

Persuasion in Global Games with Application to Stress Testing

Supplement

Nicolas Inostroza
Northwestern University

Alessandro Pavan
Northwestern University and CEPR

January 24, 2018

Abstract

This document contains additional results for the manuscript “Persuasion in Global Games with Application to Stress Testing”. All numbered items (i.e., sections, subsections, lemmas, conditions, propositions, and equations) in this document contain the prefix “S”. Any numbered reference without the prefix “S” refers to an item in the main text. Please refer to the main text for notation and definitions.

Section S1 extends Theorem 1 to a fairly general class of economies in which (a) agents’ prior beliefs need not be consistent with a common prior, nor be generated by signals drawn independently conditionally on θ , (b) the number of agents is arbitrary (in particular, finitely many agents), (c) payoffs can be heterogenous across agents, (d) agents have a level-K degree of sophistication, (e) the policy maker need not be able to distinguish perfectly the state of nature (i.e., may possess imperfect information about the payoff state and/or the agents’ beliefs). Section S2 contains the proof of Property 1 used in the main text to establish Lemma 1 in the Appendix. Section S3 contains the proof of Example 2 in the main text showing the suboptimality of monotone policies when Condition 1 in the main text is violated.

Section S1

The model in Section 2 in the main text is modified as follows.

Agents and exogenous information. Let N denote the set of agents; N is assumed to be measurable and can be finite or infinite. For each $i \in N$, let X_i denote a measurable set and define $\mathcal{X} \equiv \prod_{i \in N} X_i$. The set \mathcal{X} is endowed with the product topology. For each $i \in N$, let $\Lambda_i : X_i \rightarrow \Delta(\Theta \times \mathcal{X})$ be a measurable function (with respect to the Borel sigma-algebra associated with X_i). The profile $\mathbf{x} = (x_i)_{i \in N} \in \mathcal{X}$ indexes the hierarchy of the agents' exogenous beliefs about θ and the beliefs of other agents.

The state of Nature in this environment is denoted by $\omega = (\theta, \mathbf{x}) \in \Omega \equiv \Theta \times \mathcal{X}$ and comprises the realization of the payoff fundamental θ and the exogenous profile of the agents' beliefs \mathbf{x} . Note that no restriction on the agents' belief profile \mathbf{x} is imposed. In particular, the agents' beliefs need not be consistent with a common prior, nor be generated by signals drawn independently conditionally on θ .

Payoffs. Each agent's payoff from attacking is normalized to zero, whereas the agents' payoff from not attacking are given by

$$u_i(\theta, A) = \begin{cases} g_i(\theta, A) & \text{if } R(\theta, A) = 0 \\ b_i(\theta, A) & \text{if } R(\theta, A) = 1, \end{cases}$$

$i \in N$, where A denotes the aggregate size of attack (in case of finitely many agents, A coincides with the number of agents attacking). The functions g_i and b_i are continuously differentiable and satisfy the same monotonicity assumptions as in the main text. That is, for any $i \in N$, any (θ, A) : (a) $\frac{\partial}{\partial \theta} g_i(\theta, A), \frac{\partial}{\partial \theta} b_i(\theta, A) \geq 0$, (b) $\frac{\partial}{\partial A} g_i(\theta, A), \frac{\partial}{\partial A} b_i(\theta, A) \leq 0$; and (c) $g_i(\theta, A) > 0 > b_i(\theta, A)$.

Disclosure Policies. Let \mathcal{S} be a compact metric space defining the set of possible disclosures to the agents. Let $m : N \rightarrow \mathcal{S}$ denote a *message function*, specifying, for each individual $i \in N$, the endogenous signal $m_i \in \mathcal{S}$ disclosed to the individual. As in the main text, let $M(\mathcal{S})$ denote the set of all possible message functions with codomain \mathcal{S} . Let \mathcal{P} be a partition of Ω and $h(\omega)$ the information set (equivalently, the cell) in \mathcal{P} containing the state $\omega \in \Omega$. A *disclosure policy* $\Gamma = (\mathcal{S}, \mathcal{P}, \pi)$ consists of a set \mathcal{S} along with a mapping $\pi : \Omega \rightarrow \Delta(M(\mathcal{S}))$ measurable with respect to the σ -algebra defined by the partition \mathcal{P} .¹ For each ω , $\pi(\omega)$ denotes the lottery whose realization yields the message function used by the policy maker to communicate with the agents. The case in which the partition \mathcal{P} coincides with Ω corresponds to the case in which the policy maker is able to distinguish any two states in Ω (in this case the σ -algebra associated with Ω is the Borel σ -algebra).

¹That is, by the collection of \mathcal{P} -saturated sets. Let \mathcal{B} be the standard Borel σ -algebra associated with the primitive set Ω . A set $A \in \mathcal{B}$ is \mathcal{P} -saturated if $\omega \in A$ implies $h(\omega) \subseteq A$. Thus $A = \cup_{\omega \in A} h(\omega)$.

