

Persuasion in Global Games with Application to Stress Testing*

Nicolas Inostroza
University of Toronto

Alessandro Pavan
Northwestern University and CEPR

March 12, 2020

Abstract

We study robust/adversarial information design in global games. We show that the optimal policy coordinates all market participants on the same course of action. Importantly, while it removes any “strategic uncertainty,” it preserves heterogeneity in “structural uncertainty.” We identify conditions under which the optimal policy is a “pass/fail” test, show that the test may be highly non-monotone in fundamentals, but also identify conditions under which the optimal test is monotone. Finally, we show how the effects of an increase in market uncertainty on the toughness of the optimal stress test depend on the securities issued by the banks.

JEL classification: D83, G28, G33.

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

*Emails: nicolasinostroza2018@u.northwestern.edu, alepavan@northwestern.edu. For comments and useful suggestions, we thank the Editor, anonymous referees, Marios Angeletos, Gadi Barlevy, Eddie Dekel, Tommaso Denti, Laura Doval, Yaron Leitner, Stephen Morris, Marciano Siniscalchi, Bruno Strulovici, Jean Tirole, Xavier Vives, Nora Wagner, and seminar participants at various conferences and institutions where the paper was presented. Matteo Camboni provided excellent research assistance. Pavan acknowledges financial support from the NSF under the grant SES-1730483 and thanks Bocconi University for its hospitality during the 2017-2018 academic year. The usual disclaimer applies.

1 Introduction

Differing opinions on how stress tests should be undertaken are welcome and important...We need to move away from simple pass/fail policies (Piers Haben, Director, European Banking Authority, Financial Times, August 1, 2016).

Coordination plays a major role in many socio-economic environments. The damages to society of mis-coordination can be severe and often call for government intervention. Think of a major financial institution such as MPS (Monte dei Paschi di Siena, the oldest bank on the planet and the Italian third largest) trying to raise capital from multiple investors (mutual funds, creditors, and other major financial institutions), despite concerns about the size of the bank's non-performing loans. A default by an institution such as MPS can trigger a collapse in financial markets, and ultimately a deep recession in the Eurozone and beyond (The Economist, July 7, 2016).

Confronted with such prospects, governments and supervising authorities have incentives to intervene. However, a government's ability to calm the market by injecting liquidity into a troubled bank can be limited. For example, in Europe, legislation passed in 2015 prevents Eurozone member states from rescuing banks by purchasing assets or, more generally, by acting on the banks' balance sheets. In such situations, interventions aimed at influencing market beliefs, for example through the design of stress tests, or other targeted information policies, play a fundamental role. The questions policy makers face in designing such information policies are the following: (a) What disclosures minimize the risk of default? (b) Should all the information collected through the stress tests be passed on to the market, or should the supervising authorities commit to coarser policies, for example, a simple announcement of whether or not a bank under scrutiny passed the tests? (c) Should stress tests pass institutions with strong fundamentals and fail the rest, or are there benefits to non-monotone rules? (d) What are the effects of an increase in market uncertainty on the structure of the optimal tests, and how do they depend on the banks' recapitalization strategies, i.e., on the type of security issued by the banks?

In this paper, we develop a framework that permits us to investigate the above questions. We study the design of optimal information policies in markets in which a large number of receivers (e.g., market investors) must choose whether to play an action favorable to the designer (e.g., pledging to a bank, or refraining from speculating against it), or an "adversarial" action (e.g., refraining from pledging, or engaging in predatory trading, for example by short-selling securities linked to the bank's assets, or buying credit-default swaps, which are known to put strain on illiquid banks). Market participants are endowed with heterogenous private information about relevant economic fundamentals, such as a bank's non-performing loans, the long-term profitability of its assets, or other elements of the bank's balance sheet not in the public domain. A cash-constrained policy maker (e.g., a benevolent government, or a supervising authority such as the European Banking Authority, or the Federal Reserve Bank) can act so as to influence the market's beliefs (for example,

by designing a stress test), but is unable to use financial instruments to influence directly the market outcome.¹

While motivated by the design of stress tests, the analysis delivers results that are relevant also for other applications, including currency crises, technology and standards adoption, and political change.² The novelty relative to the rest of the Bayesian persuasion literature is that we explicitly account for the role that coordination plays among multiple, heterogeneously informed, receivers.³ Coordination plays a key role in the funding of solvent but illiquid banks (see, among others, Diamond and Dybvig (1983) and Goldstein and Pauzner (2015) for runs on deposits, Copeland et al. (2014) and Gorton and Metrick (2012) for runs on repos, Covitz et al. (2013) for runs on asset-backed commercial paper, and Pérignon et al. (2018) for dry-ups on certificates of deposit).

The backbone of our analysis is a global game of regime change in which, prior to receiving information from the *information designer* (the policy maker), each agent is endowed with an exogenous private signal about the strength of the underlying fundamentals. In the absence of additional information, such a game admits a unique rationalizable strategy profile, whereby agents play the action favorable to the policy maker (i.e., pledge to the bank) if, and only if, they assign sufficiently high probability to the underlying fundamentals being strong, and whereby regime change (i.e., default) occurs only for sufficiently weak fundamentals. In such settings, the design of the optimal persuasion strategy must account for 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).⁴ 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 economic fundamentals), and the agents' *strategic uncertainty* (i.e., the agents' beliefs about other agents' behavior).

Another point of departure with respect to most of the Bayesian persuasion literature is that 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 information disclosed, 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

¹For an account of key institutional details of the stress tests conducted in Europe, see, for example, Henry and Christoffer (2013) and Homar et al. (2016).

²For 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.

³The rest of the Bayesian persuasion literature has confined attention to settings with either a single (possibly privately informed) receiver, or multiple but uninformed receivers. See the discussion of the related literature at the end of the Introduction.

⁴See also Mathevet et al. (2019) for similar observations.

departure from both the mechanism design and the persuasion literature, where the designer is typically assumed to be able to coordinate the market on the continuation equilibrium most favorable to her. Given the type of applications the analysis is meant for, such “robust approach” appears more appropriate.⁵

Our first result shows that the optimal policy has the “*perfect coordination property*.” It induces all market participants to take the same action, irrespective of the heterogeneity in the agents’ first- and higher-order (posterior) 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 predict the actions of any other agent, but not the beliefs that rationalize such actions. In the context of our application, an investor who is induced to pledge need not be able to predict whether other investors pledge because they expect the bank’s fundamentals to be so strong that the bank will never collapse, irrespective of what other investors do, or because they expect other investors to pledge.

The optimality of policies satisfying the perfect coordination property should not be taken for granted given the robustness requirement. When the designer trusts her ability to select the agents’ strategy profile most favorable to her, the optimality of the perfect coordination property is fairly straightforward and follows from arguments similar to those establishing the Revelation Principle. This is not the case under adversarial design, for information policies that facilitate perfect coordination may also favor rationalizable profiles in which some of the agents play adversarially (in our application, refrain from pledging).

Our second result identifies primitive conditions under which the optimal policy takes the form of a simple “pass/fail” test, with no further information disclosed to the market. We show that the optimality of such policies hinges on a certain co-movement between fundamentals and beliefs, namely on the property that states of Nature in which the fundamentals are strong are also states in which most agents expect the fundamentals to be strong, other agents to expect the fundamentals to be strong, and so on.⁶ This property is consistent with what is typically assumed in the literature on coordination under incomplete information. Importantly, when such a property is not satisfied, the policy maker may be strictly better off disclosing information to the agents in addition to whether or not the bank passed the test.⁷

⁵If the designer could choose the rationalizable profile, she would fully disclose the fundamentals, and then recommend that all agents pledge, unless the bank is bound to collapse irrespective of the agents’ behavior. This is both uninteresting and unrealistic.

⁶Formally, when the agents’ beliefs are parametrized by a uni-dimensional signal, this amounts to assuming that the distribution from which the signals are drawn is log-supermodular or, equivalently, satisfies the monotone likelihood ratio property.

⁷This is another point of departure with respect to the pertinent literature. When the designer trusts her ability to select the continuation equilibrium, optimal disclosure policies always take the form of action recommendations (and hence pass/fail policies are optimal, irrespective of the agents’ exogenous beliefs). This is not the case under adversarial/robust design. See also Example 1 in the online Supplement and the discussion around it.

The above two results contribute to the debate about the (sub)optimality of European stress tests. Such tests have been criticized for not disclosing the details of the simulations (see, e.g., “Stress tests do little to restore faith in European banks,” *Financial Times*, August 1, 2017). Our results indicate that simple pass/fail policies might actually be optimal. Importantly, optimal stress tests should be *transparent*, in the sense of facilitating coordination among the relevant investors, but should not generate consensus among market participants about the soundness of the financial institutions under scrutiny. Preserving heterogenous beliefs over a bank’s fundamentals is instrumental to the minimization of default risk.

Our third result is about the optimality of monotone rules that pass institutions whose fundamentals are strong and fail those whose fundamentals are weak. We show that such policies need not be optimal. In particular, we show that when (a) the precision of the agents’ exogenous information is high, and (b) the policy maker’s objective is to minimize the probability of default of solvent but potentially illiquid banks, irrespective of which banks are saved and which are failed, then the optimal policy is highly non-monotone over the subset of the “critical region” in which default would happen in the absence of any disclosure (the critical region is the set of fundamentals over which the fate of a bank depends on the market’s behavior). For any fundamental receiving a pass grade there is a nearby fundamental receiving a fail grade. The reason why, under adversarial design, non-monotone policies may outperform monotone ones is that they make it more difficult for the agents to commonly learn the precise fundamentals when hearing that a bank passed the test and hence help reduce the risk of the market responding adversarially to the disclosed information. In turn, this permits the policy maker to give a pass grade to more banks, while guaranteeing that, after a pass grade is announced, the unique rationalizable strategy profile features all agents pledging.⁸

Monotone rules, however, can be optimal when the policy maker has strong preferences over the type of banks that she saves, that is, when saving banks with stronger fundamentals is significantly more valuable than saving those with weaker ones. We identify sharp conditions relating the policy maker’s preferences to the other primitives of the model under which such monotone rules are optimal. We also explain that the conditions that guarantee the optimality of monotone rules are more demanding when the policy maker faces multiple privately-informed receivers than when she faces either a single (possibly privately-informed) receiver, or multiple receivers who possess no exogenous private information.

We also show how the results extend to settings in which the policy maker faces uncertainty about the fate of the financial institutions under scrutiny, for example because default may be determined

⁸See also Goldstein and Leitner (2017) for an alternative reason why non-monotone policies may be optimal. In their environment, the optimality of non-monotone policies stems from the possibility of implementing superior risk sharing. In our environment, instead, from the fact that the designer faces multiple receivers with heterogenous private information playing adversarially to the designer. If the receivers played favorably to the designer, the optimal policy would be a simple monotone pass/fail rule, passing all institutions except those with the weakest fundamentals, for which default is unavoidable.

also by variables orthogonal to, or imperfectly correlated with, those measurable by the policy maker (e.g., by the behavior of noisy/liquidity traders, or by macroeconomic events only imperfectly correlated with the banks' fundamentals).

Lastly, we show how the flexibility of the (general version of the) model favors micro-foundations in which the banks under scrutiny issue equity or debt to fund their short-term liquidity obligations, and where the (market-clearing) price of the securities is endogenous and depends on the information revealed through the stress tests. We use such micro-foundations to show how the model can be used for comparative statics. As an example, we investigate the effects of an increase in market uncertainty on the toughness of the optimal stress tests and show how the latter depends on the type of security issued by the banks.

Throughout the analysis, we restrict attention to situations in which the policy maker is constrained to disclose the same information to all market participants, which appears to be the most relevant case in practice. In the online Supplement, however, we consider the possibility of discriminatory disclosures and discuss why, when feasible, such disclosures may improve upon their non-discriminatory counterparts, and relate the optimality of discriminatory disclosures to the sensitivity of the agents' payoffs to the underlying fundamentals.

The rest of the paper is organized as follows. Below, we wrap up the introduction with a brief review of the pertinent literature. Section 2 presents a stylized version of the model which we use to illustrate the key ideas in the simplest possible framework. Section 3 introduces the general model, discusses the robustness of the results in the stylized model, and provides a foundation for monotone policies. Section 4 shows how the general model of Section 3 admits as special cases micro-foundations that favor comparative statics. i. Section 5 concludes. All proofs are either in the Appendix at the end of the document or in the online Supplement.

(Most) pertinent literature. The paper is related to different strands of the literature. The first strand is the literature on *information design* (see Bergemann and Morris (2019) for an excellent overview). This literature traces back to Myerson (1986), who introduced the idea that, in a general class of multi-stage games of incomplete information, the designer can restrict attention to private incentive-compatible action recommendations to the agents. Recent developments include Rayo and Segal (2010), Kamenica and Gentzkow (2011), Gentzkow and Kamenica (2016), and Ely (2017). These papers consider persuasion with a single receiver. The case of multiple receivers is less studied. Calzolari and Pavan (2006a) consider an auction setting in which the sender is the initial owner of a good and where the different receivers are privately-informed bidders in an upstream market who then resell in a downstream market (see also Dworzak (2019) for an analysis of persuasion in other mechanism design environments with aftermarkets).⁹ More recent papers with multiple receivers include Alonso and Camara (2016a), Bardhi and Guo (2017), Che and Hörner (2018), Doval and

⁹Related is also Calzolari and Pavan (2006b). That paper studies information design in a model of sequential contracting with multiple principals.

Ely (2017), Basak and Zhou (2019), Li et al. (2019), Mathevet et al. (2019), and Taneva (2019). Importantly, these last few papers assume that the receivers are initially uninformed about the underlying fundamentals. Persuasion with ex-ante privately informed receivers has been examined primarily in settings with a single receiver (see, among others, Kolotilin et al. (2017), Alonso and Camara (2016b), Chan et al. (2019), and Guo and Shmaya (2019)). Allowing for multiple privately-informed receivers has important implications for the structure of the optimal policy. As discussed in Sections 2.2 and 3, under the assumptions of our model, the optimal stress test both in the case of a single receiver and in the case of multiple receivers with no exogenous private information is a simple monotone pass/fail test, whereas this is not necessarily the case with multiple privately-informed receivers.

See also Laclau and Renou (2017) and Gitmez and Molavi (2019) for recent papers with multiple privately-informed receivers. These papers though do not look at the implications of (adversarial) coordination among the receivers for the structure of the optimal information policy, which is the focus of the present paper.¹⁰ Bergemann and Morris (2016a) and Bergemann and Morris (2016b) provide a general characterization of the set of outcome distributions that can be sustained in such games as Bayes-Nash equilibria under arbitrary information structures consistent with a given common prior.

The present paper contributes to this literature by characterizing the properties of optimal persuasion policies in settings with multiple receivers who play a (global) coordination game under dispersed information. The approach is also different in that we consider “robust design,” that is, we assume the designer evaluates any disclosure rule on the basis of the rationalizable strategy profile that is least advantageous to her (see also Dworzak and Pavan (2019) and Morris et al. (2019) for recent papers also adopting an adversarial approach).

