\right) - \Psi\left(\frac{x' - \theta''}{\sigma}\right) \right], \quad (\text{A11})$$

$$1 - \Psi\left(\frac{\hat{x} - \theta'}{\sigma}\right) = \underline{r} + \left[r^* \rho + r^* (1 - \rho) \exp\left(\frac{r^* - \underline{r}}{\eta}\right) - \underline{r} \right] \left[\Psi\left(\frac{\hat{x} - \theta'}{\sigma}\right) - \Psi\left(\frac{\hat{x} - \theta''}{\sigma}\right) \right], \quad (\text{A12})$$

$$\begin{aligned} \theta' &= \rho \left[1 - \exp\left(-\frac{r^* - \underline{r}}{\eta}\right) \right] \Psi\left(\frac{x' - \theta'}{\sigma}\right) \\ &\quad + \left[1 - \rho + \rho \exp\left(-\frac{r^* - \underline{r}}{\eta}\right) \right] \Psi\left(\frac{\hat{x} - \theta'}{\sigma}\right) + C(r^*), \end{aligned} \quad (\text{A13})$$

$$C(r^*) = (1 - \rho) \left[1 - \exp\left(-\frac{r^* - \underline{r}}{\eta}\right) \right] \left[\Psi\left(\frac{x' - \theta''}{\sigma}\right) - \Psi\left(\frac{\hat{x} - \theta''}{\sigma}\right) \right], \quad (\text{A14})$$

$$C(r^*) \leq (1 - \rho) \left[1 - \exp\left(-\frac{r^* - \underline{r}}{\eta}\right) \right] \left[\Psi\left(\frac{x' - \theta'}{\sigma}\right) - \Psi\left(\frac{\hat{x} - \theta'}{\sigma}\right) \right]. \quad (\text{A15})$$

Then the strategies in part *b* of the proposition are part of an equilibrium.

These conditions have the following interpretation. An agent is indifferent between attacking and not attacking if and only if his posterior probability of regime change, $\mu(\theta < \theta' | x, z)$, equals his posterior expected cost of attacking, $\mathbb{E}[r | x, z]$. The combination of the particular noise structure we have assumed with the property that the equilibrium policy takes only two values (namely \underline{r} or r^*) implies that, for any given x , the net expected payoff from attacking is a step function of z with a single downward discontinuity at $z = z^*$. Moreover, net expected payoffs from attacking are decreasing in x . Indifference for the agents can then be summarized by two signal thresholds \hat{x} and x' , with $\hat{x} < x'$, such that an agent who observes a policy signal $z < z^*$ ($z \geq z^*$) is indifferent between attacking and not attacking at $x = x'$ ($x = \hat{x}$). These thresholds are characterized by conditions (A11) and (A12). Conditions (A13) and (A14), on the other hand, mean that the policy maker is indifferent between \underline{r} and r^* at $\theta = \theta'$ and at $\theta = \theta''$. The agents' strategy implies that $A(\theta, r)$ is insensitive to r for $r > r^*$, whereas the particular noise structure and the linearity of C ensure that the policy maker's payoff is convex in r for $r \in [\underline{r}, r^*]$. It follows that the problem of the policy maker reduces to a choice between \underline{r} and r^* . Condition (A15) then guarantees that r^* is preferred to \underline{r} if and only if $\theta \in [\theta', \theta'']$.

In the proof of lemma 1 below we derive these conditions and show that they appropriately define the policy maker's and agents' indifference conditions. We also show that with our particular noise structure, $A(\theta, r)$ is constant in r for all $r \geq r^*$ and concave in r for $r \leq r^*$, which in turn ensures that any $r \notin \{\underline{r}, r^*\}$ is dominated by $r \in \{\underline{r}, r^*\}$.

Proof of Lemma 1

We prove the lemma in two steps. Step 1 considers the policy maker; step 2 considers the agents.

