

Persuasion in Global Games with Application to Stress Testing*

Nicolas Inostroza
Northwestern University

Alessandro Pavan
Northwestern University and CEPR

April 5, 2017

PRELIMINARY AND INCOMPLETE

Abstract

We study information design in global games of regime change. We consider both the case in which the designer is constrained to disclose the same information to all market participants, as well as the case in which discriminatory disclosures are possible. In both cases, we show that the optimal policy has the “*perfect coordination property*”: it coordinates all market participants on the same course of action. Importantly, while the optimal policy removes any “strategic uncertainty,” it preserves (and in some cases, it enhances) heterogeneity in “structural uncertainty”. Under the optimal policy, each agent can perfectly predict the actions of any other agent, but not the beliefs that rationalize such actions. Preserving heterogeneity in structural uncertainty is key to minimizing the risk of regime change. When the policy maker is constrained to public disclosures, the optimal policy takes the form of a simple “pass/fail” test. More generally, it has a “divide-and-conquer” flavor: It combines a pass/fail public announcement with discriminatory disclosures that enhance the dispersion of beliefs among market participants about the underlying economic fundamentals. Lastly, we identify primitive conditions under which the optimal test is monotone i.e., it fails with certainty institutions with weak fundamentals and passes those with strong ones.

JEL classification: D83, G28, G33.

Keywords: Global Games, Bayesian Persuasion, Information Design, Stress Tests.

*Emails: nicolasinostroza2018@u.northwestern.edu, alepavan@northwestern.edu. We are grateful to Laura Doval, Steve Morris, Jean Tirole, and Xavier Vives for helpful comments and suggestions. The usual disclaimer applies.

1 Introduction

Coordination plays a major role in many socio-economic environments. The damages to society of mis-coordination can be severe and call for government intervention. Think of a major financial institution such as MPS (Monte dei Paschi di Siena, the oldest bank on the planet) trying to convince its creditors to refrain from pulling their money out of the troubled bank in response to rumors about the size of the bank’s non-performing loans. The bank may be financially sound, but faced with a large run, it is forced to liquidate most of its long-term investments. A default by a large financial institution such as MPS, in turn, may trigger a sequence of “domino effects,” leading to a freeze in credit, a collapse in financial markets, and ultimately a deep recession in the entire Eurozone (The Economist, July 7, 2017).

Confronted with such prospects, a government has incentives to intervene. However, a government’s ability to calm the market by injecting liquidity into the troubled bank can be limited. Regulations passed in 2015 prevent Eurozone member states from rescuing banks by purchasing toxic assets or, more generally, by acting on banks’ balance sheets. In such situations, a government’s last resort often takes the form of interventions aimed at influencing market beliefs, for example through the design of stress tests, or other targeted information disclosures. The questions the government faces are then (a) What type of disclosures minimize the risk of coordination failures? (b) Should all the information collected through the stress test be passed on to the market, or should the government commit to a coarser policy, for example one that simply announces whether or not the bank under scrutiny passed the test? (c) Should the government be specific about the level or recapitalization asked to the bank, or simply announce that the bank needs further recapitalization, leaving it to the market to figure out the details? (d) Are there benefits from discriminatory disclosures, whereby different pieces of information are disclosed to different groups of market participants?

In this paper, we develop a framework that permits us to investigate the above questions. We study the design of optimal information disclosures in markets in which a large number of agents must choose whether to play a socially desirable action (e.g., roll over their loans), or “speculate” against a status quo regime (e.g., pull the money out of the troubled bank). Market participants are endowed with heterogenous private information about relevant economic fundamentals, such as the size of the bank’s non-performing loans. A cash-constrained policy maker (e.g., a benevolent government) can act on the agents’ information (for example, by designing a stress test), but does not possess any other financial instrument to influence the market outcome.

While motivated by the design of stress tests, we abstain from many institutional details, and, instead, cast the analysis in a broader class of games of regime change that can be used to shed light on similar questions also in other applications. For example, in the context of currency crises, the policy maker may represent a central bank attempting to convince speculators to refrain from short-selling the domestic currency by releasing information about the bank’s reserves and/or about domestic economic fundamentals. Alternatively, the policy maker may represent the owners of an intellectual

property, or more broadly the sponsors of an idea, choosing among different certifiers in the attempt to persuade heterogenous market users (buyers, developers, or other technology adopters) of the merits of a new product, as in Lerner and Tirole (2006)’s analysis of forum shopping.

The key novelty relative to the rest of the persuasion literature is that we explicitly account for the role that coordination plays among the receivers.¹ Furthermore, the latter are allowed to possess heterogenous private information prior to receiving additional information from the designer. At the theoretical level, these properties imply that, to derive the optimal persuasion strategy, one needs to study the effects of information disclosure not just on the agents’ first-order beliefs, but also on their *higher-order beliefs* (that is, the agents’ beliefs about other agents’ beliefs, their beliefs about other agents’ beliefs about their own beliefs, and so on). Equivalently, the optimal policy must be derived by accounting for how different information disclosures affect both the agents’ *structural uncertainty* (i.e., their beliefs about the underlying fundamentals), as well as the agents’ *strategic uncertainty* (i.e., the agents’ beliefs about other agents’ behavior).

The backbone of the analysis is a flexible global game of regime change in which, prior to receiving information from the policy maker (the information designer), each agent is endowed with an exogenous private signal about the strength of the regime (the critical size of attack above which the status quo collapses). In the absence of additional information, such a game admits a unique rationalizable strategy profile, whereby agents attack if, and only if, they assign sufficiently high probability to the underlying fundamentals being weak, and whereby regime change occurs only for sufficiently weak fundamentals.²

We take a “robust approach” to the design of the optimal information structure. We assume that, when multiple rationalizable strategy profiles are consistent with the disclosed information, the policy maker expects the agents to play according to the “most aggressive” strategy profile (the one that minimizes the policy maker’s payoff over the entire set of rationalizable profiles). This is an important departure from both the mechanism design and the persuasion literature, where the designer is typically assumed to be able to coordinate the market on her most preferred continuation equilibrium. Given the type of applications the analysis is meant for, such “robust approach” appears more appropriate.³

¹For models of persuasion with a single receiver, see, among others, Calzolari and Pavan (2006b), Kamenica and Gentzkow (2011), Bergemann et al. (2015), Kolotilin et al. (2016), Ely (2017), and Mensch (2015). For models with multiple receivers, see, among others, Calzolari and Pavan (2006b), Bergemann and Morris (2013), Chan et al. (2016), Alonso and Camara (2015), Bardhi and Guo (2016), Bergemann and Morris (2016), Goldstein and Huang (2016), Mathevet et al. (2016), Taneva (2016). We refer the reader to Bergemann and Morris (2017) for an excellent overview of this literature.

²Games of regime change have been used to model, among other things, currency crises, debt crises, political change, and standards adoption. See, among others, Morris and Shin (2006), Angeletos et al. (2006, 2007), and Angeletos and Pavan (2013) for earlier references, and Szkup and Trevino (2015), Yang (2015), Denti (2015), and Morris and Yang (2016), for recent developments.

³If the designer could choose the continuation equilibrium, she would fully disclose the fundamentals, and then recommend that all agents refrain from attacking, unless the regime is bound to collapse irrespective of the agents’

Our first result shows that the optimal policy has the “*perfect coordination property*.” It induces all market participants to take the same action, despite heterogeneity in the agents’ first- and higher-order beliefs. In other words, the optimal policy completely removes any strategic uncertainty, while retaining heterogeneity in structural uncertainty. Under the optimal policy, each agent is able to perfectly predict the actions of any other agent, but not the beliefs that rationalize such actions. In particular, an agent who expects all other agents to refrain from attacking need not be able to predict whether most other agents do so because it is dominant for them not to attack or simply because they expect others to refrain from attacking. Such residual heterogeneity in structural uncertainty is key to minimizing the probability of regime change, irrespective of equilibrium selection.

Our second result identifies primitive conditions under which the optimal policy can be implemented by a simple “pass/fail” test that passes with certainty all institutions whose fundamentals are strong and fails, with certainty, all institutions whose fundamentals are weak. In the context of stress test design, the government simply announces whether the financial institution under scrutiny is sound (meaning that default can be avoided if the bank succeeds in raising capital according to a pre-specified recapitalization plan), without getting into the details of the institution’s balance sheet.

The above results pertain to situations in which the policy maker is constrained to disclose the same information to all market participants which is empirically the most relevant case. In Section 4, however, we also investigate properties of optimal policies when the designer can disclose different pieces of information to different market participants. We show that the optimality of the perfect-coordination property extends to discriminatory policies. To see this note that starting from any collection of beliefs (formally, from any subset of the universal type space), informing the agents that the realized state (fundamentals and beliefs) is such that regime change does not occur under the *most aggressive rationalizable* strategy profile consistent with the original type space leads to a new hierarchy of beliefs (formally, a new subset of the universal type space) in which all agents refrain from attacking under any rationalizable profile.

While the optimal policy removes any strategic uncertainty, in general, it enhances the dispersion of posterior beliefs about the underlying economic fundamentals. In the case of stress tests, the government may need to combine a public disclosure about the soundness of the financial institution under scrutiny with targeted disclosures of the bank’s balance sheet geared to different groups of investors. Importantly, the optimality of discriminatory disclosures is not a mere consequence of the fact that agents are endowed with heterogenous prior beliefs. It also holds in settings in which market participants are initially symmetrically informed. The intuition is similar to the one for the *divide-and-conquer* strategy in the contracting-with-externalities literature (e.g., Segal (2003)).⁴ The policy maker makes it dominant for certain agents not to attack, and then leverages on such a property by making it *iteratively dominant* for all other agents to refrain from attacking.

behavior. This is both uninteresting and unrealistic.

⁴See also Moriya et al. (2017) for an analysis of the benefits of discriminatory disclosures in team-production problems.

Discriminatory disclosures, while potentially advantageous, are more difficult to sustain in practice than their nondiscriminatory counterparts. We then investigate primitive conditions under which non discriminatory disclosures are optimal. A precise characterization remains elusive. However, preliminary investigations conducted by restricting attention to Gaussian information structures indicate that whether discriminatory disclosures dominate non-discriminatory ones crucially depends on the sensitivity of the agents' payoffs to the underlying fundamentals in case of regime change vis-a-vis in case the status quo is preserved.

The rest of the paper is organized as follows. Below, we wrap up the introduction by discussing briefly the most pertinent literature. Section 2 presents the baseline model. Section 3 studies properties of optimal policies. Section 4 extends the analysis to discriminatory disclosures. Section 5 concludes by discussing venues for future research. Proofs omitted in the text are in the Appendix at the end of the document.

(Most) pertinent literature. The paper is related to different strands of the literature. The first strand is the literature on *information design*. This literature traces back at least to Myerson (1986), who introduced the idea that, in a persuasion setting, the sender can always restrict attention to incentive-compatible private recommendations to the agents. More recent developments of this literature include (Aumann and Maschler, 1995), Calzolari and Pavan (2006a), and Kamenica and Gentzkow (2011). These papers consider persuasion with a single receiver. The case of multiple receivers is less studied. Calzolari and Pavan (2006b) consider an auction setting in which the sender is the initial owner of a good and where the different receivers are bidders in an upstream market who then resell in a downstream market (see also Dworzak (2016) for the analysis of persuasion in other mechanism design environments with aftermarkets). Alonso and Camara (2015) and Bardhi and Guo (2016) consider persuasion in a voting context, whereas Mathevet et al. (2016) and Taneva (2016) study persuasion in more general multi-receiver settings. Importantly, these papers assume that the receivers are homogeneously informed (share a common prior) about the underlying payoff-relevant parameters. Persuasion with ex-ante heterogeneously informed receivers is examined in Bergemann and Morris (2016, a), Kolotilin et al., 2016, Alonso and Camara (2016) and Chan et al. (2016). Bergemann and Morris (2016, a) characterize the set of outcome distributions that can be sustained as a Bayesian Nash Equilibrium under arbitrary information structures consistent with a given common prior. Alonso and Camara (2016) study public persuasion in a context of multiple receivers with heterogeneous priors. Kolotilin et al., 2016 consider a screening environment whereby the designer elicits the agents' private information prior to disclosing further information to them. Chan et al. (2016) study pivotal persuasion in a voting environment similar to the one in Alonso and Camara (2015), but where the sender is allowed to communicate privately with the voters.⁵

The present paper contributes to the persuasion literature by illustrating the benefits of discrim-

⁵Discriminatory persuasion in a voting setting is also examined in Wang (2015). That paper, however, restricts the sender to using conditionally i.i.d. signals.

inatory disclosures in a coordination environment. However, contrary to the works cited above, the benefits of private persuasion do not stem from the possibility of tailoring the information disclosed to each agent to the latter’s prior beliefs, but from the possibility to “divide-and-conquer” the market by increasing the uncertainty each player faces about other players’ beliefs. The approach we follow is also different. Instead of focusing on the sender’s most preferred continuation equilibrium, we adopt a robust-design approach by which the sender expects the rationalizable strategy profile that is worse for her.

The paper also contributes to the literature on regulatory disclosure in the financial system, reviewed in Goldstein and Sapra (2014). Close in spirit to the present paper is the work by Goldstein and Leitner (2015). That paper studies the design of stress tests by a regulator facing a competitive market, where agents hold homogeneous beliefs about the bank’s balance sheet. In contrast, in the present paper, we consider the design of stress tests by a policy maker facing a continuum of investors with heterogenous private beliefs. We also model explicitly the coordination game among the market participants. Bouvard et al. (2015) study a setting similar to ours where a policy maker must choose between transparency (full disclosure) and opacity (no disclosure) but cannot commit to a disclosure policy. They find that a regulator with an informational advantage relative to the market, (a) chooses excess financial opacity and (b) induces a non monotonic regime outcome with respect to the aggregate level of non-performing loans. In contrast, we assume the policy maker can fully commit to her disclosure policy and allow for flexible information structures. Related is also Goldstein and Huang (2016). That paper studies persuasion in a coordination setting similar to ours, but restricts the designer to announcing whether or not the fundamentals fall below a given threshold. We allow for flexible information structures, but then also identify conditions under which monotone persuasion is optimal.

The paper is also related to the literature on global games with endogenous information. Angeletos et al. (2006), and Angeletos and Pavan (2013) consider settings whereby a policy maker, endowed with private information, engages in costly actions to influence the agents’ behavior. Edmond (2013) considers a similar setting but assumes the cost of policy interventions is zero and agents receive noisy signals of the policy maker’s action. Lastly, Angeletos et al. (2007) consider a dynamic model in which agents learn from the accumulation of private signals over time and from the (possibly noisy) observation of past outcomes. The key difference relative to these works is that, in the present paper, we assume the policy maker chooses the disclosure policy prior to observing the underlying fundamentals and fully commits to it.

2 Model

Players and Actions . The economy is populated by a big player, the policy maker, who seeks to influence the fate of a regime, and a (measure-one) continuum of atomistic agents, who must choose whether or not to attack the regime. We index the agents by i and assume they are distributed

uniformly over $[0, 1]$. We denote by $a_i = 1$ the decision by agent $i \in [0, 1]$ to attack, and by $a_i = 0$ the decision by the same agent to not attack. We then denote by $A \in [0, 1]$ the aggregate size of the attack.

Fundamentals. The payoff structure is parameterized by the random variable $\theta \in \mathbb{R}$. This variable parametrizes both the strength of the status quo (i.e., the critical size of the aggregate attack above which the status quo collapses) and the agents’ preferences. We will refer to θ as the “underlying fundamentals.” It is common knowledge that θ is drawn from an absolutely continuous distribution F , with a smooth density f strictly positive over \mathbb{R} , and first and second moment given by μ_θ and σ_θ^2 , respectively.