The second strand is the literature on *stress test design* and *regulatory disclosures* in the financial system. For an excellent overview of this literature see, e.g., Goldstein and Sapra (2014).¹¹ Close in spirit is the work by Goldstein and Leitner (2018). 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.¹² 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. Importantly, we model explicitly the coordination among market participants. Bouvard et al. (2015) study a credit rollover setting where a policy maker must choose between transparency (full disclosure) and opacity (no disclosure) but cannot commit to a disclosure policy. In contrast, we assume the policy maker can fully commit to her disclosure policy and allow for flexible information structures. Alvarez and

¹⁰See also Gick and Pausch (2012), Shimoji (2017), and Arieli and Babichenko (2019). These papers, however, abstract from strategic interactions among the receivers.

¹¹See also Morgan et al. (2014), Flannery et al. (2017), and Petrella and Resti (2013). The first two papers provide evidence that the tests conducted in the US revealed information not already in the market system, whereas the second paper provides similar evidence for stress tests conducted in the EU.

¹²See also Williams (2017) for a related analysis of stress test design in a bank-run model a’ la Allen and Gale (1998), with homogenous investors.

Barlevy (2015) study the incentives of banks to disclose balance sheet (hard) information in a setting where the market is not able to observe how banks are exposed to each others' risks.¹³ Orlov et al. (2017) and Inostroza (2019) consider the joint design of stress tests and capital requirements. The latter paper also considers the interplay between information disclosures and the policy maker's role as a lender of last resort.¹⁴ Closely related is Goldstein and Huang (2016). That paper studies persuasion in a stylized coordination setting similar to the one in our baseline model but restricts the designer to announcing whether or not the fundamentals fall below a given threshold. We allow for flexible information structures, show that optimal stress tests need not be monotone under the assumptions in Goldstein and Huang (2016), but also identify primitive conditions under which, in richer settings, monotone policies are optimal.

The present paper contributes to this literature in a few important ways: (a) it shows that optimal stress tests should not be expected to create conformism in investors' beliefs about banks' fundamentals but should be sufficiently transparent to eliminate any ambiguity about the market response to the tests; (b) it identifies conditions under which optimal stress tests take the form of simple pass/fail announcements; (c) it provides conditions for optimal tests to be monotone; and (d) it discusses how the toughness of optimal stress tests relates to the type of securities issued by the banks.

Finally, the paper is 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. 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. Cong et al. (2016) consider a dynamic setting similar to the one in Angeletos et al. (2007) but allowing for policy interventions. Denti (2020), Szkup and Trevino (2015), Yang (2015) and Morris and Yang (2019) consider global games where, prior to committing their actions, agents acquire information about payoff-relevant variables at a small cost.

2 A Stylized (Global Game) Model of Stress Testing under Adversarial Coordination

To illustrate the key ideas in the simplest possible terms, we start by considering a highly stylized game of regime change in the spirit of Rochet and Vives (2004) that abstracts from many institutional details relevant for stress testing but captures the essence of the effects of (adversarial) coordination

¹³See also Corona et al. (2017) for an analysis of how stress tests disclosures may favor banks' coordinated risk taking in the spirit of Farhi and Tirole (2012).

¹⁴See also Faria-e Castro et al. (2016) and Garcia and Panetti (2017) for a joint analysis of stress tests and government bailouts.

among the bank’s investors on the design of the optimal disclosure policy.¹⁵

2.1 The environment

Players and Actions. A policy maker designs a grading system (a stress test) to evaluate the profitability, the liquidity, and the solvency of a representative bank.

To meet its short-term liquidity obligations, the bank may need funding from the market. The latter is populated by a (measure-one) continuum of investors distributed uniformly over $[0, 1]$. Each investor may either take a “friendly” action, $a_i = 1$, or an “adversarial” action, $a_i = 0$. The friendly action is interpreted as the decision to pledge to the bank (alternatively, to abstain from speculating against the bank by short-selling its assets or by engaging in predatory trading, e.g., by purchasing credit-default swaps). The adversarial action is interpreted as the decision to not pledge (alternatively, to speculate against the bank). We denote by $A \in [0, 1]$ the size of the aggregate pledge.

Fundamentals and Exogenous Information. The bank’s fundamentals are parameterized by $\theta \in \mathbb{R}$. Before the bank is scrutinized, it is commonly believed by the policy maker and the investors alike that θ is drawn from a distribution F , absolutely continuous over $\Theta \supsetneq [0, 1]$, with a smooth density f strictly positive over Θ . In addition, each investor $i \in [0, 1]$ is endowed with a private signal $x_i \in \mathbb{R}$ drawn independently across agents from the conditional cumulative distribution function $P(x|\theta)$ with density $p(x|\theta)$. We denote by $\mathbf{x} \equiv (x_i)_{i \in [0,1]}$ a profile of private signals and by $\mathbf{X}(\theta)$ the collection of all $\mathbf{x} \in \mathbb{R}^{[0,1]}$ that are consistent with the fundamentals be equal to θ . Consistently with the rest of the literature, in this baseline version of the model, we assume that $x_i = \theta + \sigma \varepsilon_i$, with ε_i drawn independently across agents from a standard Normal distribution, thus implying that $P(x|\theta) = \Phi((x - \theta)/\sigma)$, where Φ is the cumulative distribution function of the standard Normal distribution.

Default. The bank’s fundamentals θ parametrize the critical size of the aggregate pledge that is necessary for the bank to continue operating in a profitable manner. If $A > 1 - \theta$, the bank remains profitable. If, instead, $A \leq 1 - \theta$, the bank ends up in distress and defaults. We denote by $r = 0$ the event that the bank defaults, and by $r = 1$ the complement event in which the bank remains profitable.

Dominance Regions. Clearly, for any $\theta \leq 0$ the bank defaults, whereas for any $\theta > 1$ the bank remains profitable, irrespective of the size of the aggregate pledge. For $\theta \in (0, 1]$, instead, the fate of the bank is determined by the behavior of the market.

Payoffs. Each investor’s payoff differential between choosing the friendly and the adversarial action is equal to $g > 0$ in case the bank avoids default (i.e., in case $r = 1$) and $b < 0$ otherwise (i.e., in case $r = 0$).

¹⁵Rochet and Vives (2004) consider a three-period economy a’ la Diamond Dybvig (but with heterogenous investors) in which banks can liquidate assets to boost liquidity and may fail early or late. As shown in that paper, the full model admits a reduced-form version very similar to our benchmark model.

The policy maker’s objective consists in minimizing the probability of default. Her payoff is equal to $W > 0$ in case default is avoided, and $L < 0$ in case of default.

Stress Tests. Let \mathcal{S} be a compact metric space defining the set of possible grades/scores. A *stress test* $\Gamma = (\mathcal{S}, \pi)$ consists of the set \mathcal{S} along with a mapping $\pi : \Theta \rightarrow \mathcal{S}$ specifying, for each θ , the corresponding score.¹⁶

Timing. The sequence of events is the following:

1. The policy maker publicly announces the policy $\Gamma = (\mathcal{S}, \pi)$ and commits to it.
2. The fundamentals θ are drawn from the distribution F and the agents’ exogenous signals $\mathbf{x} = (x_i)_{i \in [0,1]}$ are drawn from the distribution $P(x|\theta)$.
3. The score $s = \pi(\theta)$ is publicly announced.
4. Agents simultaneously choose whether or not to pledge.
5. The fate of the bank is determined and payoffs are realized.

Adversarial Coordination and Robust Design. The policy maker does not trust the market to play according to the strategy profile that is most advantageous to her (i.e., pledge to the bank whenever the latter is solvent, i.e., whenever $\theta > 0$). If she did, a simple monotone test revealing whether or not $\theta > 0$ would be optimal. Instead, the policy maker adopts a robust approach to the design of the stress test. She evaluates any policy Γ under the “worst-case” scenario. That is, given any policy Γ , the policy maker expects the market to play according to the rationalizable action profile most adversarial to her.

Definition 1. Given any policy Γ , the **most aggressive rationalizable profile** (MARP) consistent 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 over all profiles surviving *iterated deletion of interim strictly dominated strategies* (henceforth IDISDS).

In the IDISDS procedure leading to MARP, agents update their beliefs about the fundamentals θ and the other agents’ exogenous signals $\mathbf{x} \in \mathbf{X}(\theta)$ using the common prior, F , the signal distribution, $P(x|\theta)$, the disclosure policy, Γ , and Bayes rule. Under MARP, given (x, s) , each agent $i \in [0, 1]$, after receiving exogenous information x and endogenous information s , then refrains from pledging whenever there exists at least one conjecture over (θ, A) consistent with the above Bayesian updating and supported by all other agents playing strategies surviving IDISDS, under which refraining from pledging is a best response for the individual.

¹⁶In the baseline model, we assume that, through the stress test, the policy maker perfectly learns all information that is relevant for the fate of each bank, for her payoff, and for the payoffs of all market participants. We also confine attention to deterministic policies, which seem the most relevant ones for practical reasons. We relax both these assumptions in Section 3.

Hereafter, we confine attention to policies for which MARP exists. As it will become clear from the analysis below, when this is the case, the strategy profile a^Γ is, in fact, a Bayes-Nash equilibrium (BNE) of the continuation game among the agents given Γ , and minimizes the policy maker's payoff state-by-state, and not just in expectation.

2.2 Properties of Optimal Policies

We now turn to the properties of optimal policies.

2.2.1 Perfect-coordination property

Definition 2. A policy $\Gamma = (\mathcal{S}, \pi)$ satisfies the **perfect-coordination property** if, for any $\theta \in \Theta$, any distribution of exogenous information $\mathbf{x} \in \mathbf{X}(\theta)$, and any pair of individuals $i, j \in [0, 1]$, $a_i^\Gamma(x_i, \pi(\theta)) = a_j^\Gamma(x_j, \pi(\theta))$, where $a^\Gamma = (a_i^\Gamma)_{i \in [0, 1]}$ is the most aggressive rationalizable profile (MARP) consistent with the policy Γ .

Hence, a disclosure policy has the perfect-coordination property if it induces all market participants to take the same action, after any information it discloses. For any $\theta \in \Theta$, let $r^\Gamma(\theta) \in \{0, 1\}$ denote the default outcome when investors play according to a^Γ (i.e., $r^\Gamma(\theta)$ is the fate of the bank that prevails at θ , when the agents play according to MARP consistent with Γ). Observe that, because the exogenous signals x_i are i.i.d. draws from $P(x|\theta)$, and because there is a continuum of agents, any pair of signal realizations $\mathbf{x}, \mathbf{x}' \in \mathbf{X}(\theta)$ consistent with the fundamentals being equal to θ has the same cross-sectional distribution of signals, with the latter equal to $P(x|\theta)$. As a result, the fate of the bank under MARP is the same across any pair of signal realizations $\mathbf{x}, \mathbf{x}' \in \mathbf{X}(\theta)$ and hence depends only on Γ and θ .

Hereafter, we say that the policy Γ is *regular* if MARP under Γ is well defined and the default outcome under a^Γ is measurable in θ .

Theorem 1. *Given any (regular) policy Γ , there exists another (regular) policy Γ^* satisfying the perfect coordination property and such that, for any θ , the default probability under Γ^* is the same as under Γ .*

The policy Γ^* is obtained from the original policy Γ by disclosing, in addition to the score $s = \pi(\theta)$ disclosed by the original policy Γ , the fate of the bank $r^\Gamma(\theta) \in \{0, 1\}$ under MARP consistent with the original policy Γ . That, under Γ^* , it is rationalizable for all agents to pledge when the policy discloses the information $(s, 1)$ and to refrain from pledging when the policy discloses the information $(s, 0)$ is trivial. In fact, the announcement of $(s, 1)$ (alternatively, of $(s, 0)$) makes it common certainty among the agents that $\theta \geq 0$ (alternatively, that $\theta < 1$).

The difficult part of the proof is to show that, when the new policy Γ^* announces $(s, 1)$, the most aggressive strategy profile involves all agents pledging. This is equivalent to establishing that all agents pledging is the *unique* rationalizable profile in the continuation game that starts after

the policy Γ^* announces $(s, 1)$. If, under the original policy Γ , $r^\Gamma(\theta)$ is increasing in θ , the new piece of information that θ is such that $r^\Gamma(\theta) = 1$ is equivalent to the announcement that $\theta > \hat{\theta}$, for some $\hat{\theta}$. In this case, agents update their beliefs about θ upward when receiving the additional information that θ is such that $r^\Gamma(\theta) = 1$. That each agent is more optimistic about the strength of the fundamentals, however, per se does not imply that MARP under the new policy is less aggressive than under the original policy. In fact, the new piece of information changes not only the agent's first-order beliefs about θ but also his higher-order beliefs and these higher-order beliefs matter for the determination of the most-aggressive strategy profile. Furthermore, in more general settings, such as those considered in Section 3, the announcement that θ is such that $r^\Gamma(\theta) = 1$ need not trigger an upward revision of first-order beliefs. This is because MARP under the original policy Γ need not be in monotone strategies, thus implying that $r^\Gamma(\theta)$ need not be monotone in θ .¹⁷

The proof of Theorem 1 in the Appendix shows that that the result of Theorem 1 holds irrespective of whether or not $r^\Gamma(\theta)$ is monotone. It follows from the fact that, at any stage n of the IDISDS procedure, any agent who, under the original policy Γ would have pledged under the most aggressive strategy profile surviving $n - 1$ rounds of deletion, continues to do the same under the new policy Γ^* . We show that this last property in turn follows from the fact that Bayesian updating preserves the likelihood ratio of any two states that are consistent with no default under the original policy Γ . Formally, for any $s \in \pi(\Theta)$ and any pair of states θ' and θ'' such that (a) $\pi(\theta') = \pi(\theta'') = s$ and (b) $r^\Gamma(\theta') = r^\Gamma(\theta'') = 1$, the likelihood ratio of such two states under Γ^* is the same as under the original policy Γ . This property, together with the fact that the announcement that $r^\Gamma(\theta) = 1$ makes it common certainty among the agents that default would have not occurred under MARP consistent with the original policy Γ , then guarantees that, for any agent for whom pledging was optimal under the old policy Γ , pledging is the unique rationalizable action under the new policy Γ^* .

The policy Γ^* thus removes any strategic uncertainty. When $(s, 1)$ (alternatively, $(s, 0)$) is announced, each agent knows that all other agents will pledge (alternatively, will refrain from pledging), irrespective of their exogenous information, and then finds it optimal to do the same. Importantly, while the policy Γ^* removes any strategic uncertainty, it preserves structural uncertainty. Under Γ^* , different agents hold different beliefs about the bank's fundamentals. Preserving heterogeneity in posterior beliefs about θ is key to minimizing the risk of default. In fact, if agents knew the exact fundamentals, then, under the most aggressive rationalizable profile, they would pledge only when it is dominant for them to do so, i.e., when $\theta > 1$. More generally, if agents knew each others' beliefs, the set of fundamentals under which they would pledge when playing according to MARP would be smaller than under the proposed policy. To minimize the risk of default, it is essential that agents who pledge remain uncertain as to whether other agents also pledge because they find it dominant to do so, or because they expect other agents to pledge and hence it becomes iteratively dominant

¹⁷Furthermore, as we discuss in Section 3, in more general settings, the fate of each bank may depend also on variables other than θ for which the policy maker has imperfect information about. In this case, perfect coordination cannot be attained by announcing to the market the predicted fate of the bank under examination.

for them to pledge.