Step 1. Consider the policy maker. When agents follow the proposed strategy, the size of the attack is given by

$$\begin{aligned}
 A(\theta, r) &= \left[1 - (1 - \rho) \exp\left(-\frac{r^* - r}{\eta}\right) - \rho \exp\left(-\frac{r^* - \underline{r}}{\eta}\right) \right] \Psi\left(\frac{x' - \theta}{\sigma}\right) \\
 &\quad + \left[(1 - \rho) \exp\left(-\frac{r^* - r}{\eta}\right) + \rho \exp\left(-\frac{r^* - \underline{r}}{\eta}\right) \right] \Psi\left(\frac{\hat{x} - \theta}{\sigma}\right) \\
 &\quad \text{for } r < r^*, \\
 A(\theta, r) &= \left[\rho - \rho \exp\left(-\frac{r^* - \underline{r}}{\eta}\right) \right] \Psi\left(\frac{x' - \theta}{\sigma}\right) \\
 &\quad + \left[1 - \rho + \rho \exp\left(-\frac{r^* - \underline{r}}{\eta}\right) \right] \Psi\left(\frac{\hat{x} - \theta}{\sigma}\right) \quad \text{for } r \geq r^*.
 \end{aligned}$$

The size of the attack $A(\theta, r)$ is strictly decreasing in θ , equals $A(\theta, r^*)$ for any $r \geq r^*$, and is strictly decreasing and strictly concave in $r \in (\underline{r}, r^*)$. Together with the linearity of C , this implies that any $r \notin \{\underline{r}, r^*\}$ is dominated by either \underline{r} or r^* . For the proposed strategy to be optimal, it must be that the policy maker prefers to set $r = \underline{r}$ and abandon the status quo for $\theta < \theta'$, set $r = r^*$ and maintain it for $\theta \in [\theta', \theta'']$, and maintain it while setting $r = \underline{r}$ for $\theta > \theta''$.

Let $U_1(\theta)$ and $U_2(\theta)$ denote the payoffs from setting, respectively, r^* and \underline{r} while maintaining the status quo:

$$\begin{aligned}
 U_1(\theta) &\equiv \theta - \rho \left[1 - \exp\left(-\frac{r^* - \underline{r}}{\eta}\right) \right] \Psi\left(\frac{x' - \theta}{\sigma}\right) \\
 &\quad - \left[1 - \rho + \rho \exp\left(-\frac{r^* - \underline{r}}{\eta}\right) \right] \Psi\left(\frac{\hat{x} - \theta}{\sigma}\right) - C(r^*), \\
 U_2(\theta) &\equiv \theta - \left[1 - \exp\left(-\frac{r^* - \underline{r}}{\eta}\right) \right] \Psi\left(\frac{x' - \theta}{\sigma}\right) - \exp\left(-\frac{r^* - \underline{r}}{\eta}\right) \Psi\left(\frac{\hat{x} - \theta}{\sigma}\right).
 \end{aligned}$$

The two thresholds θ' and θ'' must thus solve $U_1(\theta') = 0$ and $U_2(\theta'') = U_1(\theta'')$; these two indifference conditions are given in (A13) and (A14).

Let

$$q(\theta) \equiv (1 - \rho) \left[1 - \exp\left(-\frac{r^* - \underline{r}}{\eta}\right) \right] \left[\Psi\left(\frac{x' - \theta}{\sigma}\right) - \Psi\left(\frac{\hat{x} - \theta}{\sigma}\right) \right].$$

Since $\hat{x} < x'$, the distribution $1 - \Psi((x' - \theta)/\sigma)$ first-order stochastically dominates the distribution $1 - \Psi((\hat{x} - \theta)/\sigma)$. It follows that $\psi((x' - \theta)/\sigma)/\psi((\hat{x} - \theta)/\sigma)$ is in-

creasing in θ , and hence there exists a unique $\hat{\theta}$ such that

$$\psi\left(\frac{x' - \theta}{\sigma}\right) \equiv \psi\left(\frac{\hat{x} - \theta}{\sigma}\right)$$

if and only if $\theta \equiv \hat{\theta}$; equivalently, $q(\theta)$ is increasing for $\theta < \hat{\theta}$ and decreasing for $\theta > \hat{\theta}$. Furthermore, $\lim_{\theta \rightarrow -\infty} q(\theta) = \lim_{\theta \rightarrow +\infty} q(\theta) = 0$, and hence (A14) admits at most two solutions, which we denote with θ_1 and θ_2 ($\theta_1 < \theta_2$).