Solution Concept. Agents have a level- K degree of sophistication. The policy maker adopts a conservative approach and evaluates the performance of any given policy on the basis of the “worse outcome” consistent with the agents playing (interim correlated) level- K rationalizable strategies. That is, for any given selected policy Γ , the policy maker expects the market to play according to the “most aggressive level- K rationalizable profile” defined as follows:

Definition S1. Given any policy Γ , any $K \in \mathbb{N} \cup \{\infty\}$, the most aggressive level- K rationalizable profile (MARP- K) associated with Γ is the strategy profile $a_{(K)}^\Gamma \equiv (a_{(K),i}^\Gamma)_{i \in [0,1]}$ that minimizes the policy maker’s ex-ante expected payoff, among all profiles surviving K rounds of *iterated deletion of interim strictly dominated strategies*.

Hereafter we use IDISDS to refer to the process of iterated deletion of interim strictly dominated strategies.

Definition S2. A policy $\Gamma = (\mathcal{S}, \mathcal{P}, \pi)$ satisfies the **perfect-coordination property (PCP)** if, for any $\omega \in \Omega$, any message function $m \in \text{supp}[\pi(\omega)]$, any $i, j \in N$, $a_{(K),i}^\Gamma(x_i, m_i) = a_{(K),j}^\Gamma(x_j, m_j)$.

Fix an arbitrary policy $\Gamma = (\mathcal{S}, \mathcal{P}, \pi)$. For any $\omega \in \Omega$, any message function $m \in \text{supp}[\pi(\omega)]$, let $r(\omega, m; a_{(K)}^\Gamma) \in \{0, 1\}$ denote the regime outcome that prevails at ω when the distribution of endogenous signals is m , and agents play according to the strategy profile $a_{(K)}^\Gamma$.

Definition S3. The disclosure policy $\Gamma = (\mathcal{S}, \mathcal{P}, \pi)$ is **regular** if for any $\omega', \omega'' \in \Omega$ for which $h(\omega') = h(\omega'')$ and any $m \in \text{supp}[\pi(\omega')] = \text{supp}[\pi(\omega'')]$, $r(\omega', m; a_{(K)}^\Gamma) = r(\omega'', m; a_{(K)}^\Gamma)$.

A disclosure policy is thus regular if the regime outcome induced by MARP- K compatible with Γ is measurable with respect to the policy maker’s information (as captured by the partition \mathcal{P}). With an abuse of notation, hereafter we will occasionally denote by $r(h(\omega), m; a_{(K)}^\Gamma) \in \{0, 1\}$ the regime outcome that prevails *at any state in* $h(\omega)$ under the message function m , when we find it useful to highlight the measurability restriction implied by the regularity of the policy. Observe that, when the policy maker can perfectly distinguish between any two states, then any policy is regular.

Theorem S1. *For any regular policy Γ , there exists another regular policy Γ^* satisfying PCP and yielding the policy maker an expected payoff weakly higher than Γ .*

Proof. Let $\mathcal{A}^\Gamma \equiv \{(a_i(\cdot) : X_i \times \mathcal{S} \rightarrow [0, 1])_{i \in N}\}$ denote the entire set of strategy profiles in the continuation game among the agents that starts with the policy maker announcing the policy Γ . For any $n \in \mathbb{N}$, let $T_{(n)}^\Gamma$ denote the set of strategies surviving n rounds of IDISDS under the original policy Γ , with $T_{(0)}^\Gamma = \mathcal{A}^\Gamma$. Denote by $\bar{a}_{(n)}^\Gamma \equiv (\bar{a}_{(n),i}^\Gamma(\cdot))_{i \in [0,1]} \in T_{(n)}^\Gamma$ the profile in $T_{(n)}^\Gamma$ that minimizes the policy maker’s ex-ante payoff. Such a profile also minimizes the policy maker’s interim payoff, as it will become clear from the arguments below. Hereafter, we refer to the profile $\bar{a}_{(n)}^\Gamma$ as the most aggressive

profile surviving n rounds of IDISDS. The profiles $(\bar{a}_{(n)}^\Gamma)_{n \in \mathbb{N}}$ can be constructed inductively as follows. The profile $\bar{a}_{(0)}^\Gamma \equiv (\bar{a}_{(0),i}^\Gamma(\cdot))_{i \in [0,1]}$ prescribes that all agents attack irrespective of their exogenous and endogenous signals; that is, each $\bar{a}_{(0),i}^\Gamma$ is such that $\bar{a}_{(0),i}^\Gamma(x_i, s) = 1$, all $(x_i, s) \in X_i \times \mathcal{S}$.² Given any strategy profile $a \in \mathcal{A}^\Gamma$, any $i \in N$, let $U_i^\Gamma(x_i, m_i; a)$ denote the payoff that agent i with exogenous signal x_i and endogenous signal m_i obtains from not attacking, when all other agents follow the behavior specified by the strategy profile a . For any $n \geq 1$, the most aggressive strategy profile surviving n rounds of IDISDS is the one specifying, for each agent i , each $(x_i, m_i) \in X_i \times \mathcal{S}$, $\bar{a}_{(n),i}^\Gamma(x_i, m_i) = 0$ if $U_i^\Gamma(x_i, m_i; \bar{a}_{(n-1)}^\Gamma) > 0$ and $\bar{a}_{(n),i}^\Gamma(x_i, m_i) = 1$ if $U_i^\Gamma(x_i, m_i; \bar{a}_{(n-1)}^\Gamma) \leq 0$. The most aggressive level- K rationalizable strategy profile (MARPK) consistent with the policy Γ is thus the profile $\bar{a}_{(K)}^\Gamma = (\bar{a}_{(K),i}^\Gamma(\cdot))_{i \in N} \in T_K^\Gamma$. The case of fully rational agents in the main text corresponds to the limit in which $K \rightarrow \infty$. To be consistent with the notation in the main text, we denote MARPK consistent with Γ by dropping the subscript K and denoting such profile by $\bar{a}^\Gamma \equiv ((\bar{a}_i^\Gamma(\cdot))_{i \in N})$, with $\bar{a}_i^\Gamma(\cdot) \equiv \lim_{K \rightarrow \infty} \bar{a}_{(K),i}^\Gamma(\cdot)$, all $i \in N$.