2.2.2 Pass/Fail

Theorem 2. *Given any policy Γ satisfying the perfect coordination property, there exists a binary policy $\Gamma^* = (\{0, 1\}, \pi^*)$ that also satisfies the perfect coordination property and such that, for any θ , the probability of default under Γ^* is the same as under Γ .*

Take any policy $\Gamma = (\mathcal{S}, \pi)$ satisfying the perfect coordination property. Given the result in Theorem 1, without loss of generality, assume that $\Gamma = (\mathcal{S}, \pi)$ is such that $\mathcal{S} = \{0, 1\} \times \hat{S}$, for some measurable set \hat{S} , and is such that (a) when the policy discloses any signal $s = (\hat{s}, 1)$, all agents pledge and default does not happen, whereas (b) when the policy discloses any signal $s = (\hat{s}, 0)$, all agents refrain from pledging and default happens. Given the policy Γ , let $U^\Gamma(x, (s, 1)|k)$ denote the expected payoff differential of an agent with exogenous information x who receives information $(s, 1)$ from the policy maker and who expects all other agents to pledge if and only if their exogenous signal exceeds a cut-off k . In the Appendix, we show that, no matter the shape of the policy Γ , because $P(x|\theta)$ has the *monotone likelihood ratio property* (in short, MLRP), MARP associated with the policy Γ is in cut-off strategies. Hence, each agent's expected payoff differential under MARP can be written as $U^\Gamma(x, (s, 1)|k)$ for some k that depends on s . That the original policy Γ satisfies the perfect-coordination property in turn implies that, for any s , and any k , $U^\Gamma(k, (s, 1)|k) > 0$. That is, no matter the cutoff k , the expected payoff differential of any agent whose private signal x coincides with the cutoff k must be strictly positive. If this was not the case, the continuation game would also admit a rationalizable profile (in fact, a continuation equilibrium) in which some of the agents refrain from pledging, contradicting the fact that pledging irrespectively of x is the unique rationalizable profile following the announcement of $(s, 1)$.

Now consider a policy Γ^* that, for any θ , discloses the same $r(\theta)$ as the original policy Γ but conceals the additional information s . By the law of iterated expectations, because $U^\Gamma(k, (s, 1)|k) > 0$ for all k , then $U^{\Gamma^*}(k, 1|k) > 0$, for all k . The last property implies that the new policy Γ^* also satisfies the perfect-coordination property. The policy maker can thus drop the additional signals s from the original policy Γ and continue to guarantee that, after $r = 1$ is announced, pledging is the unique rationalizable action for all agents. The result in the theorem thus implies that simple pass/fail policies are optimal.

2.2.3 Non-monotonicity

Having established that the optimal policy has a pass/fail structure and perfectly coordinates all market participants on the same course of action, we now turn to the question of which fundamentals receive a pass and which a fail. Given previous work (most notably, Goldstein and Huang, (2016)), one may conjecture that the optimal policy fails institutions with weak fundamentals and passes

those with strong ones. The next result shows that, when the noise in the agents' exogenous signals is small, the optimal policy is “maximally non-monotone” in the following sense. Let $u(\theta, A)$ denote the payoff from pledging when the fundamentals are θ and the size of the aggregate pledge is A . Then let $\theta^{MS} \in (0, 1)$ be implicitly defined by the unique solution to

$$\int_0^1 u(\theta^{MS}, A) dA = 0. \quad (1)$$

The threshold θ^{MS} corresponds to the value of the fundamentals at which an agent who knows θ and holds *Laplacian* beliefs with respect to the measure of agents pledging is indifferent between pledging and not pledging.¹⁸ Importantly, θ^{MS} is independent of the initial common prior F and of the distribution of the agents' signals.

Next, given any policy $\Gamma = (\{0, 1\}, \pi)$, let $D^\Gamma = \{(\underline{\theta}_i, \bar{\theta}_i] : i = 1, \dots, N\}$ denote the partition of $(\underline{\theta}, \theta^{MS}]$ induced by π , with $N \in \mathbb{N}$, $\theta_1 = \underline{\theta}$, and $\theta_N = \theta^{MS}$.¹⁹ Let $d \in D^\Gamma$ denote a generic element of the partition D^Γ and, for any $\theta \in [\underline{\theta}, \theta^{MS}]$, denote by $d^\Gamma(\theta) \in D^\Gamma$ the partition cell that contains θ . Finally, let $M(\Gamma) \equiv \max_{i=1, \dots, N} |\bar{\theta}_i - \underline{\theta}_i|$ denote the *mesh* of D^Γ , that is, the Lebesgue measure of the cell of D^Γ of maximal Lebesgue measure.

Theorem 3. *There exists a scalar $\bar{\sigma} > 0$ and a function $\mathcal{E} : (0, \bar{\sigma}] \rightarrow \mathbb{R}_+$, with $\lim_{\sigma \rightarrow 0^+} \mathcal{E}(\sigma) = 0$, such that, for any $\sigma \in (0, \bar{\sigma}]$, in the game in which the noise in the agents' information is σ , the following is true: given any pass/fail policy Γ satisfying the perfect-coordination property and such that $M(\Gamma) > \mathcal{E}(\sigma)$, there exists another pass/fail policy Γ^* with $M(\Gamma^*) < \mathcal{E}(\sigma)$ that also satisfies the perfect-coordination property and such that the ex-ante probability of default under Γ^* is strictly smaller than under Γ .*

Heuristically, non-monotone policies permit the policy maker to save more banks than monotone policies by making it difficult for the agents to commonly learn the fundamentals when the latter are above 0 but below θ^{MS} and the policy maker announces that the bank passed the test. Intuitively, if the policy maker assigned a pass grade to an interval $[\theta', \theta''] \subset [0, \theta^{MS}]$ of large Lebesgue measure, when σ is small and $\theta \in [\theta', \theta'']$, most agents would receive signals $x_i \in [\theta', \theta'']$. No matter the grade assigned to fundamentals outside the interval $[\theta', \theta'']$, in the continuation game that starts after the policy maker announces that the bank passed the test, most agents receiving signals $x_i \in [\theta', \theta'']$ would then assign high probability to the joint event that $\theta \in [\theta', \theta'']$, that other agents assign high probability to $\theta \in [\theta', \theta'']$, and so on. When this is the case, it is rationalizable for such agents to refrain from pledging. Hence, when σ is small, the only way the policy maker can guarantee that, when $\theta \in [0, \theta^{MS}]$, the agents pledge after hearing a pass grade is by dividing the subset of $[0, \theta^{MS}]$

¹⁸This means that the agent believes that the proportion of agents pledging is uniformly distributed over $[0, 1]$. See Morris and Shin (2006).

¹⁹That is, given π , either (a) $\pi(\theta) = 0$ for all $\theta \in \cup_{i=2k, k \in \mathbb{N}} (\underline{\theta}_i, \bar{\theta}_i]$ and $\pi(\theta) = 1$ for all $\theta \in \cup_{i=2k-1, k \in \mathbb{N}} (\underline{\theta}_i, \bar{\theta}_i]$, or (b) $\pi(\theta) = 1$ for all $\theta \in \cup_{i=2k, k \in \mathbb{N}} (\underline{\theta}_i, \bar{\theta}_i]$ and $\pi(\theta) = 0$ for all $\theta \in \cup_{i=2k-1, k \in \mathbb{N}} (\underline{\theta}_i, \bar{\theta}_i]$.

receiving a pass grade into a collection of disjoint intervals, each of small Lebesgue measure.²⁰

Next, suppose that the intervals $(\underline{\theta}_i, \bar{\theta}_i] \subset (0, \theta^{MS}]$, $i = 1, \dots, N$, receiving a pass grade were far apart, implying that the policy maker fails an interval $[\theta', \theta''] \subset [0, \theta^{MS}]$ of large Lebesgue measure. The proof in the Appendix then shows that, starting from Γ , the policy maker could assign a pass grade to some types in the middle of $[\theta', \theta'']$ and a fail grade to some types to the right of θ'' , in such a way that (a) pledging continues to be the unique rationalizable action for all agents after hearing that the bank passed the test, and (b) the set of fundamentals receiving a pass grade under the new policy is strictly larger than under the original policy. Intuitively, this can be done by leveraging on the fact that those agents who, under the original policy Γ , are close to be indifferent between pledging and not pledging when hearing that the bank passed the test are highly sensitive to the grade the policy assigns to fundamentals close to their own signal but little sensitive to the grade the policy assigns to fundamentals far from their signal.

Formally, suppose that, starting from the original policy Γ , the policy maker assigns a pass grade to types $\theta \in [(\theta' + \theta'')/2, (\theta' + \theta'')/2 + \xi]$ and a fail grade to types $\theta \in [\theta'' + \delta/2, \theta'' + \delta]$, with ξ and δ small and chosen so that the ex-ante probability of a pass grade is the same as under the original policy Γ . Now take any individual with signal $x < (\theta' + \theta'')/2$. Suppose that, under the original policy Γ , the individual pledges and rationalizes such behavior by expecting all individuals with signal above his to also pledge. When σ is small, the individual then expects default to occur only for $\theta < x$. Because the new policy assigns a pass grade to fundamentals $\theta > (\theta' + \theta'')/2$ closer to the individual's own signal than the original policy, and because such fundamentals are associated with no default, the individual's incentives to pledge under the new policy are stronger than under the original one.

Next consider an individual with signal $x \geq \theta'' + \delta$. Suppose again that, under the original policy Γ , such an individual pledges and rationalizes his behavior by expecting all individuals with signal higher than his to also pledge. When σ is small, such individual expects the bank to default only for $\theta < x$. Because the new policy assigns a pass grade to types $\theta < x$ farther away from x than the original policy, and because such fundamentals are associated with default, the individual's incentives to pledge are again stronger under the new policy than under the original one.

In the Appendix, we show that the above two properties in turn imply that, for those individuals with signals $x \notin [(\theta' + \theta'')/2, \theta'' + \delta]$, if pledging was the unique rationalizable action under the old policy Γ then it continues to be the unique rationalizable action under the new policy.

For those agents with signal $x \in [(\theta' + \theta'')/2, \theta'' + \delta]$, instead, the incentives to pledge may be smaller under the new policy than under the original one. However, we show in the Appendix that such individuals could not be close to be indifferent between pledging and not pledging under the original policy Γ . Hence, provided that σ, ξ, δ are small enough, pledging remains the unique

²⁰Formally speaking, a non-monotone policy guarantees that the supports of the agents' posterior beliefs after hearing that the bank passed the test are not connected. Connectedness of the supports facilitates contagion across states and hence favors rationalizable profiles where some agents refrain from pledging.

rationalizable action for such individuals as well.

Building on the above properties, we then show that the policy maker can extend the pass grade to some types to the left of $(\theta' + \theta'')/2$ and to some types to the right of $\theta'' + \delta/2$ while guaranteeing that pledging after hearing a pass grade continues to be the unique rationalizable action for all agents. That the new policy improves over the original one then follows from the fact that the probability of a pass grade is strictly higher under the new policy than under the original one. Iterating the above arguments, one can then show that starting from any policy with a mesh greater than $\mathcal{E}(\sigma)$, where $\mathcal{E}(\sigma)$ is a threshold that depends on the precision of the agents' information, there exists another policy with a mesh smaller than $\mathcal{E}(\sigma)$ that also satisfies the perfect coordination property and such that the probability of default under the new policy is smaller than under the original one.

It is also possible to show that, when σ is small, the policy maker can give a pass grade to almost all $\theta > \theta^{MS}$. Formally, for any $\varepsilon > 0$, there exists $\sigma(\varepsilon)$ such that, for any $\sigma < \sigma(\varepsilon)$ and any pass/fail policy Γ satisfying the perfect coordination property and such that $\pi(\theta) = 0$ for a F -positive measure subset of $(\theta^{MS} + \varepsilon, 1]$, there exists another pass/fail policy $\Gamma' = (\{0, 1\}, \pi')$, with $\pi'(\theta) = 1$ for all $\theta \geq \theta^{MS} + \varepsilon$, that also satisfies the perfect coordination property and such that the policy maker's payoff under Γ' is strictly higher than under Γ .

We thus conclude that, in the baseline model, when noise is small, any “optimal” policy fails all institutions with $\theta < 0$, and passes almost all institutions with $\theta \geq \theta^{MS}$, and is maximally non-monotone over the interval $(0, \theta^{MS})$, in the sense of Theorem 3 above.

Role of multiplicity of receivers and of exogenous private information. At this point, it is worth contrasting the above results to those for economies featuring either a single privately-informed receiver, or multiple receivers possessing no exogenous private information.

Consider first the case of a single privately-informed receiver (of measure 1). The optimal stress test is then a simple monotone pass/fail policy with cut-off $\theta^* = 0$. This is a direct consequence of the fact that, in this baseline model, the policy maker and the receiver have perfectly aligned preferences (they both want to avoid default when default can be avoided). Next, consider a variant of the baseline model in which the receiver's payoff differential between pledging and not pledging is equal to $-g$ in case of default and $-b$ in case of no default, with $g > 0 > b$. Note that these payoff differentials may reflect the possibility that the receiver is a speculator. In case the receiver does not speculate against the bank, his payoff is equal to zero. If, instead, he speculates (for example, by short-selling the bank's securities, as in the micro-foundations in Section 4), his payoff is positive (and equal to g) in case the bank defaults but negative (and equal to b) in case the bank survives. In this variant, the receiver and the designer thus have mis-aligned preferences. Using the results in Guo and Shmaya (2019), it is possible to verify that the optimal stress test in this case has the “*interval structure*”, meaning that each type x of the receiver is induced to play the action favorable to the policy maker (abstain from attacking) over an interval of fundamentals $[\theta_1(x), \theta_2(x)]$, with $\theta_1(x) < 1 < \theta_2(x)$, for all x , with $\theta_1(x)$ decreasing in x and $\theta_2(x)$ increasing in x . Such a policy

requires sending more than two signals, and hence cannot be implemented through a simple pass/fail test. In contrast, when the policy maker faces a continuum of heterogeneously informed receivers with the above payoffs, the optimal stress test has the same properties as in our baseline model: it is a pass-fail policy that is highly non-monotone in θ , for small σ .²¹ Furthermore, the optimal policy does not have the interval structure of Guo and Shmaya (2019), as each receiver with signal x is induced to pledge over a non-connected set of fundamentals. The reason for these differences is that, when the policy maker faces a single receiver, to convince him not to attack, the policy maker must persuade him that the fundamentals are likely to be above 1, in which case the attack is unsuccessful. With multiple receivers, instead, it suffices that the policy maker persuades each receiver that enough other receivers have been persuaded not to attack for each receiver to abstain from attacking. As explained above, this is best accomplished by a highly non-monotone policy that makes it difficult for the receivers to commonly learn that the fundamentals are below θ^{MS} when they are.