Exogenous information. Each agent $i \in [0, 1]$ is endowed with a noisy private signal x_i about the underlying fundamentals. Conditional on θ , the signals x_i are i.i.d. draws from the cdf $P(x|\theta)$ with associated density $p(x|\theta)$ *log-supermodular* in (x, θ) .⁶ The cross-sectional distribution of exogenous signals in the population is denoted by $\phi \in \mathbb{R}^{[0,1]}$.

Regime outcome. Let $r \in \{0, 1\}$ denote the regime outcome, with $r = 1$ in case the status quo is abandoned, and $r = 0$ otherwise. Regime change occurs, i.e., $r = 1$, if, and only if, $R(\theta, A) < 0$, where R is a continuous function, strictly increasing in θ , and decreasing in A .

Dominance Regions. There exist thresholds $\underline{\theta}, \bar{\theta} \in \mathbb{R}$ such that $R(\underline{\theta}, 0) = R(\bar{\theta}, 1) = 0$. Irrespective of the size of the attack, the status quo thus collapses when $\theta \leq \underline{\theta}$, and survives when $\theta > \bar{\theta}$.

Payoffs. The policy maker’s payoff is equal to W if $R(\theta, A) \geq 0$, and to $L < W$ if $R(\theta, A) < 0$, where $L, R \in \mathbb{R}$. The agents’ payoff from attacking is normalized to zero, whereas their payoff from not attacking is equal to⁷

$$u(\theta, A) = \begin{cases} g(\theta, A) & \text{if } R(\theta, A) \geq 0 \\ b(\theta, A) & \text{if } R(\theta, A) < 0. \end{cases}$$

The functions g and b are continuously differentiable and satisfy the following assumptions, for any (θ, A) :⁸ (a) $g_\theta(\theta, A), b_\theta(\theta, A) \geq 0$ and $g_A(\theta, A), b_A(\theta, A) \leq 0$; (b) $g(\theta, A) > 0$ and $b(\theta, A) < 0$. In the context of stress-test design, the first assumption means that the payoff that a creditor expects from leaving the money into the bank (weakly) increases with the bank’s size of performing loans (the fundamentals) and with the number of creditors that also keep pledging to the bank $(1 - A)$. The second assumption says that leaving the money into the bank yields a payoff higher than taking the money out from the bank in case default does not occur, whereas the opposite is true in case of default. These assumptions readily extend to other applications.

⁶Log-supermodularity trivially holds when the signals take the familiar additive form $x_i = \theta + \sigma\varepsilon_i$, with $\{\varepsilon_i\}_i$ drawn independently across agents, and independently from θ , from a log-concave distribution p_ε with $\mathbb{E}[\varepsilon] = 0$.

⁷In case of currency attacks and political change, it is customary to normalize the payoff from not attacking to 0. That is, to assume the “safe action” is not-attacking. This can be accommodated by letting $\hat{a} = 1 - a$ and then interpreting $\hat{a} = 1$ (which corresponds to $a_i = 0$) as the decision to attack (the risky action).

⁸The functions $g_\theta(\theta, A), b_\theta(\theta, A)$ are partial derivatives with respect to the θ dimension. Similarly, $g_A(\theta, A), b_A(\theta, A)$ are partial derivatives with respect to A .

Disclosure Policies. The only instrument the policy maker possesses to influence the regime outcome is the design of a disclosure policy. Let \mathcal{S} be a compact metric space defining the set of possible disclosures to the agents. A *disclosure policy* $\Gamma = (\mathcal{S}, \pi)$ consists of a mapping $\pi : \Theta \rightarrow \Delta(\mathcal{S})$ specifying, for each fundamental θ , a probability distribution over the information disclosed to the agents. Note that the formalization here assumes the disclosure policy is non-discriminatory; we consider discriminatory disclosures in Section 4. As is standard in the literature, the disclosure policy Γ itself does not convey any information about θ to the agents (in the context of stress test design the assumption reflects the idea the policy maker does not possess private information about the financial institution under scrutiny prior to conducting the test). Furthermore, the policy maker can credibly commit not to modify Γ once the latter is announced.

Timing. The sequence of events is as follows:

1. The policy maker chooses a disclosure policy $\Gamma = (\mathcal{S}, \pi)$ and publicly announces it.
2. The fundamentals of the economy θ as well as the cross-sectional distribution of exogenous information ϕ in the population are realized. Each agent $i \in [0, 1]$ then privately observes his own signal x_i .
3. A public signal $s \in \text{supp}[\pi(\theta)]$ is drawn from the distribution $\pi(\theta) \in \Delta(\mathcal{S})$ and disclosed to all market participants.
4. Agents simultaneously choose whether or not to attack.
5. The regime outcome is determined by (θ, A) and payoffs are realized.

3 Optimal Policies

In designing her disclosure policy, the policy maker adopts a conservative approach. She evaluates the performance of any given policy on the basis of the “worse outcome” consistent with the agents playing (interim correlated) rationalizable strategies. That is, for any given selected policy Γ , the policy maker expects the market to play according to the “most aggressive rationalizable profile” defined as follows.

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

As it will become clear in a moment, such strategy profile is, in fact, a Bayes-Nash equilibrium (BNE) of the continuation game that follows the announcement of Γ , and minimizes the policy maker’s payoff state-by-state, and not just in expectation.

3.1 Most aggressive rationalizable profile (MARP)

Fix $\Gamma = (\mathcal{S}, \pi)$ and, for any pair (x, s) , let $\Lambda^\Gamma(\theta|x, s)$ represent the endogenous posterior beliefs about θ of each agent receiving exogenous information x and endogenous information s . Next, let

$$U^\Gamma(x, s|k) = \int_{-\infty}^{\infty} u(\theta, P(k|\theta)) d\Lambda^\Gamma(\theta|x, s),$$

denote the payoff from not attacking of an agent observing an exogenous private signal x and an endogenous public signal s , when the rest of the agents follow a cut-off strategy with cut-off k (that is, they attack if, and only if, their private signal falls short of the cut-off k). We then have the following result:

Lemma 1. *Given any policy $\Gamma = (\mathcal{S}, \pi)$, the most aggressive rationalizable strategy profile (MARP) $a^\Gamma \equiv (a_i^\Gamma)_{i \in [0,1]}$ consistent with Γ is such that, for any $s \in \mathcal{S}$, $x \in \mathbb{R}$, $i \in [0, 1]$,*

$$a_i^\Gamma(x, s) = \mathbf{1}\{x \leq \xi_s^\infty\}$$

with

$$\xi_s^\infty = \inf\{x : U^\Gamma(x, s|x) \geq 0\}, \text{ all } s \in \mathcal{S}.$$

Moreover, the strategy profile a^* is a BNE of the continuation game that starts with the announcement of the policy Γ .

The result follows from the fact that the aggregate size of attack at any (θ, s) is weakly higher under the strategy profile $a^\Gamma \equiv (a_i^\Gamma)_{i \in [0,1]}$ than under any other (interim correlated) rationalizable profile. The formal proof in the Appendix combines standard properties of supermodular games with the fact that, when the densities $p(x|\theta)$ are log-supermodular, for any cut-off k , the payoff $U^\Gamma(\cdot, s|k)$ crosses zero only once and from below in x .

3.2 Perfect Coordination Property

Equipped with the result in Lemma 1 above, we now proceed to the characterization of properties of optimal policies.

Definition 2. A policy $\Gamma = \{\mathcal{S}, \pi\}$ satisfies the **perfect-coordination property** if, for any $s \in \mathcal{S}$, there exists $r_s \in \{0, 1\}$ such that $a_i^*(x, s) = r_s$, any $i \in [0, 1]$, any $x \in \mathbb{R}$, where a^* is the most aggressive rationalizable strategy profile consistent with the policy Γ .

Hence, a disclosure policy has the perfect-coordination property if it induces all market participants to follow the same course of action, after any signal it discloses.

Theorem 1. *Given any policy Γ , there exists another policy $\Gamma^* = \{\mathcal{S}^*, \pi^*\}$, with $\mathcal{S}^* = \{0, 1\}$ that yields the policy maker a payoff weakly higher than Γ . The policy $\Gamma^* = \{\mathcal{S}^*, \pi^*\}$ has the **perfect-coordination property**. When signal $s^* = 1$ is disclosed, all agents attack, whereas, when signal*

$s^ = 0$ is disclosed, all agents refrain from attacking, irrespective of their exogenous private information.*

The result is established in the Appendix by showing that, given any policy Γ , there exists a binary policy $\Gamma^* = (\{0, 1\}, \pi^*)$ satisfying the perfect-coordination property that yields the policy maker a payoff weakly higher than Γ . The proof is in two steps.

Step 1 shows that, starting from Γ , one can construct another policy $\hat{\Gamma}$ that, for any θ , discloses the same information s as the original policy Γ , along with the regime outcome that, under Γ , would have prevailed at (θ, s) , had the agents played according to MARP consistent with the original policy Γ . Given the new policy $\hat{\Gamma}$, under the most aggressive strategy profile consistent with $\hat{\Gamma}$, no agent attacks after receiving signal s and hearing that the underlying state θ is such that, at (θ, s) , under the original policy Γ , regime change would have not occurred under MARP consistent with Γ . Likewise, under $\hat{\Gamma}$, any agent attacks, irrespective of x , after receiving signal s and hearing that, under Γ , regime change would have occurred, had the agents played according to MARP consistent with Γ . Note that the policy $\hat{\Gamma}$ so constructed satisfies the perfect-coordination property.

The second step then shows that, starting from $\hat{\Gamma}$, one can drop the original signals s inherited from Γ and disclose only the regime outcome that would have prevailed at each θ under MARP consistent with Γ ; dropping the signals s does not change the agents' behavior. Formally, the proof establishes existence of a binary policy Γ^* that satisfies the perfect-coordination property and is such that (a), for any θ , Γ^* randomizes over only two signals, $s^* = 0$ and $s^* = 1$, (b) when signal $s^* = 1$ is disclosed, all agents attack, whereas when signal $s^* = 0$ is disclosed, all agents refrain from attacking, (c) at any θ , Γ^* discloses signal $s^* = 1$ (alternatively, $s^* = 0$) with the same total probability the original policy Γ would have disclosed signals s that, under MARP consistent with Γ , would have led to regime change (alternatively, to the regime to survive).

The policy Γ^* thus removes any strategic uncertainty. When the signal $s^* = 0$ (alternatively, $s^* = 1$) is disclosed, each agent knows that all other agents refrain from attacking (alternatively, attack), irrespective of their exogenous private information, and finds it optimal to do the same.

Importantly, while the policy Γ^* removes any strategic uncertainty, it preserves heterogeneity in structural uncertainty. No matter the announcement, different agents holds different beliefs about the underlying fundamentals. Preserving heterogeneity in posterior beliefs about θ is key to eliminating the possibility of regime change. In fact, if agents knew the exact fundamentals, under the most aggressive rationalizable profile, they would all attack for any $\theta \leq \bar{\theta}$. The policy Γ^* eliminates the possibility of regime change by leveraging on the fact that, when signal $s^* = 0$ is announced, agents remain uncertain as to whether other agents are refraining from attacking because they find it dominant to do so, or because they expect others to refrain from attacking. As we show in Section 4 below, the same property also explains why discriminatory disclosures may dominate non-discriminatory ones when the primitive heterogeneity in structural beliefs does not minimize the ex-ante probability of regime change.

The proof of Theorem 1 in the Appendix establishes the result in the theorem for a broader class of policy maker's payoffs of the form

$$U^P(\theta, A) = \begin{cases} W(\theta, A) & \text{if } R(\theta, A) \geq 0 \\ L(\theta) & \text{if } R(\theta, A) < 0, \end{cases} \quad (1)$$

with the function W continuously differentiable and satisfying the following properties, for any $(\theta, A) \in \mathbb{R} \times [0, 1]$: (a) $W_A(\theta, A) \leq 0$; (b) $W(\theta, A) - L(\theta, A) > 0$ if $R(\theta, A) > 0$. The first property says that, conditional on the status quo surviving the attack, the payoff to the policy maker decreases (weakly) with the size of the aggregate attack. The second property says that the policy maker would never prefer to see the status quo collapse when it survives.⁹

In the context of stress test design, the assumption that, in case of regime change, L is invariant in A means that, when default occurs, the government is indifferent as to the precise degree of speculation that led the financial institution into bankruptcy. When this is the case, starting from any policy Γ , one can construct another policy Γ^* satisfying the perfect coordination property such that (a), for any θ , the probability of regime change under Γ^* is the same as under Γ , and (b) when regime change does not occur, the size of the attack is zero. That Γ^* improves upon Γ then follows directly from the fact that L is invariant in A and W decreasing in A . The proof follows from the same steps establishing Theorem 1 above and hence is omitted.

We conjecture, but did not prove, that the property that L is invariant in A identifies the maximal domain of payoffs functions for the policy maker over which the perfect coordination property holds.

3.3 Monotone Disclosures

We now turn to the optimality of simple threshold policies. For any (θ, x) , let $B(\theta, x) \equiv b(\theta, P(x|\theta))$ denote the agents' payoff from refraining from attacking when regime change occurs, the fundamentals are θ , and the aggregate size of attack is $A(\theta, x) = P(x|\theta)$.

Condition 1. *The function $B(\theta, x)$ is log-supermodular. Furthermore, for any x , the function $Y(\theta; x) \equiv \Delta^P(\theta)/[p(x|\theta)|B(\theta, x)]$ is nondecreasing over $[\underline{\theta}, \hat{\theta}(x)]$, where $\hat{\theta}(x)$ is the regime threshold when agents follow cut-off strategies with cut-off x (i.e., $\hat{\theta}(x)$ solves $R(\hat{\theta}(x), P(x|\hat{\theta}(x))) = 0$).*

Theorem 2. *Suppose (a) the policy maker's payoff is consistent with the representation in (1) with the payoff differential $\Delta^P(\theta) \equiv W(\theta, 0) - L(\theta)$ nondecreasing in θ , and (b) Condition (1) holds. Given any policy Γ , there exists another policy $\Gamma^* = (\{0, 1\}, \pi^*)$ satisfying the perfect-coordination property that yields the policy maker a payoff weakly higher than Γ . The policy $\Gamma^* = (\{0, 1\}, \pi^*)$ has a threshold structure: There exists $\theta^* \in [\underline{\theta}, \bar{\theta}]$ such that, for all $\theta \leq \theta^*$, $\pi^*(1|\theta) = 1$, whereas, for all $\theta > \theta^*$, $\pi^*(0|\theta) = 1$.*

⁹This second property trivially holds when the fate of the regime is controlled directly by the policy maker, as in certain applications.

That $\Delta^P(\theta)$ is nondecreasing means that the net benefit of moving the economy away from regime change towards a situation in which no agent attacks is monotone in the fundamentals. This condition alone implies that, starting from any non-monotone, and possibly stochastic policy, there exists a deterministic monotone policy that yields the policy maker a higher payoff. In games with a single receiver sharing the same prior as the policy maker, such condition suffices to guarantee the optimality of threshold policies. This, however, is not necessarily the case with multiple receivers with heterogeneous private information. The extra condition in the theorem guarantees the possibility of constructing perturbations of the original policy by swapping the probability of saving the regime from low to high states while also preserving the agents' incentives not to attack when recommended to do so.

In the context of stress test design, the assumption that $\Delta^P(\theta)$ is non-decreasing means that the benefit of convincing investors to maintain their money into the bank increases with the size of the bank's performing loans, or, more generally, with the value and profitability of the bank's assets. Importantly, stronger supermodularity conditions such as those requiring W to be supermodular, and/or, L to be non-decreasing are not essential to the result.