Now, consider the policy $\Gamma^+ = (\mathcal{S}^+, \mathcal{P}, \pi^+)$, $\mathcal{S}^+ = \mathcal{S} \times \{0, 1\}$, obtained from the original policy Γ by replacing each message function $m : N \rightarrow \mathcal{S}$ in the support of each $\pi(\omega)$ with the message function $m^+ : N \rightarrow \mathcal{S}^+$ that discloses to each agent $i \in N$ the same message m_i disclosed by the original policy m , along with the regime outcome $r(\omega, m; \bar{a}_{(K)}^\Gamma)$ that would have prevailed at (ω, m) under Γ when all agents play according to the most aggressive level- K rationalizable strategy profile $\bar{a}_{(K)}^\Gamma$ consistent with the original policy Γ . That is, for each $\omega \in \Omega$, each $m \in \text{supp}[\pi(\omega)]$, the policy Γ^+ selects the message function m^+ obtained from the original message function m by adding to its codomain the regime outcome $r(\omega, m; \bar{a}_{(K)}^\Gamma)$ that would have prevailed at (ω, m) under MARPK $\bar{a}_{(K)}^\Gamma$, with the same probability that Γ would have selected the original message function m . Hereafter, we denote by $m_i^+ = (m_i, r(\omega, m; \bar{a}_{(K)}^\Gamma))$ the message sent to agent i under the new policy Γ^+ when the exogenous state is ω and the message function selected under the original policy Γ is m . Note that the assumption that Γ is regular implies that Γ^+ is measurable with respect to the σ -algebra generated by \mathcal{P} and hence also regular.

Now let $\mathcal{A}^{\Gamma^+} \equiv \{(a_i(\cdot) : X_i \times \mathcal{S} \times \{0, 1\}) \rightarrow [0, 1]\}_{i \in N}$ denote the set of strategy profiles in the continuation game among the agents that starts with the policy maker announcing the new policy Γ^+ . For any $n \in \mathbb{N}$, let $T_{(n)}^{\Gamma^+} \subset \mathcal{A}^{\Gamma^+}$ denote the set of strategies surviving n rounds of IDISDS under the new policy Γ^+ , with $T_{(0)}^{\Gamma^+} = \mathcal{A}^{\Gamma^+}$. Denote by $\bar{a}_{(n)}^{\Gamma^+} \equiv (\bar{a}_{(n),i}^{\Gamma^+}(\cdot))_{i \in N} \in T_{(n)}^{\Gamma^+}$ the profile in $T_{(n)}^{\Gamma^+}$ that minimizes the policy maker's ex-ante payoff, and observe that $\bar{a}_{(0)}^{\Gamma^+} \equiv (\bar{a}_{(0),i}^{\Gamma^+}(\cdot))_{i \in N}$ prescribes that all agents attack irrespective of their exogenous and endogenous signals.

Step 1. First, we prove that, for any $i \in N$,

$$\begin{aligned} & \{(x_i, m_i) \in X_i \times \mathcal{S} : U_i^\Gamma(x_i, m_i; a) > 0 \forall a \in \mathcal{A}^\Gamma\} \\ & \subseteq \{(x_i, m_i) \in X_i \times \mathcal{S} : U_i^{\Gamma^+}(x_i, (m_i, 0); a) > 0 \forall a \in \mathcal{A}^{\Gamma^+}\}. \end{aligned}$$

²Note that, to ease the notation, we let each individual strategy prescribe an action for all $(x_i, m_i) \in \mathbb{R} \times \mathcal{S}$, including those that may be inconsistent with the policy Γ .

That is, any agent who finds it dominant to refrain from attacking under Γ after receiving information (x_i, m_i) also finds it dominant to refrain from attacking under Γ^+ after receiving information $(x_i, (m_i, 0))$. To see this, first use the fact that the game is supermodular to observe that, given any policy Γ ,

$$\{(x_i, m_i) \in X_i \times \mathcal{S} : U_i^\Gamma(x_i, m_i; a) > 0 \forall a \in \mathcal{A}^\Gamma\} = \{(x_i, m_i) \in X_i \times \mathcal{S} : U_i^\Gamma(x_i, m_i; \bar{a}_{(0)}^\Gamma) > 0\}.$$

Likewise,

$$\{(x_i, m_i) \in X_i \times \mathcal{S} : U_i^{\Gamma^+}(x_i, (m_i, 0); a) > 0 \forall a \in \mathcal{A}^{\Gamma^+}\} = \{(x_i, m_i) \in X_i \times \mathcal{S} : U_i^{\Gamma^+}(x_i, (m_i, 0); \bar{a}_{(0)}^{\Gamma^+}) > 0\}.$$