Next consider the case of multiple receivers with no exogenous private information. Under MARP, each receiver then plays the friendly action only if it is dominant for him to do so. The optimal policy is then again a simple monotone pass/fail policy with cut-off equal to $\theta^* = 0$ in case of aligned preferences and with cut-off equal to $\theta^* \in (0, 1)$ in case of mis-aligned preferences. The reason why the optimal policy is monotone when the receivers possess no exogenous private information is that the policy maker simply needs to convince each of them that θ is above 1 with sufficiently high probability to induce them to play the friendly action. When, instead, the receivers possess precise private information, the only way the policy maker can discourage them from attacking when $\theta \in (0, \theta^{MS})$ is to make it difficult for them to commonly believe that $\theta \in (0, \theta^{MS})$ when hearing that the bank passed the test which requires using a highly non-monotone policy, as discussed above.²²

3 Robustness and Extensions

3.1 General Model

We now introduce a more flexible model for which the one in the previous section is a special case. As anticipated above, the value of the generalization is twofold: (a) it permits us to discuss the robustness of the results in the baseline model, and (b) it permits us to introduce various enrichments that we expect to be relevant for stress testing, but also for other applications as well (e.g., the possibility that the information that the policy maker may collect through the test be only imperfectly correlated with other variables that are also relevant for the default of the banks, and the possibility that the agents' and the policy maker's payoffs depend on all relevant variables in a rich manner, as in the

²¹This is because, under MARP, all agents play the friendly action if and only if it is iteratively dominant for them to do so, irrespective of the alignment of payoffs.

²²By commonly believe that $\theta \in (0, \theta^{MS})$, we mean that each agent assigns a high probability to the joint event that $\theta \in (0, \theta^{MS})$, that other agents assign high probability to the event that $\theta \in (0, \theta^{MS})$, and so on.

case of the two micro-founded economies in the next section).

The relevant fundamentals are given by (θ, z) . As in the baseline model, θ is commonly believed to be drawn from the distribution F . Given θ , z is drawn from $[\underline{z}, \bar{z}]$ according to the cdf $Q_\theta(z)$, with $Q_\theta(z)$ weakly decreasing in θ , for any z (first order stochastic dominance). The variable θ continues to parametrize the maximal information the policy maker can collect about the relevant fundamentals through the stress test. The only information that the policy maker can collect about z is the one contained in θ . Likewise, any information the agents possess about z is encoded in the signals x they receive about θ . Conditional on θ , the private signals $(x_i)_{i \in [0,1]}$ are i.i.d. draws from an arbitrary (absolutely continuous) cumulative distribution function $P(x|\theta)$, with associated density $p(x|\theta)$ strictly positive over \mathbb{R} . The additional variable z parametrizes risk that the agents and the policy maker face at the time of the stress test (e.g., macroeconomic variables that are only imperfectly correlated with the bank’s fundamentals, and/or the exogenous supply of funds to the bank from sources other than the agents under consideration, as in the applications in the next section).

Both payoffs and default outcomes depend on (θ, A, z) , where A continues to denote the aggregate size of the pledge. There exists a function $R : \Theta \times [0, 1] \times [\underline{z}, \bar{z}] \rightarrow \mathbb{R}$ such that, given (θ, A, z) , default occurs (i.e., $r = 0$) if, and only if, $R(\theta, A, z) \leq 0$. The function R thus implicitly defines the critical pledge A necessary for the bank to avoid default. The function R is continuous, strictly increasing in (θ, z, A) , and such that $R(\underline{\theta}, 1, \bar{z}) = R(\bar{\theta}, 0, \underline{z}) = 0$, for some $\underline{\theta}, \bar{\theta} \in \mathbb{R}$, with $\underline{\theta} < \bar{\theta}$. The thresholds $\underline{\theta}$ and $\bar{\theta}$ thus define the “critical region” $(\underline{\theta}, \bar{\theta}]$ where the fate of the bank depends on the response of the market. The policy maker’s payoff is equal to

$$\hat{U}^P(\theta, A, z) = \begin{cases} \hat{W}(\theta, A, z) & \text{if } r = 1 \\ \hat{L}(\theta, A, z) & \text{if } r = 0. \end{cases} \quad (2)$$

The agents’ payoff differential between playing the “friendly” action (which we continue to interpret as pledging to the bank, or abstaining from speculating against it) and the “adversarial” action (refusing to pledge, or speculating against the bank) is equal to

$$\hat{u}(\theta, A, z) = \begin{cases} \hat{g}(\theta, A, z) & \text{if } r = 1 \\ \hat{b}(\theta, A, z) & \text{if } r = 0. \end{cases}$$

We assume that the payoff differential is positive in case of no default and negative otherwise: $\hat{g}(\theta, A, z) > 0 > \hat{b}(\theta, A, z)$, for any (θ, A, z) . With a slight abuse of notation, for any (θ, A) , we then let $r(\theta, A) \equiv \Pr\{z : R(\theta, A, z) > 0 | \theta\}$ denote the probability the bank avoids default, $g(\theta, A) \equiv \mathbb{E}_{\{z:R(\theta,A,z)>0\}}[\hat{g}(\theta, A, z)|\theta]$ and $b(\theta, A) \equiv \mathbb{E}_{\{z:R(\theta,A,z)\leq 0\}}[\hat{b}(\theta, A, z)|\theta]$ the agents’ expected payoff differential in case of no default and in case of default, respectively, and $L(\theta, A) \equiv \mathbb{E}_{\{z:R(\theta,A,z)\leq 0\}}[\hat{L}(\theta, A, z)|\theta]$ and $W(\theta, A) \equiv \mathbb{E}_{\{z:R(\theta,A,z)>0\}}[\hat{W}(\theta, A, z)|\theta]$ the policy maker’s expected payoff, again in case of default and no default, respectively. With this notation in hands, the agents’

and the policy maker's expected payoffs can then be conveniently expressed as a function of θ and A only, by letting

$$u(\theta, A) \equiv r(\theta, A)g(\theta, A) + (1 - r(\theta, A))b(\theta, A)$$

denote a representative agent's expected payoff differential and

$$U^P(\theta, A) \equiv r(\theta, A)W(\theta, A) + (1 - r(\theta, A))L(\theta, A)$$

denoting the policy maker's expected payoff. We assume that the agents' expected payoff differential $u(\theta, A)$ is weakly increasing in (θ, A) , so that the usual state and action monotonicities apply.

Finally, we allow the policy maker to use stochastic disclosure policies. Let \mathcal{S} be a compact metric space defining the set of possible disclosures to the agents. A (possibly stochastic) *disclosure policy* $\Gamma = (\mathcal{S}, \pi)$ consists of the set \mathcal{S} along with a mapping $\pi : \Theta \rightarrow M(\mathcal{S})$ specifying, for each θ , a lottery whose realization yields the signal disclosed to the agents.

3.2 Results

Given any common posterior $G \in \Delta(\Theta)$ about θ induced by the policy maker through the stress test, let

$$\bar{U}^G(x) \equiv \frac{\int u(\theta, 1 - P(x|\theta))p(x|\theta)dG(\theta)}{\int p(x|\theta)dG(\theta)}$$

denote the expected payoff differential of an agent with signal x who expects all other agents to pledge if and only if their private signal exceeds x . Let ξ^G be the largest solution to $\bar{U}^G(x) = 0$, if such an equation admits a solution, $\xi^G = +\infty$ if $\bar{U}^G(x) < 0$ for all x , and $\xi^G = -\infty$ if $\bar{U}^G(x) > 0$ for all x . Finally, let $\theta^G \equiv \inf \{\theta : u(\theta, 1 - P(\xi^G|\theta)) \geq 0\}$. The interpretation of ξ^G and θ^G is the following. Suppose that the policy maker, through some policy Γ , induces a common posterior G over Θ and that $p(x|\theta)$ is log-supermodular (i.e., satisfies MLRP). Then, in the continuation game that starts after the realization s of the policy Γ induces the common posterior G , MARP is in threshold strategies and is defined by the cut-off ξ^G . The threshold θ^G is then the smallest value of the fundamentals θ under which the agents' expected payoff differential is non-negative, when agents play according to MARP, given the induced posterior G .

Condition PC. *For any distribution $\Lambda \in \Delta(\Delta(\Theta))$ over posterior beliefs consistent with the common prior F (i.e., such that $\int Gd\Lambda(G) = F$), the following condition holds:*

$$\int \left(\int_{-\infty}^{\theta^G} U^P(\theta, 0)dG(\theta) + \int_{\theta^G}^{+\infty} U^P(\theta, 1)dG(\theta) \right) d\Lambda(G) \geq \int \left(\int U^P(\theta, 1 - P(\xi^G|\theta))dG(\theta) \right) d\Lambda(G). \quad (3)$$

Intuitively, Condition PC says that, on average, the loss to the policy maker from having no agent pledging in states $\theta \leq \theta^G$ is more than compensated by the benefit from having all agents pledging in states $\theta > \theta^G$. The average is over both the posteriors induced by the policy maker and the fundamentals.

Formally, given the prior F , consider any distribution Λ over the common posteriors G that the policy maker may induce through her disclosure policy Γ . Clearly, for such a distribution to be feasible, it must satisfy the usual Bayes-plausibility requirement that $\int G d\Lambda(G) = F$. The right-hand side of (3) is the policy maker's expected payoff when the agents play according to MARP. The left-hand side, instead, is the policy maker's expected payoff when, for any (G, θ) , she induces all agents to (a) pledge, irrespective of their private signals, if the agents' expected payoff differential under MARP at (θ, G) is positive, and (b) refrain from pledging, irrespective of their signals, if their expected payoff differential is negative.²³ The condition thus requires that the policy maker's and the agents' payoffs be not too misaligned.

Note that Property PC trivially holds when the policy maker faces no aggregate uncertainty (i.e., when each distribution Q_θ is degenerate), and L is either invariant in A , or decreasing in A . This may reflect the possibility that, in case of default, the policy maker may be indifferent to the size of the pledge, or even prefer fewer agents to pledge (possibly because the social value of their money is higher under alternative investments). More generally, Condition PC also accommodates for the possibility that L is increasing in A , provided that, on average, the sensitivity of the policy maker's payoff to A in case of default is sufficiently small compared to the one in case of no default.

We then have the following result:

Theorem 4. *(a) Given any (possibly stochastic) policy Γ , there exists a (possibly stochastic) policy Γ^* satisfying the perfect coordination property and such that, for any θ , the agents' expected payoff differential under MARP a^{Γ^*} is at least as high as under MARP a^Γ . Furthermore, when, under a^Γ , θ perfectly predicts the default outcome (e.g., when, for any θ , Q_θ is a Dirac measure), the probability of default under Γ^* is the same as under Γ . (b) Suppose that $p(x|\theta)$ is log-supermodular. The policy Γ^* from part (a) is a pass/fail policy. (c) Suppose that, in addition to $p(x|\theta)$ being log-supermodular, Condition PC holds. Then the policy maker's payoff under Γ^* is at least as high as under Γ .*

Consider first part (a). When default depends on variables z the policy maker has limited information about, perfect coordination cannot be induced by communicating to the agents the fate of the bank under MARP consistent with the original policy Γ , as in the case of Theorem 1 in the baseline model. In this case, perfect coordination under the new policy Γ^* is obtained by announcing to the agents, for any θ , the *sign of their expected payoff differential* under MARP consistent with the original policy Γ . The expectation is over z , given the policy maker's limited information θ . Arguments similar to those establishing Theorem 1 in the baseline model then imply that, when, in addition to the information s disclosed by the original policy Γ , the policy maker announces that the agents' expected payoff differential is positive, then, under the unique rationalizable profile, all agents pledge. Similarly, when the policy maker announces that the agents' expected payoff differential is negative, irrespective of their private signals, under MARP, all agents refrain from pledging.

²³Because $u(\theta, 1 - P(\xi^G|\theta))$ is nondecreasing in θ , under MARP consistent with the policy Γ , the agents' expected payoff differential at (G, θ) is negative if $\theta \leq \theta^G$, and positive if $\theta > \theta^G$.

In the special case in which θ is a perfect predictor of the default outcome, because the sign of the agents' payoff differential is determined by the default outcome, the policy Γ^* de facto informs the agents of the regime outcome that would have prevailed at θ under MARP consistent with the original policy Γ , as in the baseline model. In the online Supplement, we show that, in this case, the ability to induce the same default outcome as under the original policy Γ , while coordinating all market participants to take the same action, under MARP, extends to an even richer class of economies. In particular, it extends to economies in which (i) agents' prior beliefs need not be consistent with a common prior, nor be generated by signals drawn independently conditionally on θ , (ii) the number of agents is arbitrary (in particular, finitely many agents), (iii) agents' have a level-K degree of sophistication, (iv) payoffs may be heterogenous across agents, and (v) the designer may disclose different information to different agents. As discussed in the previous sections, the result is not trivial given the robustness requirement.

Next, consider part (b) of Theorem 4. The assumption that $p(x|\theta)$ is log-supermodular is a monotone-likelihood-ratio property (MLRP) which is always satisfied when the signals are additive, i.e., $x_i = \theta + \sigma\varepsilon_i$, with the noise ε_i drawn from a log-concave distribution, as typically assumed in the literature (note that this was the case in the baseline model in the previous section). As shown in the proof of Theorem 2, under such condition, no matter the shape of the policy Γ , in the continuation game that starts after the policy discloses information s , the agents' strategies are monotone in the private signals. This monotonicity property guarantees that the new policy Γ^* that perfectly coordinates the agents does not need to reveal anything more than the sign of the agents' expected payoff differential under MARP consistent with the original policy Γ . That is, the new policy Γ^* can be obtained from the original policy Γ by pooling together all signals s that, given θ , would have led to an expected payoff differential of the same sign. Under MARP consistent with the new policy Γ^* , for any θ , the agents' behavior is then the same as under the original policy Γ , and so are their expected payoffs.

From an economic standpoint, the log-supermodularity of the signal distribution $p(x|\theta)$ implies that states with stronger fundamentals are also states in which more agents hold optimistic beliefs about θ , that is, beliefs that assign higher probability to stronger fundamentals, in the sense of first order stochastic dominance. In the online Supplement, we provide an example showing that this property of beliefs is essential for the optimality of simple pass/fail tests. When it is not satisfied, disclosing information in addition to whether or not a bank passed the test may be essential to induce the desired response by the market.²⁴

Lastly, consider part (c) in Theorem 4. The results in parts (a) and (b) imply that, starting from any (possibly stochastic) policy Γ , there exists a (possibly stochastic) pass/fail policy Γ^* that

²⁴The example also shows why, under adversarial/robust design, restricting attention to policies that take the form of action recommendations is not without loss of optimality, thus also confirming the non-validity of the Revelation Principle.

perfectly coordinates the agents and such that the agents are weakly better off under the policy Γ^* than under the original policy Γ . That the agents are weakly better off, however, does not imply that so is the policy maker. The reason is that, under the new policy Γ^* , at any θ at which, under MARP consistent with the original policy Γ , the agents' expected payoff differential was negative, under the new policy Γ^* , none of the agents pledges, which is something that the policy maker need not like. Condition PC, however, guarantees that the loss from having fewer agents pledging in such states is more than compensated by the benefit from having more agents pledging in those other states in which their expected payoff differential under MARP associated with the original policy Γ was positive. When this is the case, the new policy leads to a Pareto improvement in the sense that that it makes all agents and the policy maker alike better off.

Under the conditions in part (c) of Theorem 4, the optimal policy thus takes the form of a simple pass/fail test. However, the optimal test need not be monotone, i.e., it may fail institutions with stronger fundamentals and pass a few with weaker ones. The reasons why such non-monotone policies may be optimal are similar to those discussed in the baseline model. In the rest of this section, we investigate conditions under which the optimal policy is monotone and deterministic (i.e., it fails with certainty institutions with weak fundamentals and passes, with certainty, those with strong fundamentals).