To sustain the proposed equilibrium, it must be the case that $\theta_1 \leq \theta' \leq \theta_2$, which holds if and only if (A15) is satisfied. In this case $U_1(\theta) \geq U_2(\theta)$ if and only if $\theta \in [\theta_1, \theta_2]$ and $U_1(\theta) \geq 0$ if and only if $\theta \geq \theta'$. But then, for all $\theta > \theta_2$, $U_2(\theta) \geq U_1(\theta) \geq 0$, which implies that it is optimal to set $r = \underline{r}$ and maintain the status quo; for $\theta \in [\theta', \theta_2]$, $U_1(\theta) \geq \max\{0, U_2(\theta)\}$, which implies that it is optimal to set $r = r^*$. For $\theta < \theta'$, we have that either $\theta \geq \theta_1$ and hence $U_2(\theta) \leq U_1(\theta) < 0$ or $\theta < \theta_1$ and $U_1(\theta) \leq U_2(\theta) < U_2(\theta_1) = U_1(\theta_1) < 0$ (by monotonicity); in either case, it is optimal to set \underline{r} and abandon the status quo. If instead $\theta' < \theta_1$, setting r^* would not be optimal for $\theta \in [\theta', \theta_1)$; and if $\theta' > \theta_2$, r^* would never be optimal.

Step 2. Consider now the agents. Let $u(x, z) \equiv \Pr[\theta \leq \theta' | x, z] - \mathbb{E}[r | x, z]$ be the expected payoff from attacking. Since ζ is exponential, that is, $\Phi(\zeta) = 1 - \exp(-\zeta)$, the likelihood ratio of r^* to \underline{r} conditional on z is

$$\rho + (1 - \rho) \exp\left(\frac{r^* - \underline{r}}{\eta}\right)$$

for $z \geq z^*$ and ρ for $z < z^*$. It follows that

$$\begin{aligned} u(x, z) &= \left(1 - \Psi\left(\frac{x - \theta'}{\sigma}\right) - \left[\underline{r} - (r - \rho r^*) \left[\Psi\left(\frac{x - \theta'}{\sigma}\right) - \Psi\left(\frac{x - \theta''}{\sigma}\right)\right]\right]\right) \\ &\quad \div \left[1 - (1 - \rho) \left[\Psi\left(\frac{x - \theta'}{\sigma}\right) - \Psi\left(\frac{x - \theta''}{\sigma}\right)\right]\right] \quad \text{for } z < r^*, \\ u(x, z) &= \left(1 - \Psi\left(\frac{x - \theta'}{\sigma}\right) - \left[\underline{r} + \left[r^* \rho + r^* (1 - \rho) \exp\left(\frac{r^* - \underline{r}}{\eta}\right) - \underline{r}\right]\right.\right. \\ &\quad \left.\left. \times \left[\Psi\left(\frac{x - \theta'}{\sigma}\right) - \Psi\left(\frac{x - \theta''}{\sigma}\right)\right]\right)\right) \\ &\quad \div \left[1 + \left[\rho + (1 - \rho) \exp\left(\frac{r^* - \underline{r}}{\eta}\right) - 1\right] \left[\Psi\left(\frac{x - \theta'}{\sigma}\right) - \Psi\left(\frac{x - \theta''}{\sigma}\right)\right]\right] \\ &\quad \text{for } z \geq r^*. \end{aligned} \tag{A16}$$

For any z , $u(x, z)$ is continuous in x , $u(x, z) \rightarrow 1 - \underline{r} > 0$ as $x \rightarrow -\infty$, and $u(x, z) \rightarrow -\underline{r} < 0$ as $x \rightarrow +\infty$. Furthermore, for any x , $u(x, z)$ is a step function in z , with discontinuity at $z = r^*$. It follows that there exist thresholds x' and \hat{x} such that x' solves $u(x', z) = 0$, or equivalently (A11), for $z < r^*$; and \hat{x} solves $u(\hat{x}, z) = 0$, or equivalently (A12), for $z \geq r^*$.