The monotonicity of the optimal disclosure policy contrasts with the results in the literature on signaling in global games (see, e.g., among others, Angeletos et al. (2006), Angeletos and Pavan (2013), and Edmond (2013)). In that literature, intermediate types intervene, whereas the rest pool on the cost-minimizing level. Because interventions are costly, the policy maker chooses to intervene only when (a) the benefit of saving the status quo is large enough to compensate for the cost of the intervention and (b), in the absence of intervention, the size of attack is large enough to make the policy maker prefer the cost of intervention to the cost of facing a large, but unsuccessful, attack.

Under the conditions of Theorem (2), the choice of the optimal policy reduces to the choice of the largest θ^* such that, for all $x \in \mathbb{R}$,

$$\int_{\theta^*}^{\infty} u(\theta, P(x|\theta))p(x|\theta)f(\theta)d\theta > 0.$$

The above problem does not have a formal solution. Notwithstanding these complications, with abuse, hereafter, we refer to the threshold policy Γ^* with cut-off

$$\theta^* \equiv \inf\{\theta' : \int_{\theta'}^{\infty} u(\theta, P(x|\theta))p(x|\theta)f(\theta)d\theta > 0 \text{ for all } x \in \mathbb{R}\}$$

as to the optimal monotone policy.¹⁰

4 Discriminatory Disclosures

We now turn to situations in which the policy maker can disclose different pieces of information to different market participants. In this section, we assume the policy maker's payoff is consistent with

¹⁰That the above problem does not admit a solution was first noticed in Goldstein and Huang (2016).

the representation in (1), and Condition (1) holds. In the case most favorable to the designer, the signals disclosed to the agents may depend both on the realized fundamentals θ as well as on the cross-sectional distribution $\phi \in \mathbb{R}^{[0,1]}$ of exogenous private information in the population (that is, on individual prior beliefs). The discussion below focuses on such a case. In future work, it would be interesting to investigate also intermediate cases in which the designer must solicit information from the agents in case she wants to condition individual disclosures on individual beliefs.

We start by considering a completely unconstrained situation, in which the policy maker can choose any disclosure policy of her choice, and show that the optimal policy continues to satisfy the *perfect coordination property*. We then proceed by discussing the benefits of discriminatory disclosures. Finally, we turn to situations in which the policy maker can disclose arbitrary public signals, but is constrained to Gaussian signals when communicating privately with the agents.

Let $m : [0, 1] \rightarrow \mathcal{S}$ denote a *message function*, specifying, for each individual $i \in [0, 1]$, the endogenous signal $m_i \in \mathcal{S}$ disclosed to the individual. Let $M(\mathcal{S})$ denote the set of all possible message functions with range \mathcal{S} . A *discriminatory disclosure policy* $\Gamma = (\mathcal{S}, \pi)$ consists of a measurable set \mathcal{S} along with a mapping $\pi : \Theta \times \mathbb{R}^{[0,1]} \rightarrow \Delta(M(\mathcal{S}))$ specifying, for each fundamental θ , and each cross-sectional distribution $\phi \in \times \mathbb{R}^{[0,1]}$ of exogenous private information, a lottery whose realization yields the message function used to communicate with the agents.

4.1 On the optimality of the perfect-coordination property

In case of discriminatory disclosures, the definition of the perfect-coordination property is naturally adjusted as follows.

Definition 3. A discriminatory policy $\Gamma = \{\mathcal{S}, \pi\}$ satisfies the **perfect-coordination property** if, for any (θ, ϕ) , any message function $m \in \text{supp}(\pi(\theta, \phi))$, any $i, j \in [0, 1]$, $a_i^\Gamma(x_i, m_i) = a_j^\Gamma(x_j, m_j)$, where $a^\Gamma \equiv (a_i^*)_{i \in [0,1]}$ is the most aggressive rationalizable profile (MARP) consistent with the policy Γ .¹¹

Theorem 3. *Given any discriminatory policy Γ , there exists another policy Γ^* satisfying the **perfect coordination property** that yields the policy maker an expected payoff weakly higher than Γ .*

The proof in the Appendix shares certain similarities with the proof of Theorem 1. Starting from any disclosure policy Γ , one can construct another policy Γ^* that, in addition to the signals disclosed privately to the agents under the original policy Γ , it announces publicly the regime outcome that, at each (θ, ϕ, m) , would have prevailed, under Γ , when agents play according to MARP consistent with the original policy Γ . Relative to the case of non-discriminatory policies, the key difficulty is in establishing the result is that agents' posterior beliefs with and without the extra public piece of

¹¹The most aggressive rationalizable profile continues to be defined as the one that minimizes the policy maker's ex-ante payoff over all rationalizable strategy profiles. Its characterization, however, is significantly more complex than in the case of non-discriminatory policies. In particular, Lemma 1 does not extend to discriminatory policies.

information described above cannot be easily ranked (e.g., according to FOSD). Furthermore, the announcement that the regime would have survived under MARP consistent with the original policy Γ carries information not only about θ , but also about the distribution of first- and higher-order beliefs in the population. In case of non-discriminatory policies, the cross-sectional distribution of the agents' beliefs depends only on θ . This is not necessarily the case under discriminatory policies. We show that the key property inducing the results for the discriminatory case is the truncation of beliefs under the new signal structure Γ^* . That is, conditional on the extra piece of information that regime would have survived under MARP consistent with the Γ , agents assign 0 probability to the collection of states and beliefs across the population that would have induced regime change. In addition, the new signal preserves the likelihood associated to each (θ, ϕ, m) for which regime change would have been spared as under Γ , modulo a common normalization by the probability that the regime would have survived under the former policy. As a result, the most aggressive rationalizable profile under the new signal is less aggressive than the one induced by the former signal without the extra piece of information. Understanding that regime change would have been spared under a more aggressive strategy profile and the fact that this has become commonly understood by all other agents, each receiver optimally decides to refrain from attacking. Similarly, when the signal announces that regime change would have taken place under the former signal Γ , it becomes common certainty among the agents that $\theta < \bar{\theta}$ and therefore the most aggressive profile prescribes that all agents attack the regime. Thus, under Γ^* , all agents coordinate in the same course of action.

4.2 On the benefits of discriminatory disclosures

We now show why, in general, discriminatory disclosures improve upon non-discriminatory ones. Importantly, the optimality of discriminatory disclosures does not come from the possibility of tailoring the information disclosed to each agent to his prior beliefs. To illustrate, consider an economy in which the agents' prior beliefs are homogenous (formally, this amounts to assuming the exogenous private signals x are completely uninformative). Next notice that, for any $\hat{\theta}$ such that

$$\int u(\theta, 1) dF(\theta | \theta > \hat{\theta}) \leq 0,$$

the most aggressive rationalizable strategy profile following the public announcement that $\theta > \hat{\theta}$ is such that every agent attacks.¹² Hereafter, we follow the same convention as in Subsection 3.3 by referring to the optimal non-discriminatory policy as to the threshold policy with cut-off given by¹³

$$\hat{\theta}^* = \inf \{ \hat{\theta} \in \mathbb{R} \text{ s.t. } \int u(\theta, 1) dF(\theta | \theta > \hat{\theta}) > 0 \}. \quad (2)$$

¹²The notation $F(\theta | \theta > \hat{\theta})$ stands for the common posterior obtained from the prior F by conditioning on the event that $\theta > \hat{\theta}$.

¹³Recall that, by virtue of Theorem 2, any non-discriminatory policy Γ can be weakly improved upon by a non-discriminatory policy Γ^* satisfying the perfect coordination property and with a threshold structure. Hence, hereafter, when comparing discriminatory policies to non-discriminatory ones, we restrict attention to non-discriminatory policies satisfying the perfect coordination property and with a threshold structure.

Suppose now the policy maker, in addition to announcing whether θ is above or below some cut-off threshold $\hat{\theta}$, sends to each individual a private signal of the form $m_i = \theta + \sigma\xi_i$, where $\sigma \in \mathbb{R}_+$ is a scalar, and where the idiosyncratic terms (ξ_i) are drawn from a smooth distribution over the entire real line (e.g., a standard Normal distribution), independently across agents, and independently from θ . From standard results in the global games literature, as the private messages become infinitely precise (formally, as $\sigma \rightarrow 0$), in the absence of any public disclosure, under the most aggressive rationalizable profile, all agents attack if, and only if, their endogenous private signals fall below a threshold $\theta^{MS} \in (\underline{\theta}, \bar{\theta})$ that is implicitly defined as the unique solution to¹⁴

$$\int_0^1 u(\theta^{MS}, l) dl = 0. \quad (3)$$

The threshold θ^{MS} corresponds to the highest value of the fundamentals θ for which an agent who knows θ and holds *Laplacian* beliefs with respect to the size of the attack¹⁵ is indifferent between attacking and not attacking. Importantly, θ^{MS} is independent of the initial common prior and of the distribution of the noise terms ξ in the agents' signals. The above result thus implies that, with discriminatory disclosures, the policy maker can always guarantee that regime change never occurs for any $\theta > \theta^{MS}$.

Proposition 1. *Assume the agents possess no exogenous private information about the underlying fundamentals. Let $\hat{\theta}^*$ be the threshold in (2) and θ^{MS} be the threshold in (3). Assume $\theta^{MS} < \hat{\theta}^*$. Then discriminatory policies strictly improve upon non-discriminatory ones.*

The result follows directly from the arguments preceding the proposition. Because $\hat{\theta}^*$ can be arbitrarily close to $\bar{\theta}$ for particular prior distributions, and because θ^{MS} is invariant in the prior distribution from which θ is drawn, the result in Proposition 1 is relevant in many cases of interest.

The reason why discriminatory policies improve upon non-discriminatory ones is that they permit the policy maker to increase the dispersion in the agents' first- and higher-order beliefs about the underlying fundamentals. A higher dispersion in turn makes it difficult for the agents to coordinate on a successful attack. Formally speaking, when beliefs are sufficiently dispersed, an agent receiving a private signal indicating that the regime may collapse under a sufficiently large attack may nonetheless refrain from attacking because he is concerned that many other agents may have received more extreme signals indicating that the fundamentals are strong enough to survive an attack of any size. In this case, refraining from attacking may become *iteratively dominant* for this individual. The optimality of discriminatory policies thus follows from a “divide-and-conquer” logic reminiscent to the one in the vertical contracting literature (see, e.g., Segal (2003) and the references therein). Hence, when discriminatory policies dominate over non-discriminatory ones, this is not because they mis-coordinate market participants on different actions (recall that the optimal policy satisfies the perfect-coordination property), but because they increase heterogeneity in structural uncertainty.

¹⁴See Morris and Shin (2006).

¹⁵This means that the agent believes that the proportion of agents attacking is uniformly distributed over $[0, 1]$.

4.3 Sufficient Conditions for the Optimality of Non-discriminatory Policies

Despite their advantages, discriminatory policies are more difficult to sustain than their non-discriminatory counterparts. It is thus important to identify markets in which non-discriminatory policies are optimal. A complete characterization of such markets is elusive at the moment. However, to get some traction, consider a special environment where the policy maker can engineer any public disclosure of her choice but is constrained to use Gaussian private signals of the form

$$\tilde{m}_i = \theta + \sigma_\xi \xi_i, \text{ with } \xi_i \sim \mathcal{N}(0, 1)$$

when communicating privately with the agents. In each state (θ, ϕ) , the endogenous information $m_i = (\tilde{s}, \tilde{m}_i)$ disclosed to each agent i thus comprises a public signal \tilde{s} , along with the private signal \tilde{m}_i . The quality of the private signals is conveniently parametrized by the variance $\sigma_\xi^2 > 0$ of the noise terms.¹⁶

Further assume that the prior distribution F from which θ is drawn is an improper prior over the entire real line and the agents' exogenous private signals are given by¹⁷

$$x_i = \theta + \sigma_\eta \eta_i, \text{ with } \eta_i \sim \mathcal{N}(0, 1).$$

Also assume the agents' payoff from not attacking is invariant in A , which amounts to assuming that there exist strictly increasing functions $\bar{g}(\theta)$ and $\bar{b}(\theta)$ such that $g(\theta, A) = \bar{g}(\theta)$ and $b(\theta, A) = \bar{b}(\theta)$, all (θ, A) . Finally, assume the function R determining the regime outcome takes the linear form $R(\theta, A) = \theta - A$.¹⁸

Now observe that the information contained in each pair (x_i, \tilde{m}_i) is the same as the information contained in the sufficient statistics

$$z_i \equiv \frac{\sigma_\xi^2 x_i + \sigma_\eta^2 m_i}{\sigma_\eta^2 + \sigma_\xi^2}, \tag{4}$$

which, given θ , is normally distributed with mean θ and variance $\sigma_z^2 \equiv (\sigma_\eta^2 \sigma_\xi^2) / (\sigma_\eta^2 + \sigma_\xi^2)$. Hence, the policy maker's choice of the discriminatory component of her disclosure policy can be conveniently reduced to the choice of the standard deviation σ_z of the sufficient statistics z_i , with $\sigma_z \in (0, \sigma_\eta]$.

Arguments analogous to those establishing Lemma 1 above then imply that, for any realization s of the endogenous public signal, the most aggressive rationalizable strategy profile a^* is characterized

¹⁶As in Section 3, we assume the distribution from which the public signal is drawn is independent of the distribution ϕ describing the exogenous private signal x_i received by each agent i . This is consistent with the idea that public disclosures cannot condition on individual private information.

¹⁷The assumption that F is improper is standard in the global-game literature. It simplifies the formulas below, without any serious effect on the results. Note that, given any proper Gaussian distribution F , when the agents' exogenous signals x_i and Gaussian, and the policy maker is constrained to use Gaussian private messages m_i when communicating privately to the market, the policy Γ^* that maximizes the policy maker's payoff has a threshold structure. Under an improper prior, we thus let the optimal policy be the one under which the regime threshold is the lowest.

¹⁸The results below extend to more general payoff functions, as long as the agents' exogenous signals x are sufficiently precise.

by a unique cut-off $\bar{z}(s)$ (whose value depends on the distribution from which the public signal is drawn) such that, for all $i \in [0, 1]$, $a_i^*(x_i, (s, m_i)) = 1 \{z_i \leq \bar{z}(s)\}$. Moreover, arguments similar to those establishing Theorem 1 above imply that, for any given choice of σ_z^2 , the optimal public message is binary with $s \in \{0, 1\}$. Finally, from Theorem 3, the optimal policy has the perfect-coordination property which means that, given σ_z^2 , $\bar{z}(0) = -\infty$, and $\bar{z}(1) = +\infty$. That is, all agents attack when $s = 1$, and they all refrain from attacking when $s = 0$.

Next, let Φ denotes the cdf of the standard Normal distribution, and define

$$z_{\sigma_z}^*(\theta) = \theta + \sigma_z \Phi^{-1}(\theta), \quad (5)$$

to be the private statistics threshold such that, when all agents attack for $z_i < z_{\sigma_z}^*(\theta)$ and refrain from attacking for $z_i > z_{\sigma_z}^*(\theta)$, regime change occurs when the fundamentals fall below θ and does not occur when they are above θ .¹⁹

Let

$$\psi(\theta_0, \hat{\theta}, \sigma_z)$$

denote the payoff from not attacking of an agent with private statistics $z_{\sigma_z}^*(\theta_0)$, when the total precision of private information is σ_z^{-2} , regime change occurs for all $\theta \leq \theta_0 \in [0, 1]$, and public information reveals that $\theta \geq \hat{\theta}$. Then let

$$\theta_{\sigma_z}^{inf} \equiv \inf \left\{ \hat{\theta} : \psi(\theta_0, \hat{\theta}, \sigma_z) > 0 \text{ all } \theta_0 \right\}.$$

Note that, for any $\hat{\theta} > \theta_{\sigma_z}^{inf}$, under the most aggressive rationalizable strategy profile, no agent attacks after the public signal reveals that $\theta \geq \hat{\theta}$. Hereafter, we assume that all agents refrain from attacking also when public disclosures reveal that $\theta \geq \theta_{\sigma_z}^{inf}$. This simplifies the exposition below by permitting us to talk about the ‘‘optimal policy.’’ As discussed above, the latter does not formally exist when agents are expected to play according to the most aggressive rationalizable profile. However, because the policy maker can always guarantee that, no matter the selection of the rationalizable strategy profile, no agent attacks for any $\theta > \theta_{\sigma_z}^{inf}$, we find the abuse justified.