Condition M. *The following properties hold:*

1. *The function $|u(\theta, 1 - P(x|\theta))|$ is log-supermodular over $\{(\theta, x) \in [\underline{\theta}, \bar{\theta}] \times \mathbb{R} : u(\theta, 1 - P(x|\theta)) \leq 0\}$;*
2. *For any $\theta \in [\underline{\theta}, \bar{\theta}]$, $U^P(\theta, 1) > U^P(\theta, 0)$. Furthermore, for any $x \in \mathbb{R}$ and $\theta_0, \theta_1 \in [\underline{\theta}, \bar{\theta}]$, with $\theta_0 < \theta_1$, such that $u(\theta_1, 1 - P(x|\theta_1)) \leq 0$,*

$$\frac{U^P(\theta_1, 1) - U^P(\theta_1, 0)}{U^P(\theta_0, 1) - U^P(\theta_0, 0)} > \frac{p(x|\theta_1)u(\theta_1, 1 - P(x|\theta_1))}{p(x|\theta_0)u(\theta_0, 1 - P(x|\theta_0))}.$$

Property 1 in Condition M says that the (percentage) increase in the agents' expected payoff differential from higher fundamentals is larger when agents follow a more conservative (monotone) strategy. Formally, for any $\theta' < \theta''$ and $x' < x''$ such that $u(\theta'', 1 - P(x'|\theta'')) < 0$, $u(\theta'', 1 - P(x''|\theta''))/u(\theta', 1 - P(x''|\theta')) > u(\theta'', 1 - P(x'|\theta''))/u(\theta', 1 - P(x'|\theta'))$. Property 2 in turn says that the benefit that the policy maker derives from inducing all agents to pledge starting from a situation in which no agent pledges increases with the fundamentals at a sufficiently high rate, with the critical rate determined by a combination of the agents' payoffs and beliefs.

Theorem 5. *Suppose that $p(x|\theta)$ is log-supermodular, Condition PC holds, and Condition M holds. Given any (possibly stochastic) policy Γ , there exists a deterministic pass/fail monotone policy $\Gamma^* = (\{0, 1\}, \pi^*)$ satisfying the perfect-coordination property and yielding the policy maker a payoff weakly higher than Γ . The policy $\Gamma^* = (\{0, 1\}, \pi^*)$ is defined by a threshold $\theta^* \in [\underline{\theta}, \bar{\theta}]$ such that all $\theta \leq \theta^*$*

receive a fail grade (i.e., $\pi^*(\theta) = 0$), whereas all $\theta > \theta^*$ receive a pass grade (i.e., $\pi^*(\theta) = 1$).²⁵

As we show in the Appendix, under the assumptions in Theorem 5, perturbations of the original policy Γ that swap the probability of inducing all agents to pledge from low to high fundamentals raise the policy maker's payoff, while also preserving the uniqueness of the rationalizable profile after the policy announces that the bank passed the test.

Monotone policies are thus optimal when the value the policy maker derives from saving banks with stronger fundamentals is significantly larger than the value she derives from saving banks with weaker fundamentals. Theorem 5 identifies a sufficient condition under which such higher value is large enough to compensate for the possibility that, from an ex-ante perspective, the probability of default may be larger under monotone policies as discussed after Theorem 3. The condition is fairly sharp in the sense that, when violated, one can construct examples where the optimal policy is non-monotone.²⁶

When Condition M holds, the choice of the optimal policy reduces to the choice of the smallest threshold θ^* such that, when agents commonly learn that $\theta > \theta^*$, under the unique rationalizable profile, all agents pledge, irrespective of their signals. For this to be the case, it must be that, for any $x \in \mathbb{R}$, $\int_{\theta^*}^{\infty} u(\theta, 1 - P(x|\theta))p(x|\theta)f(\theta)d\theta > 0$. The above problem, however, does not have a formal solution, due to the lack of upper-hemicontinuity of the designer's payoff in θ^* .²⁷ Notwithstanding these complications, with abuse, hereafter, we refer to the monotone policy Γ^* with cut-off

$$\theta^* \equiv \inf\{\hat{\theta} : \int_{\hat{\theta}}^{\infty} u(\theta, 1 - P(x|\theta))p(x|\theta)f(\theta)d\theta \geq 0 \text{ for all } x \in \mathbb{R}\} \quad (4)$$

as the ‘‘optimal monotone policy’’. The reason why this is an abuse is that, under the monotone policy with cut-off θ^* , in the continuation game that starts after the policy maker announces that the bank passed the test, there exists a rationalizable profile in which some of the agents refrain from pledging. However, there exists a monotone policy with cut-off arbitrarily close to the threshold θ^* in (4) such that, after the policy maker announces that the bank passed the test, the unique rationalizable profile features all agents pledging. Because the policy maker's payoff under the latter policy is arbitrarily close to the one she obtains when all agents pledge for $\theta > \theta^*$ and refrain from pledging when $\theta \leq \theta^*$, the abuse appears justified.

²⁵Consistently with the notation in the previous section, $\pi(\theta)$ here denotes the deterministic score the policy assigns to the bank when the fundamentals are equal to θ .

²⁶Note that the property of Condition M which is violated under the specification in the baseline model is the second one. In fact, under such a specification, for any $\theta \in [0, 1]$, and any x , the policy maker's payoff differential is equal to $U^P(\theta, 1) - U^P(\theta, 0) = W - L$ which is invariant in θ , whereas the agents' payoff differential is equal to $u(\theta, 1 - P(x|\theta)) = b < 0$ if $\theta < \hat{\theta}(x)$ and to $u(\theta, 1 - P(x|\theta)) = g$ if $\theta > \hat{\theta}(x)$, with the default threshold $\hat{\theta}(x)$ implicitly defined by $P(x|\hat{\theta}(x)) = \hat{\theta}(x)$. As a result, under the specification in the baseline model, for any $x \in \mathbb{R}$, and any $\theta_0, \theta_1 \in [0, 1]$ such that $x > (\theta_0 + \theta_1)/2$ and either (a) $\theta_1 < \hat{\theta}(x)$, or (b) $\theta_0 > \hat{\theta}(x)$, $[U^P(\theta_1, 1) - U^P(\theta_1, 0)] / [U^P(\theta_0, 1) - U^P(\theta_0, 0)] = 1$, whereas, when, x is drawn from a Normal distribution, $[p(x|\theta_1)u(\theta_1, 1 - P(x|\theta_1))] / [p(x|\theta_0)u(\theta_0, 1 - P(x|\theta_0))] = p(x|\theta_1)/p(x|\theta_0) > 1$, thus violating part 2 of Condition M.

²⁷This problem was first noticed in Goldstein and Huang (2016).

4 Micro-foundations and Comparative Statics

We conclude by showing how the general model above accommodates as special cases an economy in which banks fund themselves with equity issuances, and one in which they fund themselves with debt issuances. After showing how these two economies are nested into the general framework of the previous section, we illustrate, by means of an example, how the model can be used for comparative statics analysis.

Consider a representative bank that, at the beginning of period 1, has former liabilities in the amount of D which need to be repaid by the end of the period for the bank to continue operating. The bank has legacy assets that deliver liquid funds $l(\theta) \in \mathbb{R}_+$ at the end of period 1 and, conditional on the bank repaying its period-1 liabilities, a cash flow $V(\theta) \in \mathbb{R}_+$ in period 2. In addition, in case of default, the liquidation of the assets in period 1 delivers an extra cash flow equal to $\gamma(\theta)$, where both V and γ are bounded functions. Additionally, the bank has outstanding shares whose total amount is normalized to 1.

In order to pay for its former liabilities, the bank can either issue new shares or new short-term debt. We study each of these two cases separately. In both cases, we assume that each potential investor is endowed with 1 unit of capital and has to decide whether to "invest" by purchasing the security issued by the bank, "bet against" the bank by short-selling the security, or do nothing. Depending on the case of interest, the decision to do nothing may correspond to the decision to invest in other securities, or, in case of an existing stakeholder, to maintain the existing portfolio. To keep the portfolio decision simple, we assume that each investor is constrained in the position he can take and let that position be normalized to 1. That is, each investor can either buy or sell at most one unit of the security issued (see Albagli et al. (2015) and Brunnermeier and Pedersen (2005) for similar assumptions). We also simplify the analysis by assuming that investors submit market orders (see, e.g., Kyle (1985)). This allows us to abstain from the role of the market as an aggregator of the investors' information, which is interesting but beyond the scope of the analysis here.

Each investor $i \in [0, 1]$ is endowed with an exogenous private signal $x_i = \theta + \sigma\epsilon_i$ of the bank's underlying fundamentals θ , with the noise ϵ_i drawn independently across investors (and independently of θ) from a log-concave distribution.

We also assume that the policy maker confines attention to monotone policies, which, by virtue of Theorem 5, is without loss of optimality provided that the agents' expected payoffs differentials $u(\theta, 1 - P(x|\theta))$ between purchasing and selling the bank's securities are log-supermodular, and the policy maker's payoff satisfies Condition PC and part 2 of Condition M.

Finally, to simplify the exposition, we assume that θ is drawn from an improper uniform prior over \mathbb{R} . This assumption is inconsequential to our results. The agents' hierarchies of beliefs over θ (and hence their expected payoffs) are well defined despite the improperness of the prior. Furthermore, given the focus on monotone rules, optimal policies are also well defined (as discussed at the end of Section 3, a monotone policy is optimal if and only if its threshold θ^* satisfies Condition 4 and such

a condition is well defined despite the impropriety of the prior).

4.1 Equity issuances

The bank issues $q > 1$ new shares at a price p which is determined in equilibrium. After observing their private signals, all investors simultaneously decide whether to submit a *market order* to purchase one share of the bank ($a_i = 1$), short-sell the bank's equity ($a_i = 0$), or do nothing ($a_i = \emptyset$). Let A denote the amount of investors who decide to purchase the shares. We assume that the aggregate demand for the bank's shares is given by $A + Y_E(p, z)$, where $Y_E(p, z)$ represents additional demand coming from sources exogenous to the model (e.g., a combination of high-frequency traders submitting limit-orders and of short-term liquidity traders submitting market orders). The variable z parametrizes residual uncertainty that may correlate with the bank's fundamentals (e.g., the "amount" of liquidity traders, and/or the short-term value the high-frequency traders derive from purchasing the shares). We assume that $Y_E(\cdot, \cdot)$ is a non-increasing function of the price of the bank's shares, p , and a non-decreasing function of z .

We assume that investors are risk-neutral. Along with the fact that investors submit market orders and face constraints on their positions, this last assumption implies that doing nothing is dominated by either purchasing or short-selling a share.²⁸ Because each investor who does not purchase a share, short-sells one, the total supply of shares is thus given by $1 - A + q$, where $1 - A$ is the amount of shares shorted by the investors. It follows that the equilibrium price of the shares, $p_E^*(A, z)$, is implicitly determined by the market-clearing condition

$$1 - A + q = A + Y_E(p, z). \quad (5)$$

Given the monotonicities of Y_E , $p_E^*(A, z)$ is increasing in A and in z , and decreasing in q . Hereafter, we assume that a solution to (5) exists for any (z, A) and is bounded over $(z, A) \in [\underline{z}, \bar{z}] \times [0, 1]$.

The bank avoids default as long as the proceeds from the period-1 equity issuances are sufficient to cover the bank's liabilities D , that is, if and only if,

$$R_E(\theta, A, z) = l(\theta) + \rho_S q p_E^*(A, z) - D > 0,$$

where $l(\theta)$ is the amount of exogenous liquid funds held by the bank at the end of period 1, and ρ_S is the short-term return on the cash $q p_E^*$ collected through the equity issuance.

The investors' payoff differential (between buying and short-selling a share) in case of no default is then equal to

$$\hat{g}_E(\theta, A, z) = 2 \left(\frac{V(\theta) + \rho_L(l(\theta) + \rho_S q p_E^*(A, z) - D)}{1 + q} - p_E^*(A, z) \right)$$

whereas the payoff differential in case of default is equal to

²⁸Whether an investor sells a share he already owns or short-sells one that he borrows makes no difference in this settings. For simplicity, hereafter, we focus on the case of short-selling.

$$\hat{b}_E(\theta, A, z) = -2p_E^*(A, z),$$

where ρ_L is the long-term return on the extra cash $l(\theta) + \rho_S qp_E^* - D$ available to the bank at the end of period 1, after the bank pays its liabilities D . Note that, in writing \hat{g} and \hat{b} , we used the fact that the long-term price of equity is equal to $[V(\theta) + \rho_L(l(\theta) + \rho_S qp_E^* - D)] / (1 + q)$ in case of no default and 0 otherwise.²⁹ That is, in case the bank does not default, the long-term equilibrium price of the bank's shares is equal to the long-term cash flow $V(\theta)$ augmented by the long-term return on the invested funds $\rho_L(l(\theta) + \rho_S qp_E^* - D)$, divided by the amount of outstanding shares, $1 + q$. The payoff from short-selling the bank's share is equal to the negative of the payoff from purchasing the share and, therefore, the payoff differential between the two actions is equal to twice the payoff from purchasing the share.

This economy is thus a special case of the general model in the previous section with the agents' expected payoff differential taking the form of

$$u_E(\theta, A) = -2 \int_{\underline{z}}^{\bar{z}} p_E^*(A, z) dQ_\theta(z) + 2 \int_{\hat{z}_E(\theta, A)}^{\bar{z}} \left[\frac{V(\theta) + \rho_L(l(\theta) + \rho_S qp_E^*(A, z) - D)}{1 + q} \right] dQ_\theta(z) \quad (6)$$

with $\hat{z}_E(\theta, A)$ denoting the critical level of z below which the bank defaults.³⁰ Provided that $u(\theta, A)$ is increasing in both θ and A , all the conclusions from the previous section apply.³¹

4.2 Debt issuances

Next, consider the case of debt issuances. The bank issues $q > 1$ bonds at the beginning of period 1. Each bond is a contract that specifies a payment of F_D in period 2, in case the bank does not default, and covenants L_D that discipline the way the proceeds from liquidation will be divided between old and new debt-holders in case of default. Investors either purchase ($a_i = 1$) or short-sell ($a_i = 0$) one unit of the bond by submitting a market order.³²

Letting A denote the fraction of investors purchasing the bond and p its price, we then have that the total demand for the bond is equal to $A + Y_D(p, z)$, where $Y_D(p, z)$ represents the exogenous

²⁹Implicit in the definition of \hat{b}_E are two assumptions: (a) in case of default, equity is junior to all other existing claims, and (b) the assets' liquidation value $\gamma(\theta)$ is small and hence insufficient to provide any funds to equity holders in case of default.

³⁰Formally, $\hat{z}_E(\theta, A)$ is implicitly defined by the solution to $l(\theta) + \rho_S qp_E^*(A, z) = D$ whenever the equation has a solution, is equal to \underline{z} when $l(\theta) + \rho_S qp_E^*(A, \underline{z}) > D$, and is equal to \bar{z} when $l(\theta) + \rho_S qp_E^*(A, \bar{z}) < D$.

³¹Note that the first term in (6) is decreasing in A , as $p_E^*(A, z)$ is increasing in A , for any z . However, the second term is increasing in A (the integrand is increasing in A and the threshold $\hat{z}_E(\theta, A)$ is decreasing in A). Hence, provided the effects from the second term prevail, $u(\theta, A)$ is increasing in A . When θ and z are independent, the first term is invariant in θ whereas the second term is increasing in θ . When, instead θ and z are positively correlated, the first term may be decreasing in θ (this is the case, for example, when higher a θ implies a FOSD shift in the distribution of z , i.e., when $Q_\theta(z)$ is decreasing in θ , for any z). However, provided the dependence of z on θ is small, $u(\theta, A)$ remains increasing in θ .