Next, observe that, because under both $\bar{a}_{(0)}^\Gamma$ and $\bar{a}_{(0)}^{\Gamma^+}$ all agents attack, regardless of their exogenous and endogenous information, under both $\bar{a}_{(0)}^\Gamma$ and $\bar{a}_{(0)}^{\Gamma^+}$, regime change occurs if, and only if, $\theta \leq \bar{\theta}$. Then, note that, under Γ^+ , for any $i \in N$, any $(x_i, m_i) \in X_i \times \mathcal{S}$,

$$\partial \Lambda_i^{\Gamma^+}(\omega, m | x_i, (m_i, 0)) = \frac{\mathbb{I}_{\{r(\omega, m; \bar{a}_{(K)}^\Gamma) = 0\}}}{\pi_i^\Gamma(0 | x_i, m_i)} \partial \Lambda_i^\Gamma(\omega, m | x_i, m_i), \quad (\text{S1})$$

where

$$\pi_i^\Gamma(0 | x_i, m_i) \equiv \int_{\{(\omega, m) : r(\omega, m; \bar{a}_{(K)}^\Gamma) = 0\}} d\Lambda_i^\Gamma(\omega, m | x_i, m_i)$$

is the total probability that, under the policy Γ , agent i with information (x_i, m_i) assigns to the event $\{(\omega, m) \in \Omega \times M(\mathcal{S}) : r(\omega, m; \bar{a}_{(K)}^\Gamma) = 0\}$. Under Bayesian learning, the agents' beliefs under the new policy policy Γ^+ thus correspond to “truncations” of their beliefs under the original policy Γ . In turn, this property of Bayesian updating implies that, for any $(x_i, m_i) \in X_i \times \mathcal{S}$ such that

$$U_i^\Gamma(x_i, m_i; \bar{a}_{(0)}^\Gamma) = \int_{(\omega, m)} \left(b_i(\theta, 1) \mathbb{I}_{\{\theta \leq \bar{\theta}\}} + g_i(\theta, 1) \mathbb{I}_{\{\theta > \bar{\theta}\}} \right) d\Lambda_i^\Gamma(\omega, m | x_i, m_i) > 0,$$

it must be that

$$\begin{aligned} U_i^{\Gamma^+}(x_i, (m_i, 0); \bar{a}_{(0)}^{\Gamma^+}) &= \frac{1}{\pi_i^\Gamma(0 | x_i, m_i)} \int_{(\omega, m)} \left(b_i(\theta, 1) \mathbb{I}_{\{\theta \leq \bar{\theta}\}} + g_i(\theta, 1) \mathbb{I}_{\{\theta > \bar{\theta}\}} \right) \times \\ &\quad \times \mathbb{I}_{\{r(\omega, m; \bar{a}_{(K)}^\Gamma) = 0\}} d\Lambda_i^\Gamma(\omega, m | x_i, m_i) \\ &> \frac{1}{\pi_i^\Gamma(0 | x_i, m_i)} \int_{(\omega, m)} \left(b_i(\theta, 1) \mathbb{I}_{\{\theta \leq \bar{\theta}\}} + g_i(\theta, 1) \mathbb{I}_{\{\theta > \bar{\theta}\}} \right) d\Lambda_i^\Gamma(\omega, m | x_i, m_i) \\ &= \frac{1}{\pi_i^\Gamma(0 | x_i, m_i)} U_i^\Gamma(x_i, m_i; \bar{a}_{(0)}^\Gamma) \\ &> 0, \end{aligned}$$

where the first equality follows from the truncation property of Bayesian updating, the first inequality from the fact that, for all $(\omega, m) \in \Omega \times M(\mathcal{S})$ such that $r(\omega, m; \bar{a}_{(K)}^\Gamma) = 1$, $\theta \leq \bar{\theta}$, and hence $r(\omega, m; \bar{a}_{(0)}^\Gamma) = 1$, implying that

$$b_i(\theta, 1) \mathbb{I}_{\{\theta \leq \bar{\theta}\}} + g_i(\theta, 1) \mathbb{I}_{\{\theta > \bar{\theta}\}} = b_i(\theta, 1) < 0,$$

the second equality from the definition of $U_i^\Gamma(x_i, m_i; \bar{a}_{(0)}^\Gamma)$, and the second inequality from the fact that $U_i^\Gamma(x_i, m_i; \bar{a}_{(0)}^\Gamma) > 0$.

The above result implies that any an agent who, under Γ , finds it dominant not to attack after receiving information (x_i, m_i) also finds it dominant not to attack under Γ^+ after receiving information $(x_i, (m_i, 0))$, as claimed.

Step 2. We now show that a property analogous to the one established in Step 1 applies to any other round of the IDISDS procedure. The result is established by induction. Take any round $n \in \{1, 2, \dots, K\}$ and assume that, for any $0 \leq k \leq n - 1$, any $i \in [0, 1]$,

$$\begin{aligned} & \{(x_i, m_i) \in X_i \times \mathcal{S} : U_i^\Gamma(x_i, m_i; a) > 0 \quad \forall a \in T_{(k-1)}^\Gamma\} \\ & \subseteq \{(x_i, m_i) \in X_i \times \mathcal{S} : U_i^{\Gamma^+}(x_i, (m_i, 0); a) > 0, \quad \forall a \in T_{(k-1)}^{\Gamma^+}\}. \end{aligned} \tag{S2}$$