Let $N(x, z)$ and $D(x, z) > 0$ be, respectively, the numerator and the denominator

of $u(x, z)$ in (A16). When $z < r^*$, $\rho < \underline{r}/r^*$ suffices for $N(x, z)$ to be strictly decreasing in x , and hence for $\partial u(x', z)/\partial x < 0$. When $z \geq r^*$, on the other hand, $N(x, z) = [1 - \Psi((x - \theta')/\sigma)]H(x)$, where

$$H(x) = 1 - \frac{\underline{r}}{1 - \Psi\left(\frac{x - \theta'}{\sigma}\right)} - \frac{\left[r^*\rho + r^*(1 - \rho) \exp\left(\frac{r^* - \underline{r}}{\eta}\right) - \underline{r}\right] \left[\Psi\left(\frac{x - \theta'}{\sigma}\right) - \Psi\left(\frac{x - \theta''}{\sigma}\right)\right]}{1 - \Psi\left(\frac{x - \theta'}{\sigma}\right)}.$$

At $x = \hat{x}$, necessarily $H(\hat{x}) = 0$, which implies that $\partial N(\hat{x}, z)/\partial x = H'(\hat{x}) < 0$ since ψ (and hence $1 - \Psi$) is log concave, and therefore $\partial u(\hat{x}, z)/\partial x < 0$. It follows that, given θ' and θ'' , x' and \hat{x} are the *unique* solutions to (A11) and (A12). QED

LEMMA 2. For any $r^* \in (\underline{r}, \tilde{r})$ and $\epsilon > 0$, there exist $\bar{\eta} > 0$ and $\bar{\rho} < \underline{r}/r^*$ such that, for any $(\eta, \rho) < (\bar{\eta}, \bar{\rho})$, conditions (A13) and (A12) admit a solution $(x', \hat{x}, \theta', \theta'')$ that satisfies $\theta' \leq \theta''$, $|x' - x^*| < \epsilon$, $|\theta' - \theta^*| < \epsilon$, $|\theta'' - \theta^{**}| < \epsilon$, and $\hat{x} < -1/\epsilon$.

The proof of this lemma is technical and is published in online Appendix B. Combining the two lemmas completes the proof. QED

Proof of Proposition 5

We prove the result in three steps. The proposed strategies are shown to be part of an equilibrium in step 3; steps 1 and 2 construct the set Θ^{**} and the function X and establish conditions that are useful for step 3.

Step 1. Fix $r^* \in (\underline{r}, \tilde{r})$ and let

$$\begin{aligned} \theta^* &= \min \{ \theta : R(\theta, 0, r^*) \geq 0 \text{ and } U(\theta, 0, r^*) \geq 0 \} \\ &= \min \{ \theta \geq 0 : U(\theta, 0, r^*) \geq 0 \}. \end{aligned}$$

Next, define the correspondence $S: \mathbb{R} \rightarrow 2^{\mathbb{R}_+}$ and the function $m: \mathbb{R}^2 \rightarrow [0, 1]$ as

$$\begin{aligned} S(x) &\equiv \left\{ \theta \geq \theta^* : R\left(\theta, \Psi\left(\frac{x - \theta}{\sigma}\right), \underline{r}\right) < 0 \text{ or} \right. \\ &\quad \left. U\left(\theta, \Psi\left(\frac{x - \theta}{\sigma}\right), \underline{r}\right) \leq U(\theta, 0, r^*) \right\}, \\ m(x, x') &\equiv \frac{1 - \Psi\left(\frac{x' - \theta^*}{\sigma}\right)}{1 - \int_{S(x)} \frac{1}{\sigma} \psi\left(\frac{x' - \theta^*}{\sigma}\right) d\theta}, \end{aligned}$$