³²Investors may also do nothing ($a_i = 0$). As in the case of equity issuances, such a decision is dominated by either purchasing or short-selling one unit of the bond. The arguments are the same as with equity issuances.

(net) demand for the bonds by high-frequency and noisy traders. As in the case of equity issuances, Y_D is non-increasing in p and non-decreasing in z . We assume that, for all $z \in [\underline{z}, \bar{z}]$, if $p > F_D$, then $Y_D(p_D, z) < 0$, which guarantees that the equilibrium price of debt, p_D^* , is smaller than its face value F_D .

As in the case of equity issuances, the bank avoids default as long as its liquid funds at the end of period 1, $l(\theta) + \rho_S q p_D^*$, exceed the amount of former liabilities, D . That is, the bank avoids default if, and only if, $R_D(\theta, A, z) = l(\theta) + \rho_S q p_D^*(A, z) - D > 0$, where the equilibrium price for the newly issued bonds $p_D^*(A, z)$ is implicitly given by the market-clearing condition

$$q + 1 - A = A + Y_D(p_D^*, z). \quad (7)$$

As in the case of equity issuances, we assume that a solution to (7) exists for any (z, A) and is bounded over $(z, A) \in [\underline{z}, \bar{z}] \times [0, 1]$.

The payoff differential between purchasing the bond versus short-selling it is then equal to

$$\hat{g}_D(\theta, A, z) = 2 \left(\min \left\{ F_D, \frac{V(\theta) + \rho_L (l(\theta) + \rho_S q p_D^*(A, z) - D)}{q} \right\} - p_D^*(A, z) \right)$$

in case the bank survives and is equal to

$$\hat{b}_D(\theta, A, z) = 2 \left(\frac{L_D}{q L_D + D} (\gamma(\theta) + l(\theta) + \rho_S q p_D^*(A, z)) - p_D^*(A, z) \right)$$

in case of default. That is, in case the bank is able to repay its short-term liabilities, investors that purchased the bond receive in period 2 the minimum between the bond's face value, F_D , and the bank's cash-flows, net of the repayments of the period-1 liabilities,

$$\frac{1}{q} [V(\theta) + \rho_L (l(\theta) + \rho_S q p_D^*(A, z) - D)].$$

If, instead, the bank is unable to repay its short-term liabilities, and hence defaults, the amount that each debt-holder receives is equal to a fraction $q L_D / [q L_D + D]$ of the total cash $l(\theta) + \rho_S q p_D^*$ available at the end of period 1, plus the additional funds $\gamma(\theta)$ obtained by liquidating the bank's assets, divided by the amount of new debt, q . In other words, the available cash is divided between old and new debt holders in a pro-rated manner. Hereafter, we assume that, in case the bank does not default, the amount of cash $V(\theta)$ generated by the bank's legacy asset in period 2 is sufficiently large to cover the bond's face value F_D for all (θ, A, z) .

This economy too is thus a special case of the general model of the previous section with the agents' expected payoff differential between purchasing and short-selling the bond taking the form of

$$\begin{aligned} u_D(\theta, A) = & -2 \int_{\underline{z}}^{\bar{z}} p_D^*(A, z) dQ_\theta(z) + 2F_D \int_{\hat{z}_D(\theta, A)}^{\bar{z}} dQ_\theta(z) \\ & + 2 \int_{\underline{z}}^{\hat{z}_D(\theta, A)} \frac{L_D}{q L_D + D} (\gamma(\theta) + l(\theta) + \rho_S q p_D^*(A, z)) dQ_\theta(z) \end{aligned} \quad (8)$$

with the critical level $\hat{z}_D(\theta, A)$ below which the bank defaults defined as in the case of equity issuances. Provided that $u(\theta, A)$ is increasing in both θ and A , all the conclusions from the previous section apply.³³

4.3 Comparative statics

The general model described above can also be used for comparative-statics analysis relating the properties of optimal stress tests to the primitives of the model. In this last subsection, we illustrate such a possibility by showing how the toughness of optimal stress tests is affected by an increase in risk about the bank's fundamentals. Specifically, we use the two micro-foundations above to investigate how an increase in risk (formally captured by an increase in the parameter σ scaling the noise in the agents' signals $x_i = \theta + \sigma\epsilon_i$ about the underlying fundamentals) affects the critical threshold θ^* below which the policy maker fails the bank under examination. To make the arguments as simple as possible, we also assume that (a) θ and z are independent, (b) the short-term gross return on cash is one (i.e., $\rho_S = 1$), (c) the noise in the agents' signals is Gaussian (that is, ϵ_i is drawn from a standard Normal distribution, independently of (θ, z)), and (d) the bank's short-term exogenous liquidity is invariant in θ (that is, $l(\theta) = l$ for all θ , with $l \in \mathbb{R}_+$).

Let $\theta_E^*(\sigma)$ and $\theta_D^*(\sigma)$ denote the thresholds that characterize the optimal monotone policies when the precision of the agents' exogenous information is σ^{-2} and the bank funds itself with equity and debt, respectively. Also let θ_E^{MS} and θ_D^{MS} be the Laplacian thresholds for the two economies under consideration (defined by $\int_0^1 u_h(\theta_h^{MS}, A)dA = 0$, $h = E, D$) and recall that, in the absence of any informative policy, when $\sigma \rightarrow 0^+$, purchasing (alternatively, short-selling) security h is the unique rationalizable action for $x > \theta_h^{MS}$ (alternatively, for $x < \theta_h^{MS}$). Lastly, let $\theta_h^\#$ be defined by the solution to $u_h(\theta_h^\#, 1/2) = 0$, $h = D, E$, and note that the threshold $\theta_h^\#$ is the critical value of the fundamentals at which an investor who knows the fundamentals and expects 1/2 of the other investors to buy security h and 1/2 to short-sell it is indifferent between purchasing and short-selling the security, for $h = E, D$. Hereafter, we assume that the problem the policy maker faces is "severe" in the sense that $\theta_h^{MS} \geq \theta_h^\#$, $h = E, D$. We then have the following result:

Proposition 1. *Suppose the environment satisfies the assumptions above. There exists $\bar{\sigma} > 0$ such that the following is true for any $\sigma, \sigma' \in (0, \bar{\sigma}]$, with $\sigma' > \sigma$: $\theta_E^*(\sigma') < \theta_E^*(\sigma)$ and $\theta_D^*(\sigma') > \theta_D^*(\sigma)$.*

The result in the proposition says the following. Take an economy in which the precision in the agents' exogenous information is sufficiently high (that is, σ is small) and consider the effect of an increase in risk (formally captured by the transition to $\sigma' > \sigma$) on the toughness of the optimal stress test (formally captured by the threshold θ_h^* below which the policy maker fails the bank when the latter funds itself by issuing security $h = D, E$). More risk leads to a reduction in the toughness of

³³As in the case of equity issuances, that $u(\theta, A)$ is increasing in (θ, A) may well be consistent with the property that the first term in (8) is always decreasing in A and is decreasing in θ if θ and z are positively affiliated.

the optimal stress test when the bank finances itself with equity but to an increase in the toughness of the optimal stress test when the bank finances itself with debt.

Intuitively, the reason why, under the assumed specification, risk is beneficial to the bank in case of equity financing but not in case of debt financing is the following. Under equity financing, investors are exposed to variations in fundamentals primarily through upside risk. Their payoff differential (between purchasing and short-selling equity) is increasing in θ in case of no default and is equal to $-2p_E^*$ in case of default. Provided that the price of equity does not vary much with (θ, z) , which is the case under the assumed specification, an increase in risk then makes investors more willing to purchase equity. The policy maker can then decrease the critical threshold θ_E^* below which she fails the bank while guaranteeing that, after announcing that the bank passed the test, the unique rationalizable profile continues to feature all investors pledging by purchasing equity.

Under debt financing, instead, investors are exposed to variations in fundamentals primarily through downside risk. When the liquidation value $\gamma(\theta)$ is increasing in θ and p_D^* is not very sensitive to (θ, z) , investors' payoff differential (between purchasing and short-selling debt) is increasing in θ in case of default but constant in fundamentals in case the bank survives. An increase in risk then makes investors less willing to pledge. The policy maker must then increase the critical threshold θ_D^* below which she fails the bank if she wants to guarantee that, after announcing that the bank passed the test, the unique rationalizable profile features all investors pledging by purchasing the newly issued debt.

That risk is beneficial to the bank in case of equity financing but detrimental in case of debt financing need not extend to alternative specifications of the investors' payoffs under the two securities. What seems true more generally is the following single-crossing property. Whenever more risk is beneficial to the bank in case of debt financing, the same tends to be true under equity financing.³⁴

5 Conclusions

The results in the present paper show that, in a large class of economies in which the market cannot be trusted to coordinate on a desirable course of action (e.g., pledging to a solvent but illiquid bank), the optimal persuasion policy completely removes any strategic uncertainty, while retaining 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. The results also show that, when fundamentals and beliefs co-move, in the sense of MLRP, the optimal policy takes the form of a simple pass/fail test. However, the optimal policy need not be monotone in the fundamentals. It is monotone when the value the policy maker assigns to saving institutions with strong fundamentals (relative to the value she assigns to saving institutions with weak fundamentals) is large enough to

³⁴Clearly, for this single-crossing result to hold, one needs to make sure that the two cases are comparable. This requires, among other things, that the exogenous demand for the bank's securities is the same in the two cases, i.e., that $Y_D(p, z) = Y_E(p, z)$, for any (p, z) .

compensate for the possibility that the ex-ante probability of default may be larger under a monotone rule.

The above results are worth extending in a few directions. The analysis in the present paper assumes the policy maker knows how the distribution of market beliefs correlates with the banks' fundamentals. Such knowledge may come from previous experience with banks of similar characteristics, polls, data on professional forecasters, the IOWA betting markets, and the like. While this is a natural starting point, there are many environments in which it is more appropriate to assume that the designer is ambiguous about the joint distribution of the underlying fundamentals and market beliefs. In future work, it would be interesting to investigate the structure of the optimal policy in such situations.³⁵

The analysis in the present paper is also static. However, many of the applications of interest are intrinsically dynamic, with agents coordinating on multiple attacks and learning over time (see the discussion in Angeletos et al. (2007)). In future work, it would be interesting to extend the characterization of optimal persuasion policies 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.³⁶

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 banks, 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 banks prior to communicating with the market. This creates an interesting screening+persuasion problem in the spirit of what is examined in the literature on privacy in sequential contacting (e.g., Calzolari and Pavan (2006a) Calzolari and Pavan (2006b), and Dworzak (2019)).³⁷

Appendix

Proof of Theorem 1. Given any regular policy $\Gamma = (\mathcal{S}, \pi)$ and any $n \in \mathbb{N}$, let $T_{(n)}^\Gamma$ be the set of strategies surviving n rounds of IDISDS, with $T_{(0)}^\Gamma$ denoting the entire set of strategy profiles $a = (a_i(\cdot))_{i \in [0,1]}$, where for any $i \in [0, 1]$, $a_i : \mathbb{R} \times \mathcal{S} \rightarrow [0, 1]$, with $a_i(x, s)$ denoting the probability agent i pledges, given (x, s) . Let $a_{(n)}^\Gamma \equiv \left(a_{(n),i}^\Gamma(\cdot) \right)_{i \in [0,1]} \in T_{(n)}^\Gamma$ denote the most aggressive profile surviving n rounds of IDISDS (that is, the profile in $T_{(n)}^\Gamma$ that minimizes the policy maker's ex-ante payoff). The profiles $\left(a_{(n)}^\Gamma \right)_{n \in \mathbb{N}}$ can be constructed inductively as follows. The profile $a_{(0)}^\Gamma \equiv \left(a_{(0),i}^\Gamma(\cdot) \right)_{i \in [0,1]}$

³⁵For some recent work in this direction, see Dworzak and Pavan (2019).

³⁶For some recent work in this direction, see Basak and Zhou (2020).

³⁷See also Inostroza (2019) for recent developments in this direction.

prescribes that all agents refrain from pledging, irrespective of (x, s) . Next, let $U_i^\Gamma(x_i, s; a)$ denote the payoff differential between pledging and not pledging for agent i when, under Γ , all other agents follow the strategy in a . Then, $a_{(n),i}^\Gamma(x_i, s) = 0$ if $U_i^\Gamma(x_i, s; a_{(n-1)}^\Gamma) \leq 0$ and $a_{(n),i}^\Gamma(x_i, s) = 1$ if $U_i^\Gamma(x_i, s; a_{(n-1)}^\Gamma) > 0$. MARP consistent with Γ is then the profile $a^\Gamma = (a_i^\Gamma(\cdot))_{i \in [0,1]}$ given by $a_i^\Gamma(\cdot) = \lim_{n \rightarrow \infty} a_{(n),i}^\Gamma(\cdot)$, all $i \in [0, 1]$.

Next, consider the policy $\Gamma^+ = (\mathcal{S}^+, \pi^+)$, $\mathcal{S}^+ \equiv \mathcal{S} \times \{0, 1\}$, that, for each θ , discloses the same score $\pi(\theta)$ as the policy Γ , along with the default outcome $r^\Gamma(\theta)$ that would have prevailed at θ when agents play according to MARP consistent with Γ ; that is, for any θ , $\pi^+(\theta) = (\pi(\theta), r^\Gamma(\theta))$.

Define $T_{(n)}^{\Gamma^+}$ and $a_{(n)}^{\Gamma^+}$ analogously to $T_{(n)}^\Gamma$ and $a_{(n)}^\Gamma$, but with respect to the policy Γ^+ .

Step 1. First, we prove that, for each $i \in [0, 1]$,

$$\{(x_i, s) : U_i^\Gamma(x_i, s; a) > 0 \forall a\} \subseteq \{(x_i, s) : U_i^{\Gamma^+}(x_i, (s, 1); a) > 0 \forall a\}.$$

That is, any agent i who, under Γ , finds it dominant to pledge, given the information (x_i, s) , also finds it dominant to pledge under Γ^+ when receiving information $(x_i, (s, 1))$.