Recall that this means that any agent who, under Γ , finds it optimal not to attack when his opponents play *any* strategy surviving k rounds of IDISDS under Γ continues to find it optimal not to attack when expecting his opponents to play *any* strategy surviving k rounds of IDISDS under Γ^+ . Below we show that that the same property extends to strategies surviving n rounds of IDISDS. That is,

$$\begin{aligned} & \{(x_i, m_i) \in X_i \times \mathcal{S} : U_i^\Gamma(x_i, m_i; a) > 0 \quad \forall a \in T_{(n-1)}^\Gamma\} \\ & \subseteq \{(x_i, m_i) \in X_i \times \mathcal{S} : U_i^{\Gamma^+}(x_i, (m_i, 0); a) > 0, \quad \forall a \in T_{(n-1)}^{\Gamma^+}\}. \end{aligned} \tag{S3}$$

To see this, use again the fact that the game is supermodular, to observe that

$$\{(x_i, m_i) \in X_i \times \mathcal{S} : U_i^\Gamma(x_i, m_i; a) > 0 \quad \forall a \in T_{(n-1)}^\Gamma\} = \{(x_i, m_i) \in X_i \times \mathcal{S} : U_i^\Gamma(x_i, m_i; \bar{a}_{(n-1)}^\Gamma) > 0\} \tag{S4}$$

and, likewise,

$$\begin{aligned} & \{(x_i, m_i) \in X_i \times \mathcal{S} : U_i^{\Gamma^+}(x_i, (m_i, 0); a) > 0, \quad \forall a \in T_{(n-1)}^{\Gamma^+}\} \\ & = \{(x_i, m_i) \in X_i \times \mathcal{S} : U_i^{\Gamma^+}(x_i, m_i; \bar{a}_{(n-1)}^{\Gamma^+}) > 0\}, \end{aligned} \tag{S5}$$

where recall that $\bar{a}_{(n-1)}^\Gamma$ (alternatively, $\bar{a}_{(n-1)}^{\Gamma^+}$) is the most aggressive profile surviving $n - 1 < K$ rounds of IDISDS under Γ (alternatively, Γ^+).

Now let $A(\omega, m; a)$ denote the aggregate size of attack that, under Γ , prevails at (ω, m) , when agents play according to $a \in \mathcal{A}^\Gamma$. Then take any $i \in N$ and any $(x_i, m_i) \in X_i \times \mathcal{S}$ such that

$$U_i^\Gamma(x_i, m_i; \bar{a}_{(n-1)}^\Gamma) = \int_{(\omega, m)} u_i(\theta, A(\omega, m; \bar{a}_{(n-1)}^\Gamma)) d\Lambda_i^\Gamma(\omega, m | x_i, m_i) > 0.$$

Because $\bar{a}_{(n-1)}^\Gamma$ is more aggressive than $\bar{a}_{(K)}^\Gamma$, in the sense that, for any $i \in N$, any $(x_i, m_i) \in X_i \times \mathcal{S}$, $\bar{a}_{(n-1), i}^\Gamma(x_i, m_i) \geq \bar{a}_{K, i}^\Gamma(x_i, m_i)$, then for all (ω, m) ,

$$r(\omega, m; \bar{a}_{(K)}^\Gamma) = 1 \Rightarrow r(\omega, m; \bar{a}_{(n-1)}^\Gamma) = 1.$$

This implies that

$$\begin{aligned} \int_{(\omega, m)} u_i(\theta, A(\omega, m; \bar{a}_{(n-1)}^\Gamma)) \mathbb{I}_{\{r(\omega, m; \bar{a}_{(K)}^\Gamma) = 1\}} d\Lambda_i^\Gamma(\omega, m | x_i, m_i) &= \\ \int_{(\omega, m)} b_i(\theta, A(\omega, m; \bar{a}_{(n-1)}^\Gamma)) \mathbb{I}_{\{r(\omega, m; \bar{a}_{(K)}^\Gamma) = 1\}} d\Lambda_i^\Gamma(\omega, m | x_i, m_i) &< 0 \end{aligned} \quad (\text{S6})$$

This observation, together with the truncation property in (S1), imply that, for any $i \in N$, any $(x_i, m_i) \in X_i \times \mathcal{S}$ such that $U_i^\Gamma(x_i, m_i; \bar{a}_{(n-1)}^\Gamma) > 0$,

$$\begin{aligned} U_i^{\Gamma+}(x_i, (m_i, 0); \bar{a}_{(n-1)}^\Gamma) &= \int_{(\omega, m)} u_i(\theta, A(\omega, m; \bar{a}_{(n-1)}^\Gamma)) d\Lambda_i^{\Gamma+}(\omega, m | x_i, m_i) \\ &= \frac{1}{\pi_i^\Gamma(0 | x_i, m_i)} \int_{(\omega, m)} u(\theta, A(\omega, m; \bar{a}_{(n-1)}^\Gamma)) \mathbb{I}_{\{r(\omega, m; \bar{a}_{(K)}^\Gamma) = 0\}} d\Lambda_i^\Gamma(\omega, m | x_i, m_i) \\ &> \frac{1}{\pi_i^\Gamma(0 | x_i, m_i)} \int_{(\omega, m)} u(\theta, A(\omega, m; \bar{a}_{(n-1)}^\Gamma)) d\Lambda_i^\Gamma(\omega, m | x_i, m_i) \\ &= \frac{1}{\pi_i^\Gamma(0 | x_i, m_i)} U_i^\Gamma((x_i, m_i); \bar{a}_{(n-1)}^\Gamma) \\ &> 0 \end{aligned} \quad (\text{S7})$$

where the first and third equalities are by definition, the second equality follows from (S1), the first inequality follows from (S6), and the last inequality from the fact that $U_i^\Gamma(x_i, m_i; \bar{a}_{(n-1)}^\Gamma) > 0$, by assumption.