Proposition 2. *Suppose the policy maker is constrained to use Gaussian private signals when communicating privately with the market. Let*

$$\sigma_z^* \equiv \operatorname{argmin}_{\sigma_z \in (0, \sigma_\eta]} \theta_{\sigma_z}^{inf}$$

The optimal disclosure policy has the following structure. The policy maker publicly announces whether $\theta < \theta_{\sigma_z^}^{inf}$, or whether $\theta \geq \theta_{\sigma_z^*}^{inf}$. In addition, when $\theta \geq \theta_{\sigma_z^*}^{inf}$, the policy maker sends a Gaussian private signal to each agent of precision $\sigma_\xi^{-2} = [\sigma_\eta^2 - (\sigma_z^*)^2]/(\sigma_z^*)^2 \sigma_\eta^2$.*

The result follows from the arguments preceding the proposition – note that the precision of the endogenous private information σ_ξ^{-2} in the proposition is the one that, together with the precision of

¹⁹Given that $R(\theta, A) = \theta - A$, $z_{\sigma_z}^*(\theta)$ is implicitly defined by the equation $\Phi\left(\frac{z^* - \theta}{\sigma_z}\right) = \theta$.

the exogenous signals σ_η^{-2} yields a total precision σ_z^{-2} for the sufficient statistics z_i that minimizes the threshold $\theta_{\sigma_z}^{inf}$.

The value of Proposition 2 is in illustrating how one can use this type of results to identify primitive conditions under which the optimal policy is non-discriminatory. Given the result in Proposition 2, discriminatory disclosures dominate non-discriminatory ones if, and only if, $\sigma_z^* < \sigma_\eta$ (equivalently, if, and only if, there exists $\sigma_z < \sigma_\eta$ such that $\theta_{\sigma_z}^{inf} < \theta_{\sigma_\eta}^{inf}$). For any precision σ_z^{-2} of the agents' private statistics, let $\theta_{\sigma_z}^\#$ denote the unique solution to the equation

$$\psi(\theta_{\sigma_z}^\#, \theta_{\sigma_z}^{inf}, \sigma_z) = 0.$$

Note that $\theta_{\sigma_z}^\#$ identifies the threshold below which regime change occurs when the total precision of the agents' private information is σ_z^{-2} , and the endogenous disclosure of public information reveals that $\theta \geq \theta_{\sigma_z}^{inf}$. Let²⁰

$$D(\theta, \theta_{\sigma_z}^\#) \equiv \begin{cases} \bar{b}'(\theta) & \text{if } \theta < \theta_{\sigma_z}^\# \\ \bar{g}'(\theta) & \text{if } \theta \geq \theta_{\sigma_z}^\#. \end{cases}$$

Proposition 3. *Suppose that, for any $\sigma_z \in [0, \sigma_\eta]$,*

$$\mathbb{E}[D(\theta, \theta_{\sigma_z}^\#)(\theta - \theta_{\sigma_z}^\#) | z^*(\theta_{\sigma_z}^\#), \theta \geq \theta_{\sigma_z}^{inf}] > 0. \quad (6)$$

Then the optimal policy is non-discriminatory.

The condition in Proposition 3 is a measure of the sensitivity of the marginal agent's net payoff from not attacking to the underlying fundamentals.²¹ To see this, note that the condition is equivalent to

$$\frac{\mathbb{E}[\bar{g}'(\theta)(\theta - \theta_{\sigma_z}^\#) | z^*(\theta_{\sigma_z}^\#), \theta \geq \theta_{\sigma_z}^\#]}{\mathbb{E}[\bar{g}(\theta) | z^*(\theta_{\sigma_z}^\#), \theta \geq \theta_{\sigma_z}^\#]} > \frac{\mathbb{E}[\bar{b}'(\theta)(\theta - \theta_{\sigma_z}^\#) | z^*(\theta_{\sigma_z}^\#), \theta \in (\theta_{\sigma_z}^{inf}, \theta_{\sigma_z}^\#)]}{\mathbb{E}[\bar{b}(\theta) | z^*(\theta_{\sigma_z}^\#), \theta \in (\theta_{\sigma_z}^{inf}, \theta_{\sigma_z}^\#)]}.$$

The left-hand side is the elasticity of the marginal agent's expected net payoff from attacking with respect to the underlying fundamentals in case the status quo is preserved, whereas the right-hand side is the corresponding elasticity in case of regime change.²²

To gather some intuition, consider the case in which, when the regime collapses, the payoff differential between not attacking and attacking is constant in the underlying fundamentals (i.e., $\bar{b}'(\theta) = 0$ for all θ). In this case, the marginal agent faces only *upside risk*. Hence, when the quality of private information decreases (which amounts to a mean-preserving increase in risk), the agent's expected net payoff from not attacking increases. Starting from any policy that discloses private information to the agents (i.e., for which $\sigma_z < \sigma_\eta$), the policy maker can then do better by reducing the precision of the agents' private information. In this case, the optimal policy is non-discriminatory.

In the context of stress test design, the above condition holds, for example, when investors are equity holders. In this case, when regime change occurs (i.e., when the bank defaults), their claims are

²⁰Here b' and g' denote the derivatives of the b and g functions, respectively.

²¹See Iachan and Nenov (2015) for a similar condition in a related class of games of regime change.

²²Recall that, for the marginal agent with signal $z^*(\theta_{\sigma_\eta}^\#)$, $\mathbb{E}[\bar{g}(\theta) | z^*(\theta_{\sigma_\eta}^\#), \theta \geq \theta_{\sigma_\eta}^\#] = \mathbb{E}[\bar{b}(\theta) | z^*(\theta_{\sigma_\eta}^\#), \theta \in (\theta_{\sigma_\eta}^{inf}, \theta_{\sigma_\eta}^\#)]$.

junior (i.e., subordinated) with respect to those from other stake holders with higher seniority (e.g., bond holders). In case of default, their payoff then amounts to a liquidation value that is typically little sensitive to the exact amount of the bank's performing loans (the bank's fundamentals). On the contrary, when regime change does not occur (i.e., when the government succeeds in persuading the bank's equity holders to stay put), the value of the equity-holders' claims reflect the bank's long-term profitability, which is sensitive to the amount of the bank's performing loans.

5 Conclusions

The above results thus indicate that, in a large class of games of regime change, the optimal policy completely removes any strategic uncertainty, while retaining, and, in some cases, enhancing structural uncertainty (that is, the dispersion of beliefs about the underlying fundamentals). Under the optimal policy, each agent can perfectly predict the actions of any other agent, but not the beliefs that rationalize such actions. We also identify situations in which the optimal policy can be implemented with a simple deterministic test that fails institutions with weak fundamentals and saves those with strong ones. Finally, we document the benefits of disclosing different information to different market participants, and discuss instances in which public disclosures are optimal.

We conclude with a few lines for future research. The analysis in this paper presumes that the designer knows the distribution of market beliefs when choosing her disclosure policy. Such information may come from polls, data on professional forecasters, the IOWA betting markets, and the like. While this is a useful benchmark, there are many environments in which it is more appropriate to assume that the designer lacks information about market beliefs. In future work, it would be interesting to investigate the optimal disclosure policy in such situations. The idea is to apply a robust (i.e., *undominated* max-min) approach to the designer's problem, whereby the designer expects (a) Nature to select the information structure that minimizes the planner's payoff, and (b) the market to coordinate on the most-aggressive rationalizable strategy profile. The emphasis on the designer's policy to be undominated is key here. Given any disclosure policy, the policy maker's payoff always attains its minimum when the fundamentals are common knowledge among market participants. A policy is undominated if there exists no other policy that yields a weakly higher payoff *across all possible market beliefs*. The characterization of the set of undominated policies is highly relevant both from a theoretical standpoint and for the associated policy implications.

The analysis above is also static. However, many of the applications of interest are intrinsically dynamic, with agents coordinating on multiple attacks over time and learning from past attacks (see the discussion in Angeletos et al. (2007)). In future work, it would be interesting to extend the analysis in this direction. When the fundamentals are partially persistent over time, the optimal policy must specify the timing of information disclosures and how the information at each period depends on the agents' behavior in previous periods. Furthermore, an unsuccessful attack in one period may make the status quo more vulnerable in subsequent periods. There are difficulties in

extending the analysis to dynamic environments, but the returns are worth the effort.

Finally, the analysis in this paper is conducted by assuming that the fundamentals can be observed by the information designer at no cost. When θ represents information that is private to the financial institution under scrutiny, such an assumption need not be appropriate. In future work, it would also be interesting to investigate the problem of a designer who must solicit information from the financial institution prior to communicating with the market. This creates an interesting screening+persuasion problem in the spirit of what examined in the literature on privacy in sequential contacting (e.g., Calzolari and Pavan (2006a,b)).

Appendix

Proof of Lemma 1. We first establish the following property:

Property 1. Assume the function $g(x, \theta)$ with domain \mathbb{R}^2 and codomain \mathbb{R} is log-supermodular in (x, θ) and the real-valued function $h(\theta)$ crosses 0 only once from below. Choose any subset $\Omega \subseteq \mathbb{R}$. Define $\Psi(x; \Omega) \equiv \int_{\Omega} h(\theta)g(x, \theta)d\theta$. Suppose there exists x^* such that $\Psi(x^*; \Omega) = 0$. Then, necessarily

$$\Psi(x; \Omega) > 0 \text{ for all } x > x^*.$$

Proof of Property 1. Let θ_0 denote the unique solution to $h(\theta) = 0$. Pick any $x > x^*$. Observe that:

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

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. \square

Now fix the policy Γ . For any $s \in \mathcal{S}$, then let

$$\xi_s^1 \equiv \inf\{x : U^\Gamma(x, s|\infty) \geq 0\}.$$

Given the public signal s , it is dominant for any agent with private signal x exceeding ξ_s^1 not to attack. This follows from the fact that the expected payoff differential $\int u(\theta, 1)d\Lambda^\Gamma(\theta|x, s)$ between not attacking and attacking crosses 0 only once and from below (in x). The single crossing property of $\int u(\theta, 1)d\Lambda^\Gamma(\theta|x, s)$ in turn is a consequence of Property 1 above. In fact, for any $A \in [0, 1]$, $u(\theta, A)$ crosses 0 only once and from below (in θ), and the density $p(x|\theta)$ is log-supermodular.

Next, let T_s^1 denote the set of strategy profiles that, given the public signal s , survive the first round of IDISDS and denote by $\bar{a} \equiv (\bar{a}_i^1)_{i \in [0, 1]}$ the most aggressive strategy profile in T_s^1 , that is, the strategy profile in T_s^1 that maximizes the ex-ante probability of regime change (clearly, \bar{a} depends not only on s but also on Γ ; we drop the dependence on Γ to ease the notation). Then observe that such a strategy profile is given by

$$\bar{a}_i^1(x, s) = \mathbf{1}\{x \leq \xi_s^1\}, \text{ all } x \in \mathbb{R}, \text{ all } i \in [0, 1],$$

and that such a profile maximizes the probability of regime change, not just in expectation, but state by state, i.e., for any θ .

For any $n > 1$, then let T_s^n denote the set of strategy profiles that survive n rounds of IDISDS and $\bar{a}^n \equiv (\bar{a}_i^n)_{i \in [0,1]}$ the most aggressive strategy profile in T_s^n (that is, the strategy profile in T_s^n that maximizes the ex-ante probability of regime change). That the continuation game is supermodular implies that there exists a unique sequence $\{\xi_s^n\}_n$ such that, for any $n \geq 1$,

$$\bar{a}_i^n(x, s) = \mathbf{1}\{x \leq \xi_s^n\}, \text{ all } x \in \mathbb{R}, \text{ all } i \in [0, 1],$$

with ξ_s^1 as defined above, and with all other cut-offs ξ_s^n , $n > 1$, defined inductively by²³

$$\xi_s^n \equiv \inf\{x : U^\Gamma(x, s | \xi_s^{n-1}) \geq 0\}.$$

Next, let $T_s^\infty \equiv \bigcap_{n=1}^\infty T_s^n$ denote the set of strategy profiles that, given s , are *rationalizable* for the agents given the public signal s . The most aggressive strategy profile in T_s^∞ is then given by

$$\bar{a}_i^\infty(x, s) \equiv \mathbf{1}\{x \leq \xi_s^\infty\}, \text{ all } x \in \mathbb{R}, \text{ all } i \in [0, 1],$$

where $\xi_s^\infty \equiv \lim_{n \rightarrow \infty} \xi_s^n$. The sequence (ξ_s^n) is monotone and its limit is given by

$$\xi_s^\infty = \inf\{x : U^\Gamma(x, s | x) \geq 0\}.$$

Now let T^∞ denote the entire set of rationalizable strategy profile in the continuation game that starts with the policy maker's selection of the policy $\Gamma = (\mathcal{S}, \pi)$.²⁴ Given the properties above, the most aggressive strategy profile in T^∞ is given:

$$a_i^\Gamma(x, s) = \bar{a}_i^\infty(x, s) = \mathbf{1}\{x \leq \xi_s^\infty\}, \quad \text{all } i \in [0, 1], x \in \mathbb{R}, s \in \mathcal{S}.$$

This establishes the first part of the lemma. That the profile a^Γ is a BNE for the continuation game that starts with the announcement of the policy Γ follows from the fact that, given any $s \in \mathcal{S}$, when all agents follow a cut-off strategy with cutoff ξ_s^∞ , each agent $i \in [0, 1]$ finds it optimal to attack for $x_i < \xi_s^\infty$ and to not attack for $x_i > \xi_s^\infty$ (he is indifferent for $x_i = \xi_s^\infty$). Q.E.D.

Proof of Theorem 1. As explained in the main text, the proof below establishes the result in the theorem for a broader class of policy maker's payoffs consistent with the representation in (1). We proceed in two steps. Step 1 shows that, starting from Γ , one can construct another policy $\hat{\Gamma}$ that, for any θ , discloses the same information as the original policy Γ , along with the regime outcome that, under Γ , would have prevailed at θ , when agents play the most aggressive rationalizable strategy profile consistent with the original policy Γ . The second step then shows that, starting from $\hat{\Gamma}$, one can drop the original signals s in Γ and disclose only $r_s \in \{0, 1\}$, while preserving the same regime outcome at each θ .

²³Again, such a characterization follows from the fact that, at any θ , the payoff differential between not attacking and attacking is decreasing in the size of the aggregate attack, along with the fact that the density $p(x|\theta)$ is log-supermodular, which implies that if, for some $z \in \mathbb{R}$, $U^\Gamma(z, s; \xi_s^{n-1}) \leq 0$, then $U^\Gamma(x, s; \xi_s^{n-1}) < 0$ for any $x < z$.

²⁴A strategy profile $a = (a_i(\cdot))_{i \in [0,1]} \in T^\infty$ if, and only if, $(a_i(\cdot, s))_{i \in [0,1]} \in T_s^\infty$, all $s \in \mathcal{S}$.