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

$$\{(x_i, s) : U_i^\Gamma(x_i, s; a) > 0 \forall a\} = \{(x_i, s) : U_i^\Gamma(x_i, s; a_{(0)}^\Gamma) > 0\}$$

$$\text{and } \{(x_i, s) : U_i^{\Gamma^+}(x_i, (s, 1); a) > 0 \forall a\} = \{(x_i, s) : U_i^{\Gamma^+}(x_i, (s, 1); a_{(0)}^{\Gamma^+}) > 0\}.$$

Now let $\Lambda_i^\Gamma(\theta, \mathbf{x}|x_i, s)$ denote the beliefs of agent $i \in [0, 1]$ over the fundamentals, θ , and the cross-sectional distribution of signals, $\mathbf{x} \in \mathbb{R}^{[0,1]}$, when receiving information $(x_i, s) \in \mathbb{R} \times \mathcal{S}$ under Γ , and $\Lambda_i^{\Gamma^+}(\theta, \mathbf{x}|x_i, (s, 1))$ the corresponding beliefs under Γ^+ . Bayesian updating implies that

$$\partial \Lambda_i^{\Gamma^+}(\theta, \mathbf{x}|x_i, (s, 1)) = \frac{\mathbb{I}_{\{r^\Gamma(\theta)=1\}}}{\Lambda_i^\Gamma(1|x_i, s)} \partial \Lambda_i^\Gamma(\theta, \mathbf{x}|x_i, s), \quad (9)$$

where $\mathbb{I}_{\{r^\Gamma(\theta)=1\}}$ is the indicator function, taking value 1 if $r^\Gamma(\theta) = 1$, and 0 otherwise, and where

$$\Lambda_i^\Gamma(1|x_i, s) \equiv \int_{\{(\theta, \mathbf{x}) : r^\Gamma(\theta)=1\}} d\Lambda_i^\Gamma(\theta, \mathbf{x}|x_i, s).$$

Next, observe that, under both $a_{(0)}^\Gamma$ and $a_{(0)}^{\Gamma^+}$, default occurs if, and only if, $\theta \leq 1$. Take any $i \in [0, 1]$ and $(x_i, s) \in \mathbb{R} \times \mathcal{S}$ such that

$$U_i^\Gamma(x_i, s; a_{(0)}^\Gamma) = \int_{(\theta, \mathbf{x})} (b\mathbb{I}_{\{\theta \leq 1\}} + g\mathbb{I}_{\{\theta > 1\}}) d\Lambda_i^\Gamma(\theta, \mathbf{x}|x_i, s) > 0. \quad (10)$$

The aforementioned property of Bayesian updating implies that

$$\begin{aligned} U_i^{\Gamma^+}(x_i, (s, 1); a_{(0)}^{\Gamma^+}) &= \frac{1}{\Lambda_i^{\Gamma^+}(1|x_i, (s, 1))} \int_{(\theta, \mathbf{x})} (b\mathbb{I}_{\{\theta \leq 1\}} + g\mathbb{I}_{\{\theta > 1\}}) \mathbb{I}_{\{r^\Gamma(\theta)=1\}} d\Lambda_i^\Gamma(\theta, \mathbf{x}|x_i, s) \\ &\geq \frac{1}{\Lambda_i^\Gamma(1|x_i, s)} \int_{(\theta, \mathbf{x})} (b\mathbb{I}_{\{\theta \leq 1\}} + g\mathbb{I}_{\{\theta > 1\}}) d\Lambda_i^\Gamma(\theta, \mathbf{x}|x_i, s) = \frac{1}{\Lambda_i^\Gamma(1|x_i, s)} U_i^\Gamma(x_i, s; a_{(0)}^\Gamma) > 0, \end{aligned}$$

where the first equality follows from (9), the first inequality from the fact that, for all θ such that $r^\Gamma(\theta) = 0$, $b\mathbb{I}_{\{\theta \leq 1\}} + g\mathbb{I}_{\{\theta > 1\}} = b < 0$, the second equality follows from the definition of $U_i^\Gamma(x_i, s; a_{(0)}^\Gamma)$, and the second inequality from (10).

This means that any agent for whom pledging was dominant after receiving information (x_i, s) under Γ , continues to find it dominant to pledge after receiving information $(x_i, (s, 1))$ under Γ^+ .

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

$$\left\{ (x_i, s) : U_i^\Gamma(x_i, s; a) > 0 \quad \forall a \in T_{(k-1)}^\Gamma \right\} \subseteq \left\{ (x_i, s) : U_i^{\Gamma^+}(x_i, (s, 1); a) > 0, \quad \forall a \in T_{(k-1)}^{\Gamma^+} \right\}. \quad (11)$$

Arguments similar to those establishing the result in Step 1 above imply that

$$\left\{ (x_i, s) : U_i^\Gamma(x_i, s; a) > 0 \quad \forall a \in T_{(n-1)}^\Gamma \right\} \subseteq \left\{ (x_i, s) : U_i^{\Gamma^+}(x_i, (s, 1); a) > 0, \quad \forall a \in T_{(n-1)}^{\Gamma^+} \right\}. \quad (12)$$

Intuitively, the result follows from the combination of the following facts: (a) because the game is supermodular, $\left\{ (x_i, s) : U_i^\Gamma(x_i, s; a) > 0 \quad \forall a \in T_{(n-1)}^\Gamma \right\} = \left\{ (x_i, s) : U_i^\Gamma(x_i, s; a_{(n-1)}^\Gamma) > 0 \right\}$ and the same property holds for Γ^+ ; (b) $a_{(n-1)}^{\Gamma^+}$ is “less aggressive” than $a_{(n-1)}^\Gamma$, in the sense that any agent who, given (x, s) , pledges under $a_{(n-1)}^\Gamma$ also pledges under Γ^+ when receiving information $(x, (s, 1))$; and (c) the observation that $r^\Gamma(\theta) = 1$ removes from the support of the agents’ posterior beliefs states in which default would have occurred under a^Γ and hence under $a_{(n-1)}^\Gamma$ as well (observe that $a_{(n-1)}^\Gamma$ is more aggressive than a^Γ , meaning that any agent who, given (x, s) , pledges under $a_{(n-1)}^\Gamma$, also pledges under a^Γ when receiving the same information $(x, 1)$).

Step 3. Equipped with the results in steps 1 and 2 above, we now prove that, for all $\theta \in \Theta$ such that $r^\Gamma(\theta) = 1$, all $\mathbf{x} \in \mathbf{X}(\theta)$, all $i \in [0, 1]$, $a_i^{\Gamma^+}(x_i, (\pi(\theta), 1)) \equiv \lim_{n \rightarrow \infty} a_{(n),i}^{\Gamma^+}(x_i, (\pi(\theta), 1)) = 1$. This follows directly from the fact that, as shown above,

$$a_i^\Gamma(x_i, \pi(\theta)) = 1 \Rightarrow a_i^{\Gamma^+}(x_i, (\pi(\theta), 1)) = 1. \quad (13)$$

The announcement that θ is such that, no matter $\mathbf{x} \in \mathbf{X}(\theta)$, $r^\Gamma(\theta) = 1$ thus reveals to each agent that (θ, \mathbf{x}) is such that, when agents play according to a^{Γ^+} , default does not occur. Because the payoff from pledging is strictly positive when default does not occur, any agent i receiving information $(x_i, 1)$ under Γ^+ thus necessarily pledges. Under the new policy Γ^+ , all agents pledge when they learn that θ is such that $r^\Gamma(\theta) = 1$. That they all refrain from pledging when they learn that θ is such that $r^\Gamma(\theta) = 0$ follows from the fact that such announcement makes it common certainty that $\theta \leq 1$.

We conclude that the new policy Γ^+ satisfies the perfect coordination property and is such that, for any θ , the probability of default under Γ^+ is the same as under Γ . The result in the theorem then follows by taking $\Gamma^* = \Gamma^+$. Q.E.D.

Proof of Theorem 2. The proof is in 2 steps. Step 1 shows that, when $p(x|\theta)$ is log-supermodular, i.e., it satisfies MLRP (which is the case when $x_i = \theta + \sigma\varepsilon_i$, with ε_i drawn from a standard Normal distribution), then, irrespective of Γ , MARP is in cut-off strategies. Step 2 then

shows that, starting from any Γ satisfying the perfect-coordination property, one can drop any signal other than the predicted fate of the bank without changing the agents' behavior. In the baseline model, we have restricted attention to deterministic policies. The result, however, extends to stochastic policies. Because we accommodate for stochastic policies in Section 3, the proof below allows for the possibility that $\Gamma = (\mathcal{S}, \pi)$ is stochastic, that is, for any θ , $\pi(\theta) \in \Delta(\mathcal{S})$.

Step 1. Fix an arbitrary policy $\Gamma = (\mathcal{S}, \pi)$ and, for any pair $(x, s) \in \mathbb{R} \times S$, let $\Lambda^\Gamma(\theta|x, s)$ represent the endogenous posterior beliefs about θ of each agent receiving exogenous information x and endogenous information s . Also recall that $u(\theta, A)$ is the payoff differential between pledging and not pledging when the fundamentals are θ and the aggregate size of the pledge is A .

Next, let $U^\Gamma(x, s|k) = \int_{-\infty}^{\infty} u(\theta, 1 - P(k|\theta)) d\Lambda^\Gamma(\theta|x, s)$ denote the expected payoff differential of an agent with information (x, s) , when all other agents follow a cut-off strategy with cut-off k (i.e., they pledge if, and only if, their private signal exceeds k). The following result (whose proof is in the online Supplement and follows from monotone comparative statics results) establishes that, when the distribution $p(x|\theta)$ from which the signals are drawn satisfies MLRP, no matter Γ , MARP is in cut-off strategies:

Lemma 1. Suppose that $p(x|\theta)$ is log-supermodular. Given any policy $\Gamma = (\mathcal{S}, \pi)$, for any $s \in \mathcal{S}$, there exists $\xi^{\Gamma;s} \in \mathbb{R}$ such that MARP consistent with Γ is given by the strategy profile $a^\Gamma \equiv (a_i^\Gamma)_{i \in [0,1]}$ such that, for any $s \in \mathcal{S}$, $x \in \mathbb{R}$, $i \in [0, 1]$, $a_i^\Gamma(x, s) = \mathbb{I}\{x > \xi^{\Gamma;s}\}$ with $\xi^{\Gamma;s} \equiv \sup\{x : U^\Gamma(x, s|x) \leq 0\}$, if $\{x : U^\Gamma(x, s|x) \leq 0\} \neq \emptyset$, and $\xi^{\Gamma;s} \equiv -\infty$ otherwise. Moreover, the strategy profile a^Γ is a BNE of the continuation game that starts with the announcement of the policy Γ .

Step 2. Now take any policy $\Gamma = (\mathcal{S}, \pi)$ satisfying the perfect coordination property. Given the result in Theorem 1, without loss of generality, assume that $\Gamma = (\mathcal{S}, \pi)$ is such that $\mathcal{S} = \{0, 1\} \times \hat{S}$, for some measurable set \hat{S} , and is such that (a) when the policy discloses any signal $s = (\hat{s}, 1)$, all agents pledge and default does not happen, whereas (b) when the policy discloses any signal $s = (\hat{s}, 0)$, all agents refrain from pledging and default happens.

Equipped with the result in Lemma 1, we then show that, starting from $\Gamma = (\mathcal{S}, \pi)$, one can construct a binary policy $\Gamma^* = (\{0, 1\}, \pi^*)$ also satisfying the perfect coordination property and such that the fate of the bank under Γ^* is the same as under Γ . The policy $\Gamma^* = (\{0, 1\}, \pi^*)$ is such that, for any θ , $\pi^*(1|\theta) = \int_{\hat{S}} d\pi(\hat{s}, 1|\theta)$. That is, at each θ , the binary policy Γ^* recommends to pledge (equivalently, announces a “pass” grade) with the same total probability the original policy Γ discloses signals leading all agents to pledge.

We now show that, under Γ^* , when the policy announces that $s^* = 1$, the unique rationalizable action for each agent is to pledge. To see this, let $U^{\Gamma^*}(x, 1|k)$ denote the expected payoff differential for any agent with private signal x , when the policy Γ^* announces $s^* = 1$, and all other agents follow a cut-off strategy with cut-off k . From the law of iterated expectations, we have that

$$U^{\Gamma^*}(x, 1|k) = \int_{\hat{S}} U^\Gamma(x, (\hat{s}, 1)|k) d\zeta^\Gamma(\hat{s}|x, 1) \quad (14)$$

where $\varsigma^\Gamma(\cdot|x, 1)$ is the probability distribution over \hat{S} obtained by conditioning on the event $(x, 1)$, under the policy Γ . For any signal $s = (\hat{s}, 1)$ in the range of π , MARP consistent with Γ is such that $a_i^\Gamma(x, (\hat{s}, 1)) = 1$ all $x \in \mathbb{R}$, meaning that pledging is the unique rationalizable action after the policy Γ announces $s = (\hat{s}, 1)$. Lemma 1 in turn implies that, for all $s = (\hat{s}, 1)$ in the range of π , $\hat{s} \in \hat{S}$, all $k \in \mathbb{R}$, $U^\Gamma(k, (\hat{s}, 1)|k) > 0$. From (14), we then have that, for all all $k \in \mathbb{R}$, $U^{\Gamma^*}(k, 1|k) > 0$. In turn, this implies that, given the new policy Γ^* , when $s^* = 1$ is disclosed, under the unique rationalizable profile, all agents pledge, that is, $a_i^{\Gamma^*}(x, 1) = 1$ all x , all $i \in [0, 1]$. It is also easy to see that, when the policy Γ^* discloses the signal $s^* = 0$, it becomes common certainty among the agents that $\theta \leq 1$. Hence, under MARP consistent with Γ^* , after $s^* = 0$ is disclosed, all agents refrain from pledging, irrespective of their private signals. The new pass/fail policy Γ^* so constructed thus (a) satisfies the perfect-coordination property, and (b) is such that, for any θ , the probability of default under Γ^* is the same as under Γ . Q.E.D.

Proof of Theorem 3. The proof below contains the main conceptual steps. The complete proof with all technical details is in the online Supplement.

For any $\theta \in (0, 1)$, any $\sigma \in \mathbb{R}_+$, let $x_\sigma^*(\theta) \equiv \theta + \sigma\Phi^{-1}(\theta)$ denote the value of the private signal such that, when every agent $i \in [0, 1]$ pledges for $x_i > x_\sigma^*(\theta)$ and does not pledge for $x_i < x_\sigma^*(\theta)$, default occurs when the fundamentals are below θ and does not occur when they are above θ .³⁸ Also let $x_\sigma^*(0) \equiv -\infty$ and $x_\sigma^*(1) \equiv +\infty$.

For any $(\theta_0, \hat{\theta}, \sigma) \in (0, 1) \times \mathbb{R} \times \mathbb{R}_+$, let $\psi(\theta_0, \hat{\theta}, \sigma)$ denote the payoff from pledging of an agent with private signal $x_\sigma^*(\theta_0)$, when default occurs for $\theta \leq \theta_0$ and does not occur for $\theta > \theta_0$, the policy reveals that $\theta \geq \hat{\theta}$, and the precision of private information is σ^{-2} . Then let

$$\hat{\sigma} \equiv \inf \{ \sigma \in \mathbb{R}_+ : \psi(\theta_0, 0, \sigma) > 0 \text{ all } \theta_0 \in (0, 1) \}$$

if $\{ \sigma \in \mathbb{R}_+ : \psi(\theta_0, 0, \sigma) > 0 \text{ all } \theta_0 \in (0, 1) \} \neq \emptyset$ and $\hat{\sigma} = +\infty$ otherwise.³⁹

Then let $\Psi(\sigma) \equiv \inf_{\theta_0 \in (0, 1)} \psi(\theta_0, 0, \sigma)$ and note that $\lim_{\sigma \rightarrow 0^+} \Psi(\sigma) < 0$, implying that $\hat{\sigma} > 0$. For any σ such that $\psi(\theta_0, 0, \sigma) > 0$ for all $\theta_0 \in (0, 1)$, the policy-maker can avoid default for any $\theta > 0$ by using the monotone rule $\pi(\theta) = \mathbb{I}\{\theta > 0\}$ that fails all institutions with fundamentals $\theta \leq 0$ and passes the rest. This case is uninteresting. Hereafter, we thus confine attention to $\sigma < \hat{\sigma}$, which guarantees that the policy maker's problem is not trivial.