Next, note that $\bar{a}_{(n-1)}^\Gamma$ and $\bar{a}_{(n-1)}^{\Gamma+}$ are such that, for all $i \in N$, all $(x_i, m_i) \in X_i \times \mathcal{S}$, $\bar{a}_{(n-1),i}^\Gamma(x_i, m_i)$, $\bar{a}_{(n-1),i}^{\Gamma+}(x_i, (m_i, 0)) \in \{0, 1\}$ and

$$\{(x_i, m_i) \in X_i \times \mathcal{S} : \bar{a}_{(n-1),i}^\Gamma(x_i, m_i) = 0\} = \{(x_i, m_i) \in X_i \times \mathcal{S} : U_i^\Gamma(x_i, m_i; \bar{a}_{(n-2)}^\Gamma) > 0\}$$

and, likewise,

$$\{(x_i, m_i) \in X_i \times \mathcal{S} : \bar{a}_{(n-1),i}^{\Gamma+}(x_i, (m_i, 0)) = 0\} = \{(x_i, m_i) \in X_i \times \mathcal{S} : U_i^{\Gamma+}(x_i, (m_i, 0); \bar{a}_{(n-2)}^{\Gamma+}) > 0\}.$$

Together properties (S2), (S4) and (S5) imply that $\bar{a}_{(n-1)}^\Gamma$ and $\bar{a}_{(n-1)}^{\Gamma+}$ are such that, for all $i \in N$, all $(x_i, m_i) \in X_i \times \mathcal{S}$,

$$\bar{a}_{(n-1),i}^\Gamma(x_i, m_i) = 0 \Rightarrow \bar{a}_{(n-1),i}^{\Gamma+}(x_i, (m_i, 0)) = 0. \quad (\text{S8})$$

Condition (S8), along with the fact that the game is supermodular, implies that

$$U_i^{\Gamma+}(x_i, (m_i, 0); \bar{a}_{(n-1)}^\Gamma) > 0 \Rightarrow U_i^{\Gamma+}(x_i, (m_i, 0); \bar{a}_{(n-1)}^{\Gamma+}) > 0. \quad (\text{S9})$$

Together (S7) and (S9) imply the property in (S3).

Step 3. Equipped with the results in steps 1 and 2 above, we now prove that, for all $i \in N$, all $(x_i, m_i) \in X_i \times \mathcal{S}$,

$$\bar{a}_{(K),i}^{\Gamma+}(x_i, (m_i, 0)) = 0.$$

This follows directly from the fact that, for all $i \in N$, all $(x_i, m_i) \in X_i \times \mathcal{S}$,

$$\bar{a}_{(K),i}^\Gamma(x_i, m_i) = 0 \Rightarrow \bar{a}_{(K),i}^{\Gamma+}(x_i, (m_i, 0)) = 0, \quad (\text{S10})$$

which, in turn implies that, for any (ω, m) ,

$$r(\omega, m; \bar{a}_{(K)}^\Gamma) = 0 \Rightarrow r(\omega, m; \bar{a}_{(K)}^{\Gamma^+}) = 0.$$

Under Γ^+ , the announcement that $r = 0$ thus reveals to the agents that (ω, m) is such that $r(\omega, m; \bar{a}_{(K)}^{\Gamma^+}) = 0$. Because the payoff from refraining from attacking is strictly positive when the regime survives, any agent i receiving a signal $(m_i, 0)$ thus necessarily refrains from attacking. Under MARP-K consistent with the new policy Γ^+ thus all agents refrain from attacking, regardless of their exogenous and endogenous private signals, when the policy publicly announces $r = 0$. That they all attack, irrespective of the (x_i, m_i) , when the policy announces $r = 1$ follows from the fact that $r = 1$ makes it common certainty among the agents that (ω, m) is such that $r(\omega, m; \bar{a}_{(K)}^\Gamma) = 1$ and hence that $\theta \leq \bar{\theta}$. But then, irrespective of (x_i, m_i) , any agent $i \in N$ receiving exogenous information x_i and endogenous information $m_i^+ = (m_i, 1)$ finds it optimal to attack when expecting all other agents to attack no matter their exogenous and endogenous information. This implies that under MARP-K consistent with the new policy Γ^+ , all agents attack when hearing that $r = 1$.