Step 1. Fix $\Gamma = (\mathcal{S}, \pi)$. For any $\theta \in \mathbb{R}$, any public signal $s \in \text{supp}[\pi(\theta)]$, let $r(\theta, s; \Gamma) \in \{0, 1\}$ denote the (deterministic) regime outcome that prevails at θ when agents play the most aggressive rationalizable strategy profile $a_s^*(\cdot)$ consistent with the policy Γ (as characterized in Lemma 1). Then let $\hat{\Gamma} = (\hat{\mathcal{S}}, \hat{\pi})$ be the policy constructed from Γ by replacing each signal s disclosed at each state θ , with the signal $\hat{s} = (s, r(\theta, s; \Gamma))$. Note that, for any θ , $\hat{\Gamma}$ discloses the signal $\hat{s} = (s, r(\theta, s; \Gamma))$ with the same probability that Γ would have disclosed the signal s . In words, at any θ , the new policy $\hat{\Gamma}$ randomizes over the same signals as the original policy Γ , but then amends each signal realization s by adding to the signal the description of the regime outcome that would have prevailed at (θ, s) under Γ , when the agents play the most aggressive rationalizable profile a^Γ consistent with Γ .

Observe that, for any $x \in \mathbb{R}$, any $s \in \mathcal{S}$, the posterior belief distribution $\Lambda^{\hat{\Gamma}}(\theta|x, (s, 0)) \succeq_{FOSD} \Lambda^\Gamma(\theta|x, s)$, where $\Lambda^\Gamma(\theta|x, s)$ denotes the posterior belief under the original policy Γ , whereas

$$\Lambda^{\hat{\Gamma}}(\theta|x, (s, 0))$$

denotes the posterior belief under the new policy $\hat{\Gamma}$. The result follows from Milgrom (82). In fact, the information in the signal $\hat{s} = (s, 0)$ is either equivalent to the information in the signal s (this occurs when $\{\theta : r(\theta, s; \Gamma) = 0\} = \{\theta : \pi(s|\theta) > 0\}$), or there exists $\hat{\theta} > \underline{\theta}$ such that the extra information in the signal $\hat{s} = (s, 0)$ is equivalent to the announcement that $\theta \geq \hat{\theta}$, which is “good news” in the sense of Milgrom (1982) representation theorems.²⁵

Consider now the set of rationalizable strategy profiles under $\hat{\Gamma}$. For any $n \geq 1$, any $s \in \mathcal{S}$, let $\hat{T}_{(s,0)}^n$ denote the set of strategy profiles surviving n rounds of IDISDS following the announcement of the signal $\hat{s} = (s, 0)$, and $(\hat{a}_i^n(\cdot, (s, 0)))_{i \in [0,1]}$ denote the most aggressive strategy profile in $\hat{T}_{(s,0)}^n$ (the formal definition of the sets $\hat{T}_{(s,0)}^n$ and of the strategy profiles $(\hat{a}_i^n(\cdot, (s, 0)))_{i \in [0,1]}$ is as in the proof of Lemma 1 above). Let $\{\hat{\xi}_s^n\}_n$ denote the sequence of thresholds characterizing the strategy profiles $(\hat{a}_i^n(\cdot, (s, 0)))_{i \in [0,1]}$, as in the proof of Lemma 1. Because, for any $x \in \mathbb{R}$, any $s \in \mathcal{S}$, $\Lambda^{\hat{\Gamma}}(\theta|x, (s, 0)) \succeq_{FOSD} \Lambda^\Gamma(\theta|x, s)$, we have that $\hat{\xi}_{(s,0)}^1 \leq \xi_s^1$. Now, suppose that $\hat{\xi}_{(s,0)}^j \leq \xi_s^j$ for all $j \leq n$. Recall that this means that, for any $j \leq n$, the most aggressive strategy profile $(\hat{a}_i^j(\cdot, (s, 0)))_{i \in [0,1]}$ surviving j rounds of IDISDS under $\hat{\Gamma}$ is less aggressive than the corresponding strategy profile $(\bar{a}_i^j(\cdot, s))_{i \in [0,1]}$ under Γ . This property, together with the fact that $\Lambda^{\hat{\Gamma}}(\theta|x, (s, 0)) \succeq_{FOSD} \Lambda^\Gamma(\theta|x, s)$, implies that, for any $x \in \mathbb{R}$ for which $U^\Gamma(x, s; \xi_s^n) \geq 0$ it is also the case that $U^{\hat{\Gamma}}(x, (s, 0); \hat{\xi}_{(s,0)}^n) \geq 0$. In turn, this implies that $\hat{\xi}_{(s,0)}^{n+1} \leq \xi_s^{n+1}$. Thus,

$$\hat{\xi}_{(s,0)}^\infty \leq \xi_s^\infty. \quad (7)$$

By the way the policy $\hat{\Gamma}$ is constructed, each agent with private signal $x \in \mathbb{R}$ observing the public signal $\hat{s} = (s, 0)$ knows that the underlying state θ is such that, if agents were to follow the most aggressive strategy profile under the original policy Γ , then regime change would not occur. By (7),

²⁵Observe that, by virtue of Lemma 1, given any signal s , the regime outcome is monotone under the most aggressive strategy profile consistent with the original policy Γ .

under the new policy $\hat{\Gamma}$, the most aggressive strategy profile following the announcement $\hat{s} = (s, 0)$ is less aggressive than under the original policy Γ . By definition of ξ_s^∞ , for any $x < \xi_s^\infty$,

$$U^{\hat{\Gamma}}(x, (s, 0); x) \geq U^{\hat{\Gamma}}(x, (s, 0); \xi_s^\infty) > 0.$$

This means that, necessarily, $\hat{\xi}_{(s,0)}^\infty = -\infty$, meaning that no agent attacks after the signal $\hat{s} = (s, 0)$ is disclosed.

It is also easy to see that, given the original policy Γ , for any signal s , the announcement that the fundamentals θ are such that, under the most aggressive rationalizable profile $(\bar{a}_i^\Gamma(\cdot, s))_{i \in [0,1]}$, regime change occurs, makes it common certainty among the agents that $\theta < \bar{\theta}$. It is then easy to see that, given the new policy $\hat{\Gamma}$, under the most aggressive rationalizable profile, all agents attack, irrespective of their private signals x , after any signal $\hat{s} = (s, 1)$ is disclosed.

Because, for any θ , the new policy $\hat{\Gamma}$ sends each signal $(s, r(\theta, s; \Gamma))$ with the same probability the original policy Γ would have sent the signal s , we then have that, under the new policy $\hat{\Gamma}$, for any θ , (a) the probability of regime change under $\hat{\Gamma}$ is the same as under Γ . It is then immediate that the new policy $\hat{\Gamma}$ yields the policy maker the same payoff as the original policy Γ .

Step 2. Next, let

$$S_0(\theta; \hat{\Gamma}) \equiv \{s \in \mathcal{S} : \hat{\pi}((s, 0)|\theta) > 0\}$$

denote the set of signals disclosed with positive probability in conjunction with the announcement of $R = 0$, under the policy $\hat{\Gamma}$ constructed in Step 1, when the fundamentals are θ . Note that, by the way the policy $\hat{\Gamma}$ is constructed, $S_0(\theta; \hat{\Gamma})$ also coincides with the set of signals

$$\{s \in \mathcal{S} : \pi(s|\theta) > 0 \text{ and } r(\theta, s; \Gamma) = 0\}$$

disclosed by the original policy Γ with positive probability at θ which, under the most aggressive rationalizable profile consistent with Γ , would have resulted in no regime change.

Now consider the policy $\Gamma^* = (\{0, 1\}, \pi^*)$ constructed from the policy $\hat{\Gamma} = (\hat{\mathcal{S}}, \hat{\pi})$ by letting $\mathcal{S}^* = \{0, 1\}$ and then letting

$$\pi^*(0|\theta) = \hat{\pi}(S_0(\theta; \hat{\Gamma})|\theta).$$

That is, the policy Γ^* is a stochastic “pass/fail” policy that, for each θ , randomizes over the regime outcome and then perfectly informs market participants of whether or not regime change is to be expected. For each θ , the probability the policy Γ^* recommends to attack coincides with the probability the original policy Γ would have selected signals leading to regime change under the most aggressive strategy profile consistent with Γ .

Next, let $S_0 \equiv \{s \in \mathcal{S} : \exists \theta \text{ s.t. } s \in S_0(\theta; \hat{\Gamma})\}$ denote the set of signals disclosed with positive probability under the policy $\hat{\Gamma}$ at some θ , in conjunction with the announcement of $r = 0$. Again, note that, by the way the policy $\hat{\Gamma}$ is constructed, S_0 also coincides with the set of signals

$$\{s \in \mathcal{S} : \exists \theta \text{ s.t. } \pi(s|\theta) > 0 \text{ and } r(\theta, s; \Gamma) = 0\}$$

disclosed with positive probability at some θ such that, when agents play according to the most aggressive rationalizable strategy profile consistent with Γ , regime change would not have occurred at that θ .

Next, observe that, under Γ^* , for any cutoff k , any private signal x , the payoff $U^{\Gamma^*}(x, 0|k)$ that any agent with private signal x expects from refraining from attacking after the public signal $s^* = 0$ is disclosed, when all other agents follow a cut-off strategy with cut-off k , is equal to

$$U^{\Gamma^*}(x, 0|k) = \int_{S_0} U^{\hat{\Gamma}}(x, (s, 0)|k) dQ^{\hat{\Gamma}}(s|x, 0) \quad (8)$$

where $Q^{\hat{\Gamma}}(\cdot|x, 0)$ is the probability distribution over S_0 obtained by conditioning on the event $(x, 0)$, under the policy $\hat{\Gamma}$. From Step 1, under the policy $\hat{\Gamma}$, for any signal $\hat{s} = (s, 0)$ in the range of $\hat{\pi}$ (equivalently, for any signal $\hat{s} = (s, 0)$ with $s \in S_0$), the set $\hat{T}_{(s,0)}^\infty$ of rationalizable strategy profiles contains only the profile $(a_i^\Gamma(\cdot, (s, 0)))_{i \in [0,1]}$ that prescribes

$$a_i^\Gamma(x, (s, 0)) = 0 \text{ all } x \in \mathbb{R}.$$

This implies that, for all $s \in S_0$, all $k \in \mathbb{R}$, $U^{\hat{\Gamma}}(k, (s, 0)|k) > 0$. From (8), we then have that, for all $k \in \mathbb{R}$,

$$U^{\Gamma^*}(k, 0|k) > 0.$$

In turn, this implies, given the policy Γ^* , when the signal $s^* = 0$ is disclosed, the unique rationalizable profile is such that no agent attacks, which implies that $a_i^{\Gamma^*}(x, 0) = 0$ all x , all $i \in [0, 1]$.

It is also easy to see that, under the most aggressive rationalizable profile consistent with Γ^* , when the signal disclosed is $s^* = 1$, all agents attack. The policy Γ^* so constructed thus (a) satisfies the perfect-coordination property, (b) is such that, at any θ , the probability of regime change under Γ^* is the same as under $\hat{\Gamma}$, and (c), in case of no regime change, the size of attack under Γ^* is zero. The statement in the theorem then follows from the above properties. Q.E.D.

Proof of Theorem 2. Without loss of generality, assume the policy $\Gamma = (\{0, 1\}, \pi)$ satisfies the perfect coordination property and is such that (a) $\pi(0|\theta) = 0$ for all $\theta < \underline{\theta}$, and (b) $\pi(0|\theta) = 1$ for all $\theta > \bar{\theta}$. Note that arguments similar to those establishing Theorem 1 above but adapted to the fact that the policy maker's payoff satisfies the more general structure in (1) imply that, if Γ does not satisfy these properties, there exists another policy Γ' that does satisfy these properties and yields the policy maker a payoff weakly higher than Γ . The proof then follows from applying the arguments below to Γ' instead of Γ .

Clearly, if the original policy $\Gamma = (\{0, 1\}, \pi)$ is such that there exists $\theta^* \in [\underline{\theta}, \bar{\theta}]$ such that $\pi(0|\theta) = 0$ for F -almost all $\theta \leq \theta^*$ and $\pi(0|\theta) = 1$ for F -almost all $\theta \geq \theta^*$, then the threshold policy $\Gamma^* = (\{0, 1\}, \pi^*)$ such that $\pi^*(0|\theta) = 0$ for all $\theta \leq \theta^*$ and $\pi^*(0|\theta) = 1$ for all $\theta \geq \theta^*$ also satisfies the perfect coordination property and yields the policy maker the same payoff as Γ , in which case the result holds.

Suppose, instead, that Γ is such that there exists no θ^* such that $\pi(0|\theta) = 0$ for F -almost all $\theta \leq \theta^*$ and $\pi(0|\theta) = 1$ for F -almost all $\theta \geq \theta^*$. We then establish the result by showing that there exist a threshold policy Γ^* satisfying the perfect coordination property that yields the policy maker's a payoff at least as high as Γ .

Now recall that, for the policy Γ to satisfy the perfect coordination property, it must be that $U^\Gamma(x, 0|x) > 0$ all x . Now let \mathbb{G} denote the set of all policies with range $S = \{0, 1\}$ that, in addition to properties (a) and (b) above, satisfy the additional property that $U^\Gamma(x, 0|x) \geq 0$, all x . Let $\tilde{\Gamma} \in \operatorname{argmax}_{\Gamma' \in \mathbb{G}} U^P[\Gamma']$ be any policy that maximizes the policy maker's payoff over the set \mathbb{G} . Note that such a policy need not satisfy the perfect coordination property for, under $\tilde{\Gamma}$, there may exist x such that $U^{\tilde{\Gamma}}(x, 0|x) = 0$. It is also immediate to see that, because the original policy $\Gamma \in \mathbb{G}$, the policy maker's payoff under the original policy Γ cannot be greater than under $\tilde{\Gamma}$.

Below, we first show that, necessarily, under $\tilde{\Gamma} = (\{0, 1\}, \tilde{\pi})$, there exists θ^* such that $\tilde{\pi}(0|\theta) = 0$ for F -almost all $\theta \leq \theta^*$ and $\tilde{\pi}(0|\theta) = 1$ for F -almost all $\theta \geq \theta^*$. We then show that the policy maker's payoff under $\tilde{\Gamma}$ can be approximated arbitrarily well by a threshold policy $\Gamma^* \in \mathbb{G}$ that satisfies the perfect coordination property.

First observe that, under any policy $\Gamma' \in \mathbb{G}$, the function $U^{\Gamma'}(x, 0|x)$ is continuous in x . Now let

$$\mathcal{X} \equiv \left\{ x : U^{\tilde{\Gamma}}(x, 0|x) = 0 \right\}$$

and then denote by

$$\bar{x} \equiv \sup \mathcal{X} \quad \text{and} \quad \underline{x} \equiv \inf \mathcal{X}.$$

That $U^{\tilde{\Gamma}}(x, 0|x)$ is continuous in x implies that $\mathcal{X} \neq \emptyset$ for, otherwise, the policy maker could increase her payoff by increasing $\pi(0|\theta)$ over a set of F -positive measure, thus contradicting the optimality of $\tilde{\Gamma}$.

Now suppose that, under $\tilde{\Gamma}$, there exists no θ^* such that $\tilde{\pi}(0|\theta) = 0$ for F -almost all $\theta \leq \theta^*$ and $\tilde{\pi}(0|\theta) = 1$ for F -almost all $\theta \geq \theta^*$. We then show that there exists another policy in \mathbb{G} that strictly improves upon $\tilde{\Gamma}$, thus contradicting the assumption that $\tilde{\Gamma}$ maximizes the policy maker's payoff over \mathbb{G} .