where θ^* is the lowest type who is willing to raise the policy at r^* if this ensures that no agent attacks; $S(x)$ is the set of $\theta \geq \theta^*$ who prefer r^* to \underline{r} when agents

do not attack when $r = r^*$ and attack if and only if their signal is less than x when $r = \underline{r}$; $m(x, x')$ in turn is the posterior probability of regime change for an agent with signal x' when he observes \underline{r} and believes that the regime is abandoned if and only if $\theta < \theta^*$ and that the policy is $r(\theta) = \underline{r}$ if and only if $\theta \notin S(x)$.

Step 2 below shows that either (i) there exists an $x^* \in \mathbb{R}$ such that $m(x^*, x^*) = g(\underline{r})$ or (ii) it is the case that $m(x, x) > g(\underline{r})$ for all x , in which case we let $x^* = \infty$. In either case, $\Theta^{**} = S(x^*)$. The triplet $(x^*, \theta^*, \Theta^{**})$ thus identifies an equilibrium for the fictitious game in which the policy maker is restricted to set $r \in \{\underline{r}, r^*\}$ and the agents are restricted not to attack when $r = r^*$. Step 3 shows that this is also part of an equilibrium for the unrestricted game.

Step 2. Note that $S(x)$ is continuous in x , and $S(x_1) \subseteq S(x_2)$ for any $x_1 \leq x_2$ (because $R(\theta, A, r)$ and $U(\theta, A, r)$ are nonincreasing and continuous in A and $\Psi((x - \theta)/\sigma)$ is nondecreasing and continuous in x), whereas $m(x, x')$ is continuous in (x, x') , nondecreasing in x (by the monotonicity of S), and nonincreasing in x' (by the log concavity of ψ). Moreover, for any $(x, x') \in \mathbb{R}^2$, we have

$$S(x) \subseteq \bar{S} \equiv \{\theta \geq \theta^* : R(\theta, 1, \underline{r}) < 0 \text{ or } U(\theta, 1, \underline{r}) \leq U(\theta, 0, r^*)\}$$

and

$$1 \geq \frac{1 - \Psi\left(\frac{x' - \theta^*}{\sigma}\right)}{1 - \int_{\bar{S}} \frac{1}{\sigma} \psi\left(\frac{x' - \theta^*}{\sigma}\right) d\theta} \geq m(x, x') \geq 1 - \Psi\left(\frac{x' - \theta^*}{\sigma}\right).$$

It follows that, for any x , $m(x, x') \geq g(\underline{r})$ for all $x' \leq x^\#$, where $x^\# \in \mathbb{R}$ is the solution to $1 - \Psi((x^\# - \theta^*)/\sigma) = g(\underline{r})$.

Now define the sequence $\{x_k\}_{k=0}^\infty$, with $x_k \in \mathbb{R} \cup \{+\infty\}$ as follows. For $k = 0$, let $x_0 = x^\#$. For $k \geq 1$, let x_k be the solution to $m(x_{k-1}, x_k) = g(\underline{r})$ if $x_{k-1} < \infty$; if, instead, $x_{k-1} = \infty$, let $x_k = \inf\{x' : m(x_{k-1}, x') \leq g(\underline{r})\}$ if $\{x' : m(x_{k-1}, x') \leq g(\underline{r})\} \neq \emptyset$ and $x_k = \infty$ otherwise. The fact that $m(x^\#, x^\#) \geq 1 - \Psi((x^\# - \theta^*)/\sigma) = g(\underline{r})$, together with the continuity and monotonicities of m , ensures that this sequence is well defined and nondecreasing. It follows that either $\lim_{k \rightarrow \infty} x_k \in [x^\#, +\infty)$ or $\lim_{k \rightarrow \infty} x_k = +\infty$. In the former case, let $x^* = \lim x_k$ and $\Theta^{**} = S(x^*)$; in the latter, let $x^* = \infty$ and $\Theta^{**} = S(\infty) \equiv \bar{S}$.