We conclude that the new policy Γ^+ satisfies the perfect coordination property. That such a policy improves upon the original policy Γ follows from the fact that, for any $(\theta, \mathbf{x}) \in \Theta \times \mathcal{X}$, the probability of regime change under Γ^+ is the same as under Γ , but, in case the status quo survives, the aggregate attack is smaller under Γ^+ than under Γ . Q.E.D. \square

Section S2

Property 1. *Assume the function $g : \mathbb{R}^2 \rightarrow \mathbb{R}_+$ is log-supermodular in (x, θ) and the real-valued function $h : \mathbb{R} \rightarrow \mathbb{R}$ crosses 0 only once from below at $\theta = \theta_0$. Choose any (Lebesgue) measurable subset $\Omega \subseteq \mathbb{R}$ and, for any $x \in \mathbb{R}$, let $\Psi(x; \Omega) \equiv \int_\Omega h(\theta)g(x, \theta)d\theta$. Suppose there exists $x^* \in \mathbb{R}$ such that $\Psi(x^*; \Omega) = 0$. Then, necessarily $\Psi(x; \Omega) \geq 0$ for all $x > x^*$, and $\Psi(x; \Omega) \leq 0$ for all $x < x^*$, with both inequalities strict if $\Omega \neq \{\theta : h(\theta) = 0\}$, Ω has strict positive Lebesgue measure, and g is strictly positive and strictly supermodular.*

Proof of Property 1. Pick any $x > x^*$. Observe that

$$\begin{aligned} \Psi(x; \Omega) &= \int_{\Omega \cap (-\infty, \theta_0)} h(\theta)g(x, \theta)d\theta + \int_{\Omega \cap (\theta_0, \infty)} h(\theta)g(x, \theta)d\theta \\ &= \int_{\Omega \cap (-\infty, \theta_0)} h(\theta)g(x^*, \theta) \frac{g(x, \theta)}{g(x^*, \theta)} d\theta + \int_{\Omega \cap (\theta_0, \infty)} h(\theta)g(x^*, \theta) \frac{g(x, \theta)}{g(x^*, \theta)} d\theta \\ &\geq \frac{g(x, \theta_0)}{g(x^*, \theta_0)} \left(\int_{\Omega \cap (-\infty, \theta_0)} h(\theta)g(x^*, \theta)d\theta + \int_{\Omega \cap (\theta_0, \infty)} h(\theta)g(x^*, \theta)d\theta \right) \\ &= \frac{g(x, \theta_0)}{g(x^*, \theta_0)} \Psi(x^*; \Omega), \\ &= 0 \end{aligned}$$

where the inequality follows from the fact that $\frac{g(x,\theta)}{g(x^*,\theta)}$ is increasing in θ as a consequence of the fact that $g(x,\theta)$ is log-supermodular. Similar arguments imply that, for $x < x^*$, $\Psi(x;\Omega) \leq 0$. The same arguments imply that the two inequalities are strict if $\Omega \neq \{\theta : h(\theta) = 0\}$ and g is strictly positive and strictly supermodular. Q.E.D

Section S3

Example 2 in the main text. Suppose (a) the policy maker's payoff is equal to L in case of regime change and W in case of no regime change, independently of (θ, A) , (b) regime change occurs if, and only if, $A \geq \theta$, (c) agents' exogenous signals are given by $x_i = \theta + \sigma \epsilon_i$ where $\sigma \in \mathbb{R}_+$ and where each ϵ_i is drawn from a standard Normal distribution, independently across agents and independently from θ , (d) the agents' payoff are equal to $g = 1 - c$ in case of no regime change and $b = -c$ in case of regime change, with $c > 1/2$. Then, for σ small, Condition 1 is violated and the optimal non-discriminatory policy is not monotone.

Proof of Example 2. Let Φ denotes the cdf of the standard Normal distribution and ϕ its density. Let $x_\sigma^*(\theta) \equiv \theta + \sigma \Phi^{-1}(\theta)$ denote the exogenous signal threshold such that, when all agents attack for $x_i < x_\sigma^*(\theta)$ and refrain from attacking for $x_i > x_\sigma^*(\theta)$, regime change occurs when the fundamentals fall below θ and does not occur when they are above θ . Let $\psi(\theta_0, \hat{\theta}, \sigma)$ denote the payoff from not attacking of an agent with exogenous signal $x_\sigma^*(\theta_0)$, when all other agents follow a cut-off rule with threshold $x_\sigma^*(\theta_0)$, public information makes it common certainty among the agents that $\theta \geq \hat{\theta}$, and the precision of exogenous private information is σ^{-2} . Then let $\theta_\sigma^{\text{inf}} \equiv \inf \left\{ \hat{\theta} : \psi(\theta_0, \hat{\theta}, \sigma) > 0 \text{ all } \theta_0 \in [0, 1] \right\}$. When the policy publicly reveals that $\theta > \hat{\theta} > \theta_\sigma^{\text{inf}}$, in the continuation equilibrium where the agents must choose whether or not to attack, there is a unique rationalizable strategy profile and is such that all agents refrain from attacking. When, instead, the policy publicly reveals that $\theta \geq \theta_\sigma^{\text{inf}}$, there is a rationalizable strategy profile in which each agent $i \in [0, 1]$ attacks if, and only if, $x_i \leq x_\sigma^*(\theta_\sigma^\#)$, with the threshold $\theta_\sigma^\#$ implicitly defined by $\psi(\theta_\sigma^\#, \theta_\sigma^{\text{inf}}, \sigma) = 0$. Note that $\theta_\sigma^\#$ is the fundamentals threshold such that, when it is common certainty that $\theta \geq \theta_\sigma^{\text{inf}}$ and that regime change occurs if, and only if, $\theta \leq \theta_\sigma^\#$, the marginal agent with signal $x_\sigma^*(\theta_\sigma^\#)$ is just indifferent between attacking and not attacking. Angeletos et al (2007) proved that $\lim_{\sigma \rightarrow 0^+} \theta_\sigma^\# = c > 1/2$. By continuity, for σ sufficiently small, $x_\sigma^*(\theta_\sigma^\#) = \theta_\sigma^\# + \sigma \Phi^{-1}(\theta_\sigma^\#) > \theta_\sigma^\# > \theta_\sigma^{\text{inf}}$. As a result, for σ sufficiently small,