Let

$$\theta_0 \equiv \inf \{ \theta : \exists \delta > 0 \text{ s.t. } \tilde{\pi}(0|\theta') > 0 \text{ for } F\text{-almost all } \theta' \in [\theta, \theta + \delta) \}$$

and, similarly,

$$\theta_1 = \sup \{ \theta : \exists \delta > 0 \text{ s.t. } \tilde{\pi}(0|\theta') < 1 \text{ for } F\text{-almost all } \theta' \in [\theta, \theta + \delta) \}.$$

That, under $\tilde{\Gamma}$, there exists no θ^* with the aforementioned properties implies that $\theta_0 < \theta_1$. Furthermore, $[\theta_0, \theta_1] \subset [\underline{\theta}, \bar{\theta}]$.

First, suppose that $\hat{\theta}(\bar{x}) < \theta_1$, where, recall that, for any x , $\hat{\theta}(x)$ is the regime threshold when all agents follow a cut-off strategy with cut-off x (i.e., $\hat{\theta}(x)$ is the regime threshold such that, when

agents attack when their signal falls below x and refrain from attacking otherwise, regime change occurs if, and only if, $\theta \leq \hat{\theta}(x)$.

Consider the policy $\Gamma^\epsilon = (\{0, 1\}, \pi^\epsilon)$ defined by: (a) $\pi^\epsilon(\theta) = \tilde{\pi}(\theta)$ for all $\theta \notin [\theta_0, \theta_0 + \epsilon] \cup [\theta_1 - \delta(\epsilon), \theta_1]$; (b) $\pi^\epsilon(0|\theta) = 0$ for all $\theta \in [\theta_0, \theta_0 + \epsilon]$, and $\pi^\epsilon(0|\theta) = 1$ for all $\theta \in [\theta_1 - \delta(\epsilon), \theta_1]$; (c) $\delta(\epsilon)$ solves

$$\int_{\theta_0}^{\theta_0 + \epsilon} \tilde{\pi}(0|\theta) dF(\theta) = \int_{\theta_1 - \delta(\epsilon)}^{\theta_1} (1 - \tilde{\pi}(0|\theta)) dF(\theta). \quad (9)$$

Provided ϵ is small, such a policy exists and $\theta_1 - \delta(\epsilon) > \hat{\theta}(\bar{x})$. Also note that the policy Γ^ϵ improves the policy maker's payoff upon $\tilde{\Gamma}$. This follows from the fact that the policy Γ^ϵ preserves the ex-ante probability the status quo survives, along with the fact that the policy maker's payoff differential $\Delta^P(\theta)$ is nondecreasing. To establish the result in the theorem it then suffices to show that the new policy Γ^ϵ satisfies the perfect coordination property. To see this, let $B(\theta, x) \equiv b(\theta, P(x|\theta))$ and $G(\theta, x) \equiv g(\theta, P(x|\theta))$, and then let

$$p^{\tilde{\Gamma}}(x, 0) \equiv \int \tilde{\pi}(0|\theta) p(x|\theta) dF(\theta) \text{ and } p^{\Gamma^\epsilon}(x, 0) \equiv \int \pi^\epsilon(0|\theta) p(x|\theta) dF(\theta)$$

denote the total probability of $(x, 0)$ under $\tilde{\Gamma}$ and Γ^ϵ , respectively. Next, observe that, for any $x \leq \hat{\theta}^{-1}(\theta_0 + \epsilon)$ (i.e., for any x such that $\hat{\theta}(x) < \theta_0 + \epsilon$), $U^{\Gamma^\epsilon}(x, 0|x) > 0$, whereas for any $x \in [\hat{\theta}^{-1}(\theta_0 + \epsilon), \bar{x}]$,

$$\begin{aligned} U^{\Gamma^\epsilon}(x, 0|x) p^{\Gamma^\epsilon}(x, 0) &= \int_{\theta_0 + \epsilon}^{\hat{\theta}(x)} B(\theta, x) \pi^\epsilon(0|\theta) p(x|\theta) dF(\theta) + \int_{\hat{\theta}(x)}^{+\infty} G(\theta, x) \pi^\epsilon(0|\theta) p(x|\theta) dF(\theta) \\ &> \int_{\theta_0}^{\hat{\theta}(x)} B(\theta, x) \tilde{\pi}(0|\theta) p(x|\theta) dF(\theta) + \int_{\hat{\theta}(x)}^{+\infty} G(\theta, x) \pi(0|\theta) p(x|\theta) dF(\theta) \\ &= U^{\tilde{\Gamma}}(x, 0|x) p^{\tilde{\Gamma}}(x, 0) \\ &\geq 0. \end{aligned}$$

The first inequality follows from the fact that, for any $\theta \in [\theta_0 + \epsilon, \hat{\theta}(x)]$, $\pi^\epsilon(\theta) = \tilde{\pi}(\theta)$, along with the properties that, for any (θ, x) , $B(\theta, x) < 0 < G(\theta, x)$, and the fact that, for all $\theta \geq \hat{\theta}(x)$, $\pi^\epsilon(\theta) \geq \tilde{\pi}(\theta)$. Hence, for all $x \leq \bar{x}$, $U^{\Gamma^\epsilon}(x, 0|x) > 0$. Now recall that, by definition of \bar{x} , $U^{\tilde{\Gamma}}(x, 0|x) > 0$ for all $x > \bar{x}$. That $U^{\Gamma^\epsilon}(x, 0|x)$ is continuous in x , along with the fact that $U^{\Gamma^\epsilon}(\bar{x}, 0|\bar{x}) > 0$, then also implies that, for ϵ strictly positive, but small enough, $U^{\Gamma^\epsilon}(x, 0|x) > 0$ for all $x > \bar{x}$. The new policy Γ^ϵ thus satisfies the perfect coordination property, as claimed above.

Next, suppose that $\hat{\theta}(\bar{x}) \geq \theta_1$. Consider the policy $\Gamma^\epsilon = (\{0, 1\}, \pi^\epsilon)$ defined by the following properties: (a) $\pi^\epsilon(\theta) = \tilde{\pi}(\theta)$ for all $\theta \notin [\theta_0, \theta_0 + \epsilon] \cup [\theta_1 - \delta(\epsilon), \theta_1]$ with $\theta_0 + \epsilon < \theta_1 - \delta(\epsilon)$; (b) $\pi^\epsilon(0|\theta) = 0$ for all $\theta \in [\theta_0, \theta_0 + \epsilon]$, and $\pi^\epsilon(0|\theta) = 1$ for all $\theta \in [\theta_1 - \delta(\epsilon), \theta_1]$; (c) with $\delta(\epsilon)$ implicitly defined by

$$\int_{\theta_0}^{\theta_0 + \epsilon} B(\theta, \bar{x}) \tilde{\pi}(0|\theta) p(\bar{x}|\theta) dF(\theta) = \int_{\theta_1 - \delta(\epsilon)}^{\theta_1} B(\theta, \bar{x}) (1 - \tilde{\pi}(0|\theta)) p(\bar{x}|\theta) dF(\theta). \quad (10)$$

Note that, for ϵ strictly positive but small, such policy exists. Also note that property (c) implies that, under the new policy Γ^ϵ , the payoff from not attacking of an agent with signal \bar{x} who expects all other agents to follow a cut-off strategy with cutoff \bar{x} is the same as under the original policy Γ . Formally, $U^{\Gamma^\epsilon}(\bar{x}, 0|\bar{x}) = U^{\bar{\Gamma}}(\bar{x}, 0|\bar{x})$. To see this, recall that, by definition of \bar{x} , $U^{\bar{\Gamma}}(\bar{x}, 0|\bar{x}) = 0$, or equivalently,

$$\int_{\theta_0}^{\hat{\theta}(\bar{x})} B(\theta, \bar{x}) \tilde{\pi}(0|\theta) p(\bar{x}|\theta) dF(\theta) + \int_{\hat{\theta}(\bar{x})}^{+\infty} G(\theta, \bar{x}) \tilde{\pi}(0|\theta) p(\bar{x}|\theta) dF(\theta) = 0.$$

Condition (10), along with properties (a) and (b) above, imply that

$$\int_{\theta_0+\epsilon}^{\hat{\theta}(\bar{x})} B(\theta, X) \pi^\epsilon(0|\theta) p(\bar{x}|\theta) dF(\theta) + \int_{\hat{\theta}(\bar{x})}^{+\infty} G(\theta, X) \pi^\epsilon(0|\theta) p(\bar{x}|\theta) dF(\theta) = 0.$$

Because

$$U^{\Gamma^\epsilon}(\bar{x}, 0|\bar{x}) = \frac{\int_{\theta_0+\epsilon}^{\hat{\theta}(\bar{x})} B(\theta, \bar{x}) \pi^\epsilon(0|\theta) p(\bar{x}|\theta) dF(\theta) + \int_{\hat{\theta}(\bar{x})}^{+\infty} G(\theta, \bar{x}) \pi^\epsilon(0|\theta) p(\bar{x}|\theta) dF(\theta)}{p^{\Gamma^\epsilon}(0, \bar{x})}.$$

we conclude that $U^{\Gamma^\epsilon}(\bar{x}, 0|\bar{x}) > 0$, as claimed.

We now establish that, for ϵ sufficiently small, under the new policy Γ^ϵ , $U^{\Gamma^\epsilon}(x, 0|x) > 0$ for any $x \in \mathcal{X}$, with $x \neq \bar{x}$. A necessary and sufficient condition for this to be the case is that, for any such x ,

$$\lim_{\epsilon \rightarrow 0^+} \frac{\partial}{\partial \epsilon} U^{\Gamma^\epsilon}(x, 0|x) > 0. \quad (11)$$

Condition (11) holds if, and only if, for any $x \in \mathcal{X}$, $x \neq \bar{x}$,

$$\lim_{\epsilon \rightarrow 0^+} \frac{\partial}{\partial \epsilon} \{U^{\Gamma^\epsilon}(x, 0|x) p^{\Gamma^\epsilon}(0, x)\} > 0. \quad (12)$$

To see that (12) holds, implicitly differentiate the equation in Condition (10) with respect to ϵ and then take the limit as $\epsilon \rightarrow 0$ to observe that

$$\lim_{\epsilon \rightarrow 0^+} \delta'(\epsilon) = \frac{\tilde{\pi}(0|\theta_0) f(\theta_0) p(\bar{x}|\theta_0) |B(\theta_0, \bar{x})|}{(1 - \tilde{\pi}(0|\theta_1)) f(\theta_1) p(\bar{x}|\theta_1) |B(\theta_1, \bar{x})|}.$$

Now take first any $x \in [\hat{\theta}^{-1}(\theta_1), \bar{x}) \cap \mathcal{X}$ and observe that, for any such x ,

$$U^{\Gamma^\epsilon}(x, 0|x) = \frac{\int_{\theta_0+\epsilon}^{\theta_1-\delta} B(\theta, x) \tilde{\pi}(0|\theta) p(x|\theta) dF(\theta) + \int_{\theta_1-\delta}^{\theta_1} B(\theta, x) p(x|\theta) dF(\theta)}{p^{\Gamma^\epsilon}(0, x)} + \frac{\int_{\theta_1}^{\hat{\theta}(x)} B(\theta, x) \tilde{\pi}(0|\theta) p(x|\theta) dF(\theta) + \int_{\hat{\theta}(x)}^{+\infty} G(\theta, x) \tilde{\pi}(0|\theta) p(x|\theta) dF(\theta)}{p^{\Gamma^\epsilon}(0, x)}. \quad (13)$$

Now use (13) to observe that, for any $x \in [\hat{\theta}^{-1}(\theta_1), \bar{x}) \cap \mathcal{X}$,

$$\begin{aligned} & \lim_{\epsilon \rightarrow 0^+} \frac{\partial}{\partial \epsilon} \{U^{\Gamma^\epsilon}(x, 0|x) p^{\Gamma^\epsilon}(0, x)\} = -B(\theta_0, x) \tilde{\pi}(0|\theta_0) p(x|\theta_0) f(\theta_0) \\ & \quad + B(\theta_1, x) (1 - \tilde{\pi}(0|\theta_1)) p(x|\theta_1) f(\theta_1) \lim_{\epsilon \rightarrow 0} \delta'(\epsilon) \\ & = \tilde{\pi}(0|\theta_0) f(\theta_0) p(\bar{x}|\theta_0) |B(\theta_0, \bar{x})| \left(\frac{p(x|\theta_0) |B(\theta_0, x)|}{p(\bar{x}|\theta_0) |B(\theta_0, \bar{x})|} - \frac{p(x|\theta_1) |B(\theta_1, x)|}{p(\bar{x}|\theta_1) |B(\theta_1, \bar{x})|} \right) > 0, \end{aligned}$$

where the inequality follows from the log-supermodularity of $p(x|\theta)$ and $B(\theta, x)$.

Finally, consider any $x \in [\underline{x}, \hat{\theta}^{-1}(\theta_1)] \cap \mathcal{X}$. That $U^{\tilde{\Gamma}}(\underline{x}, 0|\underline{x}) = 0$ implies that $\hat{\theta}(\underline{x}) > \theta_0$. Hence, for any $x \in [\underline{x}, \hat{\theta}^{-1}(\theta_1)] \cap \mathcal{X}$, $\hat{\theta}(x) > \theta_0$, which means that

$$\begin{aligned}
U^{\Gamma^\epsilon}(x, 0|x)p^{\Gamma^\epsilon}(0, x) &= I(\hat{\theta}(x) > \theta_0 + \epsilon) \int_{[\theta_0 + \epsilon, \hat{\theta}(x)]} B(\theta, x)\tilde{\pi}(0|\theta)p(x|\theta)dF(\theta) \\
&+ \int_{[\max\{\hat{\theta}(x), \theta_0 + \epsilon\}, \theta_1 - \delta(\epsilon)]} G(\theta, x)\tilde{\pi}(0|\theta)p(x|\theta)dF(\theta) \\
&+ \int_{\theta_1 - \delta(\epsilon)}^{\theta_1} G(\theta, x)p(x|\theta)dF(\theta) + \int_{\theta_1}^{+\infty} G(\theta, x)\tilde{\pi}(0|\theta)p(x|\theta)dF(\theta) \\
&> U^{\tilde{\Gamma}}(x, 0|x)p^{\tilde{\Gamma}}(0, x) \\
&\geq 0,
\end{aligned} \tag{14}$$

$$\tag{15}$$

where $I(\hat{\theta}(x) > \theta_0 + \epsilon)$ is the indicator function, taking value 1 when $\hat{\theta}(x) > \theta_0 + \epsilon$, and 0 otherwise.

We conclude that, when ϵ is small, under the new policy Γ^ϵ , for all $x \in \mathcal{X}$, $x \neq \bar{x}$, $U^{\Gamma^\epsilon}(x, 0|x) > 0$. That, under the same policy, $U^{\Gamma^\epsilon, \theta^s}(x, 0|x) > 0$ also for $x \notin \mathcal{X}$ follows from the fact that $U^{\Gamma^\epsilon}(x, 0|x)$ is continuous in (x, ϵ) .