Note that $\theta^* \in [0, \theta]$ and $x^* > \hat{x}$, where $\tilde{\theta} \in (0, 1)$ is the solution to $R(\tilde{\theta}, 1 - g(\underline{r}), \underline{r}) = 0$ and $\hat{x} \in \mathbb{R}$ is the solution to $R(\theta^*, \Psi((\hat{x} - \theta^*)/\sigma), \underline{r}) = 0$. That $\theta^* < \theta$ follows immediately from $r^* < \tilde{r}$, and hence $U(\tilde{\theta}, 0, r^*) > U(\theta, 0, \tilde{r}) = 0$. To see that $x^* > \hat{x}$, note that, by the definitions of \hat{x} , θ , and $x^\#$,

$$R\left(\theta^*, \Psi\left(\frac{\hat{x} - \theta^*}{\sigma}\right), \underline{r}\right) = 0 = R(\tilde{\theta}, 1 - g(\underline{r}), \underline{r}) = R\left(\tilde{\theta}, \Psi\left(\frac{x^\# - \theta^*}{\sigma}\right), \underline{r}\right),$$

which together with $\theta^* < \tilde{\theta}$ implies $x^\# > \hat{x}$ and therefore $x^* \geq x^\# > \hat{x}$. This in turn implies that there exists a $\hat{\theta} \in (\theta^*, 1)$ that solves $R(\hat{\theta}, \Psi((x^* - \hat{\theta})/\sigma), \underline{r}) = 0$ and that is such that $R(\theta, \Psi((x^* - \theta)/\sigma), \underline{r}) < 0$ if and only if $\theta < \hat{\theta}$. But then we necessarily have that $D(\theta) = 1$ if and only if $\theta < \theta^*$ and that $[\theta^*, \hat{\theta}] \subseteq S(x^*)$, which gives $\min \Theta^{**} = \theta^*$.

Finally, when the noise is bounded, $m(x, x') = 0$ for all $x' \geq \theta^* + \sigma$, and

$S(x) \subseteq [\theta^*, x + \sigma]$ if $x + \sigma \geq \theta^*$ and $S(x) = \emptyset$ otherwise. It follows that, with bounded noise, $x^* < \theta^* + \sigma$, $\Theta^{**} \subseteq [\theta^*, \theta^* + 2\sigma]$, and $\hat{\theta} \in (\theta^*, \theta^* + 2\sigma)$.

Step 3. Let $X: [\underline{x}, \bar{r}] \rightarrow \mathbb{R} \cup \{\pm\infty\}$ be defined as follows: $X(\underline{x}) = x^*$; for $r \in (\underline{x}, r^*)$, $X(r) = \infty$ if $x^* = \infty$ and otherwise $X(r) \geq x^*$ is the solution to $R(\hat{\theta}, \Psi((X(r) - \hat{\theta})/\sigma), r) = 0$; finally, for $r \in [r^*, \bar{r}]$, $X(r) = \underline{x}$. We now show that the strategies in the proposition are part of an equilibrium.