$$Y(\theta; x_\sigma^*(\theta_\sigma^\#)) \equiv \frac{\Delta^P(\theta)}{p(x_\sigma^*(\theta_\sigma^\#)|\theta)|b(\theta, P(x_\sigma^*(\theta_\sigma^\#)|\theta))|} = \frac{\sigma(W - L)}{\phi\left(\frac{x_\sigma^*(\theta_\sigma^\#) - \theta}{\sigma}\right)c}$$

is strictly decreasing over $[\underline{\theta}, \theta_\sigma^{\text{inf}}]$, implying that property 3 in Condition 1 is violated.

Now recall that the optimal monotone policy is given by $\Gamma^{\text{Mon}} = \{\{0, 1\}, \pi^{\text{Mon}}\}$ with $\pi^{\text{Mon}}(0|\theta) = \mathbb{I}\{\theta > \theta_\sigma^{\text{inf}}\}$. We show that, starting from the optimal monotone policy Γ^{Mon} , one can construct a

non-monotone policy Γ that yields the policy maker a payoff strictly higher than Γ . To see this, fix σ small, take some arbitrary small value of $\epsilon > 0$, let $\delta(\epsilon)$ be implicitly defined by

$$c \left(\int_0^{\delta(\epsilon)} \phi(x_\sigma^*(\theta_\sigma^\#)|\theta) dF(\theta) + \int_{\theta_\sigma^{\text{inf}} + \epsilon}^{\theta_\sigma^\#} \phi(x_\sigma^*(\theta_\sigma^\#)|\theta) dF(\theta) \right) = (1 - c) \int_{\theta_\sigma^\#}^\infty \phi(x_\sigma^*(\theta_\sigma^\#)|\theta) dF(\theta), \quad (\text{S11})$$

and consider the non-monotone policy $\Gamma^\epsilon = \{\{0, 1\}, \pi^\epsilon\}$ defined by $\pi^\epsilon(0|\theta) = \mathbb{I}\{\theta \in [0, \delta(\epsilon)] \cup (\theta_\sigma^{\text{inf}} + \epsilon, \infty)\}$. Note that $\delta(\epsilon)$ is such that, under the policy Γ^ϵ constructed above, the payoff $U^{\Gamma^\epsilon}(x_\sigma^*(\theta_\sigma^\#), 0|x_\sigma^*(\theta_\sigma^\#))$ that an agent with signal equal to $x_\sigma^*(\theta_\sigma^\#)$ expects from not attacking when all other agents attack if and only if their signal falls below $x_\sigma^*(\theta_\sigma^\#)$ and when the public announcement that $s = 0$ reveals that $\theta \in [0, \delta(\epsilon)] \cup (\theta_\sigma^{\text{inf}} + \epsilon, \infty)$ is exactly equal to zero.

Clearly, $\delta(0) = 0$. Furthermore, using the implicit function theorem and then taking the limit as $\epsilon \rightarrow 0$, we have that

$$\lim_{\epsilon \rightarrow 0^+} \delta'(\epsilon) = \frac{\phi(x_\sigma^*(\theta_\sigma^\#)|\theta_\sigma^{\text{inf}})f(\theta_\sigma^{\text{inf}})}{\phi(x_\sigma^*(\theta_\sigma^\#)|0)f(0)}.$$

Next, observe that, for σ small, the payoff for the policy maker under the policy Γ^ϵ is equal to

$$U_P^{\Gamma^\epsilon} = L \left(F(0) + F(\theta_\sigma^{\text{inf}} + \epsilon) - F(\delta(\epsilon)) \right) + W \left(F(\delta(\epsilon)) - F(0) + 1 - F(\theta_\sigma^{\text{inf}} + \epsilon) \right).$$

Hence,

$$\begin{aligned} \lim_{\epsilon \rightarrow 0^+} \left(\frac{d}{d\epsilon} U_P^{\Gamma^\epsilon} \right) &= (W - L) \left(f(0) \lim_{\epsilon \rightarrow 0^+} \delta'(\epsilon) - f(\theta_\sigma^{\text{inf}}) \right) \\ &= \frac{f(\theta_\sigma^{\text{inf}})(W - L)}{\phi(x_\sigma^*(\theta_\sigma^\#)|0)} \left(\phi(x_\sigma^*(\theta_\sigma^\#)|\theta_\sigma^{\text{inf}}) - \phi(x_\sigma^*(\theta_\sigma^\#)|0) \right) > 0. \end{aligned}$$

The last property implies that the optimal monotone policy Γ^{Mon} can be improved upon by the non-monotone policy Γ^ϵ , for any $\epsilon > 0$ sufficiently small, thus establishing the suboptimality of monotone policies. Q.E.D.

Angeletos, G.-M., Hellwig, C., Pavan, A., 2007. Dynamic global games of regime change: Learning, multiplicity, and the timing of attacks. *Econometrica* 75 (3), 711–756.