From the arguments above, we have that the new policy $\Gamma^\epsilon \in \mathbb{G}$. We now show that, when, in addition, Condition (1) holds, the new policy strictly improves the policy maker's expected payoff relative to $\tilde{\Gamma}$. To see this, observe that, for any $\epsilon \geq 0$, the policy maker's payoff under the policy Γ^ϵ is equal to

$$\begin{aligned}
U^P[\Gamma^\epsilon] &= \int_{-\infty}^{\theta_0 + \epsilon} L(\theta)dF(\theta) + \int_{\theta_1 - \delta(\epsilon)}^{\theta_1} W(\theta, 0)dF(\theta) \\
&+ \int_{(\theta_0 + \epsilon, \theta_1 - \delta(\epsilon)) \cup (\theta_1, +\infty)} (\tilde{\pi}(0|\theta)W(\theta, 0) + (1 - \tilde{\pi}(0|\theta))L(\theta))dF(\theta)
\end{aligned}$$

Differentiating $U^P[\Gamma^\epsilon]$ with respect to ϵ and taking the limit as $\epsilon \rightarrow 0$, we have that:

$$\begin{aligned}
\lim_{\epsilon \rightarrow 0^+} \frac{dU^P[\Gamma^\epsilon]}{d\epsilon} &= f(\theta_1)(1 - \tilde{\pi}(0|\theta_1))\Delta^P(\theta_1) \left(\lim_{\epsilon \rightarrow 0} \delta'(\epsilon) \right) - f(\theta_0)\tilde{\pi}(0|\theta_0)\Delta^P(\theta_0) \\
&= f(\theta_0)\tilde{\pi}(0|\theta_0) \left(\Delta^P(\theta_1) \frac{p(\bar{x}|\theta_0)|B(\theta_0, \bar{x})}{p(\bar{x}|\theta_1)|B(\theta_1, \bar{x})} - \Delta^P(\theta_0) \right).
\end{aligned}$$

Therefore, a sufficient condition for $\lim_{\epsilon \rightarrow 0^+} \frac{dU^P[\Gamma^\epsilon]}{d\epsilon} > 0$ is that

$$\frac{\Delta^P(\theta)}{p(X|\theta)|B(\theta, X)|}$$

is nondecreasing over $[\theta_0, \theta_1]$. Condition (M) guarantees this is the case.

We conclude that, no matter whether, under the original policy $\tilde{\Gamma}$, $\hat{\theta}(X) \geq \theta_1$, or $\hat{\theta}(X) < \theta_1$, starting from $\tilde{\Gamma}$, there exist policies in \mathbb{G} that strictly improve upon $\tilde{\Gamma}$, which contradicts the optimality of $\tilde{\Gamma}$. This means that any policy that maximizes the policy maker's payoff over \mathbb{G} is such

that there exists θ^* such that $\tilde{\pi}(0|\theta) = 0$ for F -almost all $\theta \leq \theta^*$ and $\tilde{\pi}(0|\theta) = 1$ for F -almost all $\theta \geq \theta^*$.

Now recall that the original policy $\Gamma = (\{0, 1\}, \pi)$ is such that there exists no θ^* such that $\pi(0|\theta) = 0$ for F -almost all $\theta \leq \theta^*$ and $\pi(0|\theta) = 1$ for F -almost all $\theta \geq \theta^*$. This means that any threshold policy $\tilde{\Gamma}$ that maximizes the policy maker's payoff over \mathbb{G} yields a payoff strictly higher than Γ . The result in the theorem then follows from observing that, given $\tilde{\Gamma}$, there exists a nearby threshold policy $\Gamma^* \in \mathbb{G}$ with cut-off θ^{**} arbitrarily close θ^* that satisfies the perfect coordination property (i.e., such that $U^{\Gamma^*}(x, 0|x) > 0$ all x) and that yields the policy maker a payoff arbitrarily close to that of $\tilde{\Gamma}$. To see this, it suffices to note that the threshold θ^* that defines $\tilde{\pi}$ in $\tilde{\Gamma}$ is given by

$$\theta^* \equiv \inf\{\theta' : \int_{\theta'}^{\infty} u(\theta, P(x|\theta))p(x|\theta)f(\theta)d\theta > 0 \text{ for all } x \in \mathbb{R}\}.$$

The policy maker's payoff under $\tilde{\Gamma}$ can then be approximated arbitrarily well by any threshold policy with cut-off equal to $\theta^* + \varepsilon$. Because any such policy satisfies the perfect coordination policy, we then have that the result in the theorem holds. Q.E.D.

Proof of Theorem 3. The proof parallels the arguments in Step 1 in the proof of Theorem 1. Take an arbitrary policy $\Gamma = (\mathcal{S}, \pi)$, and for any $\theta \in \mathbb{R}$, any $\phi \in \mathbb{R}^{[0,1]}$, any message function $m \in \text{supp} [\pi(\theta, \phi)]$, let $r(\theta, \phi, m; \Gamma) \in \{0, 1\}$ denote the regime outcome that prevails at θ when the distribution of exogenous signals is ϕ , the distribution of endogenous signals is m , and agents play according to the *most aggressive rationalizable strategy profile* (MARP) consistent with Γ . Note that, under arbitrary policies, such a profile need not be monotone in the agents' exogenous, or endogenous, signals.

For any $n \geq 1$, let T^n be the set of strategies surviving n rounds of IDISDS under the original policy Γ , and $\bar{a}^n \equiv (\bar{a}_i^n(\cdot))_{i \in [0,1]} \in T^n$ the profile in T^n that minimizes the policy maker's ex-ante payoff (as in the case of non-discriminatory policies, such a profile also minimizes the policy maker's interim payoff, as it will become clear from the arguments below). Hereafter, we refer to this profile as the most aggressive profile surviving n rounds of IDISDS. The profiles $(\bar{a}^n)_{n \in \mathbb{N}}$ can be constructed inductively as follows. The profile $\bar{a}^0 \equiv (\bar{a}_i^0(\cdot))_{i \in [0,1]}$ prescribes that agents attack irrespective of their exogenous and endogenous signals; that is, each $\bar{a}_i(\cdot)$ is such that $\bar{a}_i(x, s) = 1$, all $(x, s) \in \mathbb{R} \times \mathcal{S}$.²⁶ With an abuse of notation, for any $n \geq 1$, then let $U_i^\Gamma((x, s); a^{n-1})$ denote the payoff that each agent i obtains from not attacking, after observing the exogenous signal x and the endogenous signal s , when all other agents follow the strategy in the profile a^{n-1} . The most aggressive strategy profile surviving n rounds of IDISDS is the one specifying, for each agent $i \in [0, 1]$, $\bar{a}_i^n(x, s) = 0$ if $U_i^\Gamma((x, s); \bar{a}^{n-1}) > 0$ and $\bar{a}_i^n(x, s) = 1$ if $U_i^\Gamma((x, s); \bar{a}^{n-1}) \leq 0$. The most aggressive rationalizable strategy profile (MARP) consistent with the policy Γ is then the profile $a^\Gamma = (a_i^\Gamma(\cdot))_{i \in [0,1]} \in T^\infty$ given by

$$a_i^\Gamma(\cdot) = \bar{a}_i^\infty(\cdot) \equiv \lim_{n \rightarrow \infty} \bar{a}_i^n(\cdot), \text{ all } i \in [0, 1].$$

²⁶Note that, to ease the notation, we let each individual strategy prescribe an action for all $(x, s) \in \mathbb{R} \times \mathcal{S}$, including those that are possibly inconsistent with the policy Γ .

Next, consider the policy Γ^+ obtained from the original policy Γ by replacing each message function $m : [0, 1] \rightarrow \mathcal{S}$ in the support of each $\pi(\theta, \phi)$ with the message function $m^+ : [0, 1] \rightarrow \mathcal{S} \times \{0, 1\}$ that discloses to each agent $i \in [0, 1]$ the same message m_i disclosed by the original policy Γ , along with the regime outcome $r(\theta, \phi, m; \Gamma)$ that would have prevailed at θ when the distribution of exogenous signals is ϕ , the distribution of endogenous signals is m , and all agents play according to the most aggressive rationalizable strategy profile (MARP) \bar{a}^Γ consistent with the original policy Γ (that is, $m_i^+ = (m_i, r(\theta, \phi, m; \Gamma))$). Note that the new policy $\Gamma^+ = (S^+, \pi^+)$ selects, for each (θ, ϕ) , the message function m^+ with the same probability as the original policy Γ selects the function m .

Now let T_+^n denote the set of strategies surviving n rounds of IDISDS under the new policy Γ^+ , and $\bar{a}_+^n = (\bar{a}_{i+}^n(\cdot))_{i \in [0, 1]} \in T_+^n$ the profile in T_+^n that minimizes the policy maker's ex-ante payoff, where $\bar{a}_+^0 = (\bar{a}_{i+}^0(\cdot))_{i \in [0, 1]}$ prescribe that all agents attack irrespective of their exogenous and endogenous signals.

Step 1. First, we prove that

$$\{(x_i, m_i) : U_i^\Gamma(x_i, m_i; a) > 0 \forall a : \mathbb{R} \times \mathcal{S} \rightarrow \{0, 1\}\} \subseteq \{(x_i, m_i) : U_i^{\Gamma^+}(x_i, (m_i, 0); a) > 0 \forall a : \mathbb{R} \times \mathcal{S} \rightarrow \{0, 1\}\}.$$

To see this, first use the fact that the game is supermodular to observe that, given any policy Γ ,

$$\{(x_i, m_i) : U_i^\Gamma(x_i, m_i; a) > 0 \forall a : \mathbb{R} \times \mathcal{S} \rightarrow \{0, 1\}\} = \{(x_i, m_i) : U_i^\Gamma(x_i, m_i; \bar{a}^0) > 0\}.$$

Likewise,

$$\{(x_i, m_i) : U_i^{\Gamma^+}(x_i, (m_i, 0); a) > 0 \forall a : \mathbb{R} \times \mathcal{S} \rightarrow \{0, 1\}\} = \{(x_i, m_i) : U_i^{\Gamma^+}(x_i, (m_i, 0); \bar{a}_+^0) > 0\}.$$

Next, recall that, under $\bar{a}^0 = \bar{a}_+^0$, all agents attack regardless of their exogenous and endogenous information and therefore, under $\bar{a}^0 = \bar{a}_+^0$, regime change occurs if, and only if, $\theta \leq \bar{\theta}$.

Take any $(x_i, m_i) \in \mathbb{R} \times \mathcal{S}$ such that

$$U_i^\Gamma((x_i, m_i); \bar{a}^0) = \int_{(\theta, \phi, m)} \left(b(\theta, 1) 1_{\{\theta \leq \bar{\theta}\}} + g(\theta, 1) 1_{\{\theta > \bar{\theta}\}} \right) d\Lambda^\Gamma((\theta, \phi, m) | x_i, m_i) > 0.$$

Then, note that, under Γ^+ ,

$$\partial \Lambda^{\Gamma^+}((\theta, \phi, m) | x_i, (m_i, 0)) = \frac{1_{\{r(\theta, \phi, m; \Gamma) = 0\}}}{\pi^\Gamma(0 | x_i, m_i)} \partial \Lambda^\Gamma((\theta, \phi, m) | x_i, m_i) \quad (16)$$

where

$$\pi^\Gamma(0 | x_i, m_i) \equiv \int_{\{(\theta, \phi, m) : r(\theta, \phi, m; \Gamma) = 0\}} d\Lambda^\Gamma((\theta, \phi, m) | x_i, m_i)$$

is the total probability assigned by agent i to the event $\{(\theta, \phi, m) : r(\theta, \phi, m; \Gamma) = 0\}$ when observing the signals (x_i, m_i) under Γ and when expecting all agents to play according to the most aggressive rationalizable profile consistent with Γ . At an intuitive level, the agents' beliefs under the new policy

policy Γ^+ corresponds to “truncations” of the posteriors beliefs under the original policy Γ . In turn, this property implies that, given any $(x_i, m_i) \in \mathbb{R} \times \mathcal{S}$ such that $U_i^\Gamma((x_i, m_i); \bar{a}^0) > 0$,

$$\begin{aligned}
U_i^{\Gamma^+}(x_i, (m_i, 0); \bar{a}_+^0) &= \frac{1}{\pi^\Gamma(0|x_i, m_i)} \int_{(\theta, \phi, m)} \left(b(\theta, 1)1_{\{\theta \leq \bar{\theta}\}} + g(\theta, 1)1_{\{\theta > \bar{\theta}\}} \right) \times \\
&\quad \times 1_{\{r(\theta, \phi, m; \Gamma) = 0\}} d\Lambda^\Gamma((\theta, \phi, m)|x_i, m_i) \\
&> \frac{1}{\pi^\Gamma(0|x_i, m_i)} \int_{(\theta, \phi, m)} \left(b(\theta, 1)1_{\{\theta \leq \bar{\theta}\}} + g(\theta, 1)1_{\{\theta > \bar{\theta}\}} \right) d\Lambda^\Gamma((\theta, \phi, m)|x_i, m_i) \\
&= \frac{1}{\pi^\Gamma(0|x_i, m_i)} U_i^\Gamma((x_i, m_i); \bar{a}^0) \\
&> 0
\end{aligned}$$

where the first equality follows from the truncation property introduced above, the first inequality from the fact that, for all (θ, ϕ, m) such that $r(\theta, \phi, m; \Gamma) = 1$,

$$b(\theta, 1)1_{\{\theta \leq \bar{\theta}\}} + g(\theta, 1)1_{\{\theta > \bar{\theta}\}} = b(\theta, 1) < 0,$$

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

This means that, any an agent for whom not attacking is dominant under Γ , continue to find it dominant not to attack after observing $(x_i, (m_i, 0))$ under Γ^+ .

Step 2. Next, take any $n > 0$. Assume that, for any $0 \leq k \leq n - 1$, any $i \in [0, 1]$,

$$\{(x_i, m_i) : U_i^\Gamma(x_i, m_i; a) > 0 \quad \forall a \in T^k\} \subseteq \{(x_i, m_i) : U_i^{\Gamma^+}(x_i, (m_i, 0); a) > 0, \quad \forall a \in T_+^k\}. \quad (17)$$

Recall that this means that any agent who 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,

$$\{(x_i, m_i) : U_i^\Gamma(x_i, m_i; a) > 0 \quad \forall a \in T^n\} \subseteq \{(x_i, m_i) : U_i^{\Gamma^+}(x_i, (m_i, 0); a) > 0, \quad \forall a \in T_+^n\}. \quad (18)$$

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

$$\{(x_i, m_i) : U_i^\Gamma(x_i, m_i; a) > 0 \quad \forall a \in T^n\} = \{(x_i, m_i) : U_i^\Gamma(x_i, m_i; \bar{a}_\Gamma^n) > 0\}$$

and, likewise,

$$\{(x_i, m_i) : U_i^{\Gamma^+}(x_i, (m_i, 0); a) > 0, \quad \forall a \in T_+^n\} = \{(x_i, m_i) : U_i^{\Gamma^+}(x_i, m_i; \bar{a}_{\Gamma^+}^n) > 0\},$$

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

Now let $A(\theta, \phi, m; a)$ and $r(\theta, \phi, m; a)$ denote, respectively, the aggregate size of attack and the regime outcome that prevail at (θ, ϕ, m) when agents play according to a . Then take any $i \in [0, 1]$ and any $(x_i, m_i) \in \mathbb{R} \times \mathcal{S}$ such that

$$U_i^\Gamma((x_i, m_i); \bar{a}^n) = \int_{(\theta, \phi, m)} u(\theta, A(\theta, \phi, m; \bar{a}^n)) d\Lambda^\Gamma((\theta, \phi, m)|x_i, m_i) > 0.$$

Because \bar{a}^n is more aggressive than a^Γ , for all (θ, ϕ, m) ,

$$r(\theta, \phi, m; \Gamma) = 1 \Rightarrow r(\theta, \phi, m; \bar{a}^n) = 1.$$

This implies that

$$\begin{aligned} \int_{(\theta, \phi, m)} u(\theta, A(\theta, \phi, m; \bar{a}^n)) 1_{\{r(\theta, \phi, m; \Gamma)=1\}} d\Lambda^\Gamma((\theta, \phi, m)|x_i, m_i) &= \\ \int_{(\theta, \phi, m)} b(\theta, A(\theta, \phi, m; \bar{a}^n)) 1_{\{r(\theta, \phi, m; \Gamma)=1\}} d\Lambda^\Gamma((\theta, \phi, m)|x_i, m_i) &< 0. \end{aligned} \quad (19)$$

This observation, together with the truncation property in (16), imply that

$$\begin{aligned} U_i^{\Gamma^+}(x_i, (m_i, 0); \bar{a}^n) &= \int_{(\theta, \phi, m)} u(\theta, A(\theta, \phi, m; \bar{a}^n)) d\Lambda^{\Gamma^+}((\theta, \phi, m)|x_i, m_i) \\ &= \frac{1}{\pi^\Gamma(0|x_i, m_i)} \int_{(\theta, \phi, m)} u(\theta, A(\theta, \phi, m; \bar{a}^n)) 1_{\{r(\theta, \phi, m; \Gamma)=0\}} d\Lambda^\Gamma((\theta, \phi, m)|x_i, m_i) \\ &> \frac{1}{\pi^\Gamma(0|x_i, m_i)} \int_{(\theta, \phi, m)} u(\theta, A(\theta, \phi, m; \bar{a}^n)) d\Lambda^\Gamma((\theta, \phi, m)|x_i, m_i) \\ &= \frac{1}{\pi^\Gamma(0|x_i, m_i)} U_i^\Gamma((x_i, m_i); \bar{a}^n) \\ &> 0, \end{aligned} \quad (20)$$

where the first and third equalities are by definition, the second equality follows from (16), the first inequality follows from (19), and the last inequality is by assumption.