Let $U_\sigma^\Gamma(x, 1|x)$ denote the payoff differential of an agent with signal x who expects all other agents to pledge if and only if their signal exceeds x , when the precision of private information is σ^{-2} , and the policy Γ announces that $s = 1$. Also let $U_\sigma^\Gamma(x_\sigma^*(0), 1|x_\sigma^*(0)) \equiv \lim_{x \rightarrow -\infty} U_\sigma^\Gamma(x, 1|x)$ and $U_\sigma^\Gamma(x_\sigma^*(1), 1|x_\sigma^*(1)) \equiv \lim_{x \rightarrow +\infty} U_\sigma^\Gamma(x, 1|x)$.

³⁸Given that default occurs if and only if $A \leq 1 - \theta$, $x_\sigma^*(\theta)$ is implicitly defined by the solution to the equation $\Phi\left(\frac{x-\theta}{\sigma}\right) = \theta$. Hence, at θ , the measure of agents pledging (which coincides with the measure of agents receiving signals above $x_\sigma^*(\theta)$) is exactly equal to $1 - \theta$.

³⁹Recall that, when the announcement that $s = 1$ reveals to the market that $\theta \geq 0$, the unique rationalizable profile features all agents pledging, irrespective of their private information, if and only if $\psi(\theta_0, 0, \sigma) > 0$ for all $\theta_0 \in (0, 1)$. This follows directly from Lemma 1.

From the proofs of Theorems 1 and 2, recall that a policy $\Gamma = (\{0, 1\}, \pi)$ satisfies the perfect coordination property only if, after signal $s = 1$ is disclosed, the unique rationalizable profile features all agents pledging, which is the case if and only if $U_\sigma^\Gamma(x, 1|x) > 0$ for all x .

Now let \mathbb{G}_σ denote the set of binary policies $\Gamma = (\{0, 1\}, \pi)$ such that (a) $\pi(\theta) = 0$ for all $\theta \leq 0$, $\pi(\theta) = 1$ for all $\theta > 1$, and (b) for all $x \in \mathbb{R}$, $U_\sigma^\Gamma(x, 1|x) \geq 0$. From the proofs of Theorems 1 and 2, given any σ , and any binary policy Γ' satisfying the perfect coordination property, there exists a binary policy $\Gamma \in \mathbb{G}_\sigma$ that also satisfies the perfect coordination property and such that the probability of default under Γ is weakly smaller than under Γ' . Hence, without loss of generality, hereafter we restrict attention to policies $\Gamma \in \mathbb{G}_\sigma$. However, note that the set \mathbb{G}_σ contains also policies that do not satisfy the perfect coordination property.⁴⁰

Proof Structure. The proof proceeds in four steps, with some of the technical details relegated to the online Supplement. Step 1 establishes that, when σ is small, any policy $\Gamma = (\{0, 1\}, \pi) \in \mathbb{G}_\sigma$ must have the property that any interval $(\theta', \theta'') \subset (0, \theta^{MS}]$ receiving a pass grade (i.e., such that $\pi(\theta) = 1$ for all $\theta \in (\theta', \theta'')$) has a sufficiently small Lebesgue measure, with the measure vanishing when $\sigma \rightarrow 0^+$. If this was not the case, for some $\theta \in (\theta', \theta'')$, $U_\sigma^\Gamma(x_\sigma^*(\theta), 1|x_\sigma^*(\theta)) < 0$, contradicting the fact that $\Gamma \in \mathbb{G}_\sigma$.

Step 2 then considers an *auxiliary game* G_σ in which the agents play less aggressively than in the original game. Namely, G_σ is the game in which (i) the policy maker's choice set is \mathbb{G}_σ and (ii) given *any* policy $\Gamma \in \mathbb{G}_\sigma$, all agents pledge after receiving the signal $s = 1$ and refrain from pledging after receiving the signal $s = 0$, irrespective of their private information. By the definition of \mathbb{G}_σ , the agents' behavior is rationalizable. However, the above action profile is MARP only for those $\Gamma \in \mathbb{G}_\sigma$ for which, for all x , $U_\sigma^\Gamma(x, 1|x) > 0$. For those $\Gamma \in \mathbb{G}_\sigma$ for which there exists x such that $U_\sigma^\Gamma(x, 1|x) = 0$, instead, the above action profile is less aggressive than MARP. We show that, when σ is small, for any given any policy $\Gamma = (\{0, 1\}, \pi) \in \mathbb{G}_\sigma$ that gives a fail grade to an interval $(\theta', \theta'') \subseteq (\underline{\theta}, \theta^{MS}]$ of large Lebesgue measure, there exists another policy $\Gamma^\# \in \mathbb{G}_\sigma$ that gives a pass grade to a F -positive measure subset of (θ', θ'') , has a mesh smaller than Γ , and is such that, when agents play as in G_σ (that is, pledge irrespective of x when hearing that $s = 1$), the probability of default under $\Gamma^\#$ is strictly smaller than under Γ .

Step 3 then combines the results from Steps 1 and 2 to shows that, when σ is small, given any policy $\Gamma \in \mathbb{G}_\sigma$ for which the mesh $M(\Gamma)$ of $(0, \theta^{MS}]$ is larger than ε , there exists another policy $\Gamma' \in \mathbb{G}_\sigma$ with a mesh $M(\Gamma')$ smaller than ε such that, when agents play as in the auxiliary game G_σ , the probability of default is strictly smaller under Γ' than under Γ . Starting from $\Gamma' \in \mathbb{G}_\sigma$ one can then construct a “nearby” policy $\Gamma^* \in \mathbb{G}_\sigma$ such that the probability of default under Γ^* is arbitrarily close to that under Γ' (and hence strictly smaller than under Γ) and such that $U_\sigma^{\Gamma^*}(x, 1|x) > 0$ for

⁴⁰These are policies Γ for which there exists x such that $U_\sigma^\Gamma(x, 1|x) = 0$; when this is the case, in the continuation game that starts after the policy Γ announces $s = 1$, in addition to the rationalizable profile under which all agents pledge irrespective of their signal, there also exists a rationalizable profile where each agent pledges if and only if his private signal exceeds x .

all x . As shown in the proof of Theorem 2, the last property implies that Γ^* satisfies the perfect coordination property: when Γ^* discloses the signal $s = 1$, the unique rationalizable profile features all agents pledging, irrespective of their private signals. The policy Γ^* thus improves upon Γ also in the original game, as claimed in the theorem.

Finally, step 4 closes the proof by showing how to construct the function \mathcal{E} in the theorem relating the noise σ in the agents' exogenous private information to the bound $\mathcal{E}(\sigma)$ on the mesh of the policies.

Step 1. We start with the following result:

Lemma 2. *For any $\varepsilon \in \mathbb{R}_{++}$, there exists $\sigma(\varepsilon) \in \mathbb{R}_{++}$ such that, for any $\sigma \in (0, \sigma(\varepsilon)]$, the following is true: for any policy $\Gamma = (\{0, 1\}, \pi) \in \mathbb{G}_\sigma$ and any $(\theta', \theta'') \in D^\Gamma$ with $|\theta'' - \theta'| > \varepsilon$, necessarily $\pi(\theta) = 0$ for some strictly positive F -measure subset of (θ', θ'') .⁴¹*

Proof of Lemma 2. While the intuition for this result (reported in the main text) is simple, the formal proof is tedious and relegated to the online Supplement. There, we first shows that, for any $\sigma > 0$, any policy $\Gamma = (\{0, 1\}, \pi) \in \mathbb{G}_\sigma$, and any cell $(\theta', \theta'') \in D^\Gamma$ such that $\pi(\theta) = 1$ for all $\theta \in (\theta', \theta'')$, if the policy maker were to replace Γ with a cutoff policy $\Gamma^{\theta'}$ that fails with certainty all types below θ' and passes with certainty all types above θ' , then for any $\theta \leq \theta''$, the payoff differential of the marginal agent with signal $x_\sigma^*(\theta) \equiv \theta + \sigma\Phi^{-1}(\theta)$ would be higher than under the original policy Γ : that is, $U_\sigma^{\Gamma^{\theta'}}(x_\sigma^*(\theta), 1|x_\sigma^*(\theta)) \geq U_\sigma^\Gamma(x_\sigma^*(\theta), 1|x_\sigma^*(\theta))$ for any $\theta \leq \theta''$. Starting from this result, the rest of the proof then shows that for any interval $(\theta', \theta'') \subset (0, \theta^{MS})$ of Lebesgue measure $|\theta'' - \theta'| > \varepsilon$, there exists $\sigma(\varepsilon) \in \mathbb{R}_{++}$ such that for all $\sigma \in (0, \sigma(\varepsilon)]$, $U_\sigma^{\Gamma^{\theta'}}(x_\sigma^*(\theta), 1|x_\sigma^*(\theta)) < 0$ for some $\theta \in (\theta', \theta'')$. Together the two results then imply that, for small σ , $U_\sigma^\Gamma(x_\sigma^*(\theta), 1|x_\sigma^*(\theta)) < 0$ for some $\theta \in (\theta', \theta'')$, thus implying that $\Gamma \notin \mathbb{G}_\sigma$ (recall that, for Γ to be in \mathbb{G}_σ , it must be that $U_\sigma^\Gamma(x, 1|x) \geq 0$ for all x). ■

Step 2. Next, we show that, for any policy $\Gamma = (\{0, 1\}, \pi) \in \mathbb{G}_\sigma$ that gives a fail grade to an interval $(\theta', \theta'') \subseteq (0, \theta^{MS})$ of large Lebesgue measure, there exists another policy $\Gamma^\# \in \mathbb{G}_\sigma$ with a smaller mesh such that, when agents play as in G_σ , the probability of default under $\Gamma^\#$ is strictly smaller than under Γ . We start with the following result:

Lemma 3. *For any $\varepsilon > 0$, there exists $\sigma^\#(\varepsilon) \in (0, \hat{\sigma})$ such that, for any $\sigma \in (0, \sigma^\#(\varepsilon)]$, and any policy $\Gamma = (\{0, 1\}, \pi) \in \mathbb{G}_\sigma$ for which there exists $(\theta', \theta'') \in D^\Gamma$ such that (a) $|\theta'' - \theta'| > \varepsilon$ and (b) $\pi(\theta) = 0$ for all $\theta \in (\theta', \theta'')$, there exists another policy $\Gamma^\# = (\{0, 1\}, \pi^\#) \in \mathbb{G}_\sigma$, with $M(\Gamma^\#) < M(\Gamma)$, such that, in the auxiliary game G_σ , the probability of default under $\Gamma^\#$ is strictly smaller than under Γ .*

Proof of Lemma 3. For any $\theta \in (0, \theta^{MS})$, $\lim_{\sigma \rightarrow 0^+} x_\sigma^*(\theta) \equiv x_{0^+}^*(\theta) = \theta$. Furthermore, for any $\varepsilon \in (0, \theta^{MS})$, the function $x_{0^+}^* : [\frac{\varepsilon}{4}, \theta^{MS}] \rightarrow \mathbb{R}$ is uniformly continuous. Hence, for any $\gamma < \varepsilon/4$,

⁴¹Recall that $D^\Gamma \equiv \{d_i = (\underline{\theta}_i, \bar{\theta}_i] : i = 1, \dots, N\}$ is the partition of $[\underline{\theta}, \theta^{MS}]$ induced by the policy Γ .

there exists $\tilde{\sigma}(\gamma) > 0$ such that, for any $\sigma \in (0, \tilde{\sigma}(\gamma)]$, any $\theta \in [\frac{\varepsilon}{4}, \theta^{MS}]$, $|x_\sigma^*(\theta) - \theta| \leq \gamma$. In turn, this implies that, for any $\varepsilon > 0$, there exists $\sigma^\#(\varepsilon) \in (0, \hat{\sigma}]$ such that, for any $\sigma \in (0, \sigma^\#(\varepsilon)]$, any $(\theta', \theta'') \in D^\Gamma$ such that $|\theta'' - \theta'| > \varepsilon$, for any $\theta \geq \theta''$, $|\theta - x_\sigma^*(\theta)| < |(\theta' + \theta'')/2 - x_\sigma^*(\theta)|$. Likewise, for any $\theta \leq \theta'$, and any $\hat{\theta} \geq \theta''$, $|\theta - x_\sigma^*(\theta)| < |x_\sigma^*(\theta) - \hat{\theta}|$ provided that $\sigma \in (0, \sigma^\#(\varepsilon)]$.

Next, pick any policy $\Gamma = (\{0, 1\}, \pi) \in \mathbb{G}_\sigma$ for which there exists $d \equiv (\theta', \theta'') \in D^\Gamma$ such that (a) $|\theta'' - \theta'| > \varepsilon$ and (b) $\pi(\theta) = 0$ for all $\theta \in (\theta', \theta'')$. If $\min_{\theta \in [0, 1]} U_\sigma^\Gamma(x_\sigma^*(\theta), 1|x_\sigma^*(\theta)) > 0$, the result in the lemma follows directly from the possibility to perturb the original policy Γ by increasing the set of fundamentals over which the policy maker assigns a pass grade, in a way that guarantees that the perturbed policy is sufficiently close to the original policy (in the L_1 norm) and hence belongs in \mathbb{G}_σ (the formal proof is in the online Supplement).

Thus assume that $\min_{\theta \in [0, 1]} U_\sigma^\Gamma(x_\sigma^*(\theta), 1|x_\sigma^*(\theta)) = 0$. Suppose first that $\min_{\theta \in [\theta'', 1]} U_\sigma^\Gamma(x_\sigma^*(\theta), 1|x_\sigma^*(\theta)) > 0$. In the online Supplement, we then show that necessarily $U_\sigma^\Gamma(x_\sigma^*(\theta), 1|x_\sigma^*(\theta)) > 0$ for all $\theta \in (\theta', \theta'')$. This follows from the fact that $U_\sigma^\Gamma(x_\sigma^*(\theta), 1|x_\sigma^*(\theta)) = 0$ for some $\theta \in (\theta', \theta'')$ is only consistent with $\pi(\theta) = 1$ for all $\theta \in (\theta', \theta'')$, contradicting the assumptions in the lemma (intuitively, when σ is small, $x_\sigma^*(\theta)$ is close to θ ; hence, for the marginal agent with signal $x_\sigma^*(\theta)$ to be indifferent between pledging and not pledging when hearing that $r = 1$, it must be that $r = 1$ is consistent with fundamentals close to θ).

Hence, $\min_{\theta \in (\theta', 1]} U_\sigma^\Gamma(x_\sigma^*(\theta), 1|x_\sigma^*(\theta)) > 0$. We show that, starting from Γ , we can then construct a policy Γ^η that continues to satisfy the property that $U_\sigma^{\Gamma^\eta}(x_\sigma^*(\theta), 1|x_\sigma^*(\theta)) \geq 0$ for all $\theta \in [0, 1]$ and such that the probability of default under Γ^η is smaller than under Γ . The policy Γ^η is obtained from Γ by giving a pass grade to a positive-measure interval of types in the middle of (θ', θ'') . Formally, take $\eta \in (0, (\theta'' - \theta')/2)$ and let $\Gamma^\eta = (\{0, 1\}, \pi^\eta)$ be the policy whose rule π^η is given by (a) $\pi^\eta(\theta) = \pi(\theta)$ for all $