Consider first the policy maker. By construction of $X(r)$, for any $r < r^*$, $A(\theta, r) \geq A(\theta, \underline{x})$ and $\text{sign}\{R(\theta, A(\theta, r), r)\} = \text{sign}\{R(\theta, A(\theta, \underline{x}), \underline{x})\}$, whereas for any $r \geq r^*$, $A(\theta, r) = A(\theta, r^*)$ and $\text{sign}\{R(\theta, A(\theta, r), r)\} = \text{sign}\{R(\theta, A(\theta, r^*), r^*)\}$. It follows that the policy maker strictly prefers \underline{x} to any $r \in (\underline{x}, r^*)$ and r^* to any $r > r^*$. For any $\theta < \theta^*$, $R(\theta, A(\theta, \underline{x}), \underline{x}) < 0$ and $U(\theta, 0, r^*) < 0$, which implies that any $\theta < \theta^*$ finds it optimal to set \underline{x} and then face regime change. On the contrary, any $\theta > \theta^*$ necessarily maintains the status quo, since setting r^* guarantees that $R(\theta, 0, r^*) > 0$ and $U(\theta, 0, r^*) > 0$. By definition of S , $\Theta^{**} = S(x^*)$ is then the set of types $\theta \geq \theta^*$ who prefer raising the policy at r^* to setting \underline{x} . We thus conclude that $r(\theta) = r^*$ if $\theta \in \Theta^{**}$ and $r(\theta) = \underline{x}$ otherwise is indeed optimal for the policy maker.

Next, consider the agents. Given $r = \underline{x}$, beliefs are necessarily pinned down by Bayes' rule since any x is consistent with either $(-\infty, \theta^*)$ or $(\theta^*, +\infty) \setminus \Theta^{**}$; this is immediate in the case of unbounded noise and is ensured by the fact that $\Theta^{**} \subseteq [\theta^*, \theta^* + 2\sigma]$ in the case of bounded noise. The posterior probability of regime change is then given by

$$\mu(\theta^*|x, \underline{x}) = \frac{1 - \Psi\left(\frac{x - \theta^*}{\sigma}\right)}{1 - \int_{\Theta^{**}} \frac{1}{\sigma} \psi\left(\frac{x - \theta^*}{\sigma}\right) d\theta} = m(x^*, x)$$

and is decreasing in x . Moreover, by definition of x^* , either $x^* < +\infty$ and $m(x^*, x^*) = g(\underline{x})$ or $x^* = +\infty$, in which case the probability of regime change is $m(x^*, x) \geq g(\underline{x})$ for all x . Hence, given \underline{x} , it is indeed optimal to attack if and only if $x < x^*$. When instead $r = r^*$, Bayes' rule implies $\mu(\theta^*|x, r^*) = 0$ for any x such that $\Theta(x) \cap \Theta^{**} \neq \emptyset$, in which case it is optimal not to attack.

For out-of-equilibrium events, we follow a construction similar to that in proposition 2. The set of types for whom a deviation to $r \notin \{\underline{x}, r^*\}$ is dominated in equilibrium is $\bar{\Theta}(r) = (-\infty, 0] \cup \{\theta \geq 0 : U(\theta, 0, r) < U^*(\theta)\}$, where $U^*(\theta) = \max\{0, U(\theta, A(\theta, r(\theta)), r(\theta))\}$ denotes the equilibrium payoff for type θ . For any $r \in (\underline{x}, r^*)$, in which case $[\theta^*, \hat{\theta}] \not\subseteq \bar{\Theta}(r)$, take any μ such that $\mu(\hat{\theta}|x, r) > g(r)$ if and only if $x < X(r)$. If $x < X(r)$, we also restrict $\mu(\theta \in \bar{\Theta}(r)|x, r) = 0$ when $\Theta(x) \not\subseteq \bar{\Theta}(r)$. If instead $x \geq X(r)$, we do not need to impose such a restriction, for it would only make the agents more aggressive and hence the deviation even less profitable. Finally, when either $r > r^*$ or $r = r^*$ and $\Theta(x) \cap \Theta^{**} = \emptyset$, necessarily $\Theta(x) \subseteq \bar{\Theta}(r)$. Take then any beliefs such that $\mu(0|x, r) = 0$ for $x \geq \underline{x}$ and $\mu(0|x, r) = 1$ otherwise. Given these beliefs and the definition of X , the strategy of the agents is sequentially rational. QED