Next, note that \bar{a}^n and \bar{a}_+^n are such that, for all $i \in [0, 1]$, all $(x_i, m_i) \in \mathbb{R} \times \mathcal{S}$, $\bar{a}_i^n(x_i, m_i), \bar{a}_{i+}^n(x_i, m_i) \in \{0, 1\}$ and

$$\{(x_i, m_i) : \bar{a}_i^n(x_i, m_i) = 0\} = \{(x_i, m_i) : U_i^\Gamma(x_i, m_i; \bar{a}_\Gamma^{n-1}) > 0\}$$

and

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

The result in Step 1, along with (17), thus imply that \bar{a}^{n-1} and \bar{a}_+^{n-1} are thus such that, for all $i \in [0, 1]$, all $(x_i, m_i) \in \mathbb{R} \times \mathcal{S}$,

$$\bar{a}_i^{n-1}(x_i, m_i) = 0 \Rightarrow \bar{a}_{i+}^{n-1}(x_i, m_i) = 0. \quad (21)$$

Condition (21), along with the fact that the game is supermodular, implies that

$$U_i^{\Gamma^+}(x_i, (m_i, 0); \bar{a}^n) > 0 \Rightarrow U_i^{\Gamma^+}(x_i, (m_i, 0); \bar{a}_+^n) > 0. \quad (22)$$

Together (20) and (22) imply the property in (18).

Step 3. Equipped with the results in steps 1 and 2 above, we now prove that, for all $i \in [0, 1]$, all $(x_i, m_i) \in \mathbb{R} \times \mathcal{S}$,

$$a_i^{\Gamma^+}(x_i, (m_i, 0)) \equiv \lim_{n \rightarrow \infty} \bar{a}_{i+}^n(x_i, (m_i, 0)) = 0.$$

This follows directly from the fact that, for all $i \in [0, 1]$, all $(x_i, m_i) \in \mathbb{R} \times \mathcal{S}$,

$$a_i^{\Gamma}(x_i, m_i) = 0 \Rightarrow a_i^{\Gamma^+}(x_i, (m_i, 0)) = 0. \quad (23)$$

which, in turn implies that, for any (θ, ϕ, m) ,

$$r(\theta, \phi, m; a^{\Gamma}) = 0 \Rightarrow r(\theta, \phi, m; a^{\Gamma^+}) = 0.$$

The announcement that $r = 0$ thus reveals to the agent that (θ, ϕ, m) is such that $r(\theta, \phi, m; a^{\Gamma^+}) = 0$. 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 the new signal structure Γ^+ , all agents thus refrain from attacking, regardless of their exogenous and endogenous private signals, when they learn that $r = 0$. That they all attack when they learn that $r = 1$ follows from the fact that $r = 1$ makes it common certainty that $\theta \leq \bar{\theta}$.

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 (θ, ϕ) , 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.

Proof of Proposition 3. We establish the result by showing that, when Condition (6) holds, for any fixed $\hat{\theta}$, the function $\Psi(\hat{\theta}, \sigma_z) \equiv \min_{\theta_0} \psi(\theta_0, \hat{\theta}, \sigma_z)$ is increasing in σ_z . Moreover, in this case, the regime threshold in the absence of any public disclosure, $\theta_{\sigma_z}^*$, is decreasing in σ_z , and $\lim_{\sigma_z \rightarrow 0^+} \theta_{\sigma_z}^* = \theta^{MS}$.

To ease the notation, let $\sigma = \sigma_z$. By the envelope theorem, we have that

$$\Psi_{\sigma}(\hat{\theta}, \sigma) = \psi_{\sigma}(\bar{\theta}_{\sigma}, \hat{\theta}, \sigma),$$

where recall that $\bar{\theta}_{\sigma} = \arg \min_{\theta_0} \psi(\theta_0, \hat{\theta}, \sigma)$.

Note that, for any $\theta_0 > \hat{\theta}$, any σ ,

$$\begin{aligned} \psi_{\sigma}(\theta_0, \hat{\theta}, \sigma) &= \frac{\partial}{\partial \sigma} \int_{\hat{\theta}}^{\infty} (\bar{b}(\theta) 1_{\theta < \theta_0} + \bar{g}(\theta) 1_{\theta \geq \theta_0}) \frac{\phi\left(\frac{z_{\sigma}^*(\theta_0) - \theta}{\sigma}\right)}{\sigma \Phi\left(\frac{z_{\sigma}^*(\theta_0) - \hat{\theta}}{\sigma}\right)} d\theta \\ &= \frac{\partial}{\partial \sigma} \frac{\int_0^{\Phi\left(\frac{z_{\sigma}^*(\theta_0) - \hat{\theta}}{\sigma}\right)} (\bar{b}(z_{\sigma}^*(\theta_0) - \sigma \Phi^{-1}(A)) 1_{A > \theta_0} + \bar{g}(z_{\sigma}^*(\theta_0) - \sigma \Phi^{-1}(A)) 1_{A \leq \theta_0}) dA}{\Phi\left(\frac{z_{\sigma}^*(\theta_0) - \hat{\theta}}{\sigma}\right)} \\ &= \frac{\int_0^{\Phi\left(\frac{z_{\sigma}^*(\theta_0) - \hat{\theta}}{\sigma}\right)} (\bar{b}'(z_{\sigma}^*(\theta_0) - \sigma \Phi^{-1}(A)) 1_{A > \theta_0} + \bar{g}'(z_{\sigma}^*(\theta_0) - \sigma \Phi^{-1}(A)) 1_{A \leq \theta_0}) (\Phi^{-1}(\theta_0) - \Phi^{-1}(A)) dA}{\Phi\left(\frac{z_{\sigma}^*(\theta_0) - \hat{\theta}}{\sigma}\right)} \\ &+ \frac{(\psi(\theta_0, \hat{\theta}, \sigma) - b(\hat{\theta})) \phi\left(\frac{z_{\sigma}^*(\theta_0) - \hat{\theta}}{\sigma}\right) (\theta_0 - \hat{\theta})}{\sigma^2 \Phi\left(\frac{z_{\sigma}^*(\theta_0) - \hat{\theta}}{\sigma}\right)} \end{aligned}$$

where the second equality follows from the change of variables $A = \Phi\left(\frac{z_\sigma^*(\theta_0) - \theta}{\sigma}\right)$ and the third equality from recalling that $z_\sigma^*(\theta) = \theta + \sigma\Phi^{-1}(\theta)$. Lastly, by reverting the change of variables we have:

$$\begin{aligned}\psi_\sigma(\theta_0, \hat{\theta}, \sigma) &= \frac{\int_{\hat{\theta}}^\infty D(\theta, \theta_0)(\theta - \theta_0)\phi\left(\frac{z_\sigma^*(\theta_0) - \theta}{\sigma}\right) d\theta + (\psi(\theta_0, \hat{\theta}, \sigma) - b(\hat{\theta}))\phi\left(\frac{z_\sigma^*(\theta_0) - \hat{\theta}}{\sigma}\right) (\theta_0 - \hat{\theta})}{\sigma^2\Phi\left(\frac{z_\sigma^*(\theta_0) - \hat{\theta}}{\sigma}\right)} \\ &= \sigma^{-1}\mathbb{E}[D(\theta, \theta_0)(\theta - \theta_0)|z_\sigma^*(\theta_0), \theta > \hat{\theta}] + \frac{(\psi(\theta_0, \hat{\theta}, \sigma) - b(\hat{\theta}))\phi\left(\frac{z_\sigma^*(\theta_0) - \hat{\theta}}{\sigma}\right) (\theta_0 - \hat{\theta})}{\sigma^2\Phi\left(\frac{z_\sigma^*(\theta_0) - \hat{\theta}}{\sigma}\right)},\end{aligned}$$

Thus, when evaluated at $\hat{\theta} = \theta_\sigma^{inf}$ and at $\theta_0 = \theta_\sigma^\#$, the above expression becomes

$$\psi_\sigma(\theta_\sigma^\#, \theta_\sigma^{inf}, \sigma) = \sigma^{-1}\mathbb{E}[D(\theta, \theta_\sigma^\#)(\theta - \theta_\sigma^\#)|z_\sigma^*(\theta_\sigma^\#), \theta > \theta_\sigma^{inf}] + \frac{|b(\theta_\sigma^{inf})|\phi\left(\frac{z_\sigma^*(\theta_\sigma^\#) - \theta_\sigma^{inf}}{\sigma}\right) (\theta_\sigma^\# - \theta_\sigma^{inf})}{\sigma^2\Phi\left(\frac{z_\sigma^*(\theta_\sigma^\#) - \theta_\sigma^{inf}}{\sigma}\right)}.\quad (24)$$

It is now easy to see that Condition (6) implies that $\psi_\sigma(\theta_\sigma^\#, \theta_\sigma^{inf}, \sigma) > 0$.

The above property implies that, for fixed $\theta_{\sigma_z}^{inf}$, an increase in σ_z increases $\Psi_{\sigma_z}(\theta_{\sigma_z}^{inf}, \sigma_z)$. Furthermore, because the threshold $\theta_{\sigma_z}^\#$ solves $\psi(\theta_{\sigma_z}^\#, \theta_{\sigma_z}^{inf}, \sigma_z) = 0$, we have that, by increasing σ_z while keeping $\theta_{\sigma_z}^{inf}$ fixed, the policy maker guarantees that, for any $\theta > \theta_{\sigma_z}^{inf}$, $\psi(\theta, \theta_{\sigma_z}^{inf}, \sigma_z) > 0$. That $\theta_{\sigma_z}^{inf}$ is decreasing in then follows from the fact that, for any σ_z , any $\theta > \hat{\theta}$, $\psi(\theta, \hat{\theta}, \sigma_z)$ is strictly increasing in $\hat{\theta}$ (this property follows from Lemma 2 in Angeletos et al. (2007)). Starting from any discriminatory policy, a reduction in the precision of the agents' private information (i.e., an increase in σ_z) lowers the fundamental threshold $\theta_{\sigma_z}^{inf}$ below which regime change occurs, thus improving the policy maker's payoff. Q.E.D.

References

- Alonso, R., Camara, O., 2015. Persuading voters. forthcoming. *American Economic Review*.
- Alonso, R., Camara, O., 2016. Bayesian persuasion with heterogeneous priors. *Journal of Economic Theory* 165, 672–706.
- Angeletos, G.-M., Hellwig, C., Pavan, A., 2006. Signaling in a global game: Coordination and policy traps. *Journal of Political Economy* 114 (3), 452–484.
- 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.
- Angeletos, G.-M., Pavan, A., 2013. Selection-free predictions in global games with endogenous information and multiple equilibria. *Theoretical Economics* 8 (3), 883–938.
- Aumann, R. J., Maschler, M., 1995. Repeated games with incomplete information. MIT press.
- Bardhi, A., Guo, Y., 2016. Modes of persuasion toward unanimous consent. Working paper, Northwestern University.
- Bergemann, D., Brooks, B., Morris, S., 2015. The limits of price discrimination. *The American Economic Review* 105 (3), 921–957.
- Bergemann, D., Morris, S., 2013. Robust predictions in games with incomplete information. *Econometrica* 81 (4), 1251–1308.
- Bergemann, D., Morris, S., 2016. Bayes correlated equilibrium and the comparison of information structures in games. *Theoretical Economics* 11 (2), 487–522.
- Bergemann, D., Morris, S., 2017. Information design: A unified perspective. DETC working paper 084-2016.
- Bouvard, M., Chaigneau, P., Motta, A. d., 2015. Transparency in the financial system: Rollover risk and crises. *The Journal of Finance* 70 (4), 1805–1837.
- Calzolari, G., Pavan, A., 2006a. Monopoly with resale. *The RAND Journal of Economics* 37 (2), 362–375.
- Calzolari, G., Pavan, A., 2006b. On the optimality of privacy in sequential contracting. *Journal of Economic theory* 130 (1), 168–204.
- Chan, J., Gupta, S., Li, F., Wang, Y., 2016. Pivotal persuasion. Mimeo. Available at SSRN.
- Denti, T., 2015. Unrestricted information acquisition. Mimeo.

- Dworczak, P., 2016. Mechanism design with aftermarkets: Cutoff mechanisms. Mimeo.
- Edmond, C., 2013. Information manipulation, coordination, and regime change. *The Review of Economic Studies*, rdt020.
- Ely, J. C., 2017. Beeps. *The American Economic Review* 107 (1), 31–53.
- Goldstein, I., Huang, C., 2016. Bayesian persuasion in coordination games. *The American Economic Review* 106 (5), 592–596.
- Goldstein, I., Leitner, Y., 2015. Stress tests and information disclosure. FRB of Philadelphia Working Paper.
- Goldstein, I., Sapra, H., 2014. Should banks' stress test results be disclosed? an analysis of the costs and benefits. *Foundations and Trends (R) in Finance* 8 (1), 1–54.
- Iachan, F. S., Nenov, P. T., 2015. Information quality and crises in regime-change games. *Journal of Economic Theory* 158, 739–768.
- Kamenica, E., Gentzkow, M., 2011. Bayesian persuasion. *American Economic Review* 101, 2590–2615.
- Kolotilin, A., Mylovanov, T., Zapechelnuk, A., Li, M., 2016. Persuasion of a privately informed receiver. Mimeo.
- Lerner, J., Tirole, J., 2006. A model of forum shopping. *The American economic review* 96 (4), 1091–1113.
- Mathevet, L., Perego, J., Taneva, I., 2016. Information design: The epistemic approach. Tech. rep., Working Paper.
- Mensch, J., 2015. Monotone persuasion. Tech. rep., Discussion Paper, Center for Mathematical Studies in Economics and Management Science.
- Moriya, F., Yamashita, T., et al., 2017. Asymmetric information allocation to avoid coordination failure. Tech. rep.
- Morris, S., Shin, H. S., 2006. Global games: Theory and applications. *Advances in Economics and Econometrics*, 56.
- Morris, S., Yang, M., 2016. Coordination under continuous choice. Tech. rep.
- Myerson, R. B., 1986. Multistage games with communication. *Econometrica: Journal of the Econometric Society*, 323–358.

- Segal, I., 2003. Coordination and discrimination in contracting with externalities: Divide and conquer? *Journal of Economic Theory* 113 (2), 147–181.
- Szkup, M., Trevino, I., 2015. Information acquisition in global games of regime change. *Journal of Economic Theory* 160, 387–428.
- Taneva, I. A., 2016. Information design. Working paper, University of Edinburgh.
- Wang, Y., 2015. Bayesian persuasion with multiple receivers. Working Paper.
- Yang, M., 2015. Coordination with flexible information acquisition. *Journal of Economic Theory* 158, 721–738.