References

- Angeletos, George-Marios, Christian Hellwig, and Alessandro Pavan. 2003. "Coordination and Policy Traps." Working Paper no. 9767 (June), NBER, Cambridge, MA.
- Angeletos, George-Marios, and Iván Werning. 2006. "Crises and Prices: Information Aggregation, Multiplicity and Volatility." *A.E.R.* 96 (forthcoming).
- Atkeson, Andrew. 2001. "Discussion." In *NBER Macroeconomics Annual 2000*, edited by Ben S. Bernanke and Kenneth Rogoff. Cambridge, MA: MIT Press.
- Calvo, Guillermo A. 1988. "Servicing the Public Debt: The Role of Expectations." *A.E.R.* 78 (September): 647–61.
- Carlsson, Hans, and Eric van Damme. 1993. "Global Games and Equilibrium Selection." *Econometrica* 61 (September): 989–1018.
- Chamley, Christophe. 1999. "Coordinating Regime Switches." *Q.J.E.* 114 (August): 869–905.
- Cho, In-Koo, and David M. Kreps. 1987. "Signaling Games and Stable Equilibria." *Q.J.E.* 102 (May): 179–221.
- Cole, Harold L., and Timothy J. Kehoe. 2000. "Self-Fulfilling Debt Crises." *Rev. Econ. Studies* 67 (January): 91–116.
- Cooper, Russell, and Andrew John. 1988. "Coordinating Coordination Failures in Keynesian Models." *Q.J.E.* 103 (August): 441–63.
- Corsetti, Giancarlo, Bernardo Guimaraes, and Nouriel Roubini. 2006. "International Lending of Last Resort and Moral Hazard: A Model of IMF's Catalytic Finance." *J. Monetary Econ.* 53 (April): 441–76.
- Dasgupta, Amil. Forthcoming. "Coordination and Delay in Global Games." *J. Econ. Theory*.
- Diamond, Douglas W., and Philip H. Dybvig. 1983. "Bank Runs, Deposit Insurance, and Liquidity." *J.P.E.* 91 (June): 401–19.
- Edmond, Chris. 2005. "Information Manipulation, Coordination and Regime Change." Manuscript, New York Univ., Stern School Bus.
- Goldstein, Itay, and Ady Pauzner. 2005. "Demand-Deposit Contracts and the Probability of Bank Runs." *J. Finance* 60 (June): 1293–1327.
- Hellwig, Christian. 2002. "Public Information, Private Information, and the Multiplicity of Equilibria in Coordination Games." *J. Econ. Theory* 107 (December): 191–222.
- Hellwig, Christian, Arijit Mukherji, and Aleh Tsyvinski. 2006. "Self-Fulfilling Currency Crises: The Role of Interest Rates." *A.E.R.* 96 (forthcoming).
- Katz, Michael L., and Carl Shapiro. 1986. "Technology Adoption in the Presence of Network Externalities." *J.P.E.* 94 (August): 822–41.
- Morris, Stephen, and Hyun Song Shin. 1998. "Unique Equilibrium in a Model of Self-Fulfilling Currency Attacks." *A.E.R.* 88 (June): 587–97.
- . 2001. "Rethinking Multiple Equilibria in Macroeconomic Modeling." In *NBER Macroeconomics Annual 2000*, edited by Ben S. Bernanke and Kenneth Rogoff. Cambridge, MA: MIT Press.
- . 2003. "Global Games—Theory and Applications." In *Advances in Economics and Econometrics: Theory and Applications*, 8th World Congress of the Econometric Society, edited by Mathias Dewatripont, Lars Peter Hansen, and Stephen J. Turnovsky. Cambridge: Cambridge Univ. Press.
- . 2004. "Liquidity Black Holes." *Rev. Finance* 8 (1): 1–18.
- Obstfeld, Maurice. 1986. "Rational and Self-Fulfilling Balance-of-Payments Crises." *A.E.R.* 76 (March): 72–81.
- . 1996. "Models of Currency Crises with Self-Fulfilling Features." *European Econ. Rev.* 40 (April): 1037–47.

Rochet, Jean-Charles, and Xavier Vives. 2004. "Coordination Failures and the Lender of Last Resort: Was Bagehot Right after All?" *J. European Econ. Assoc.* 2 (December): 1116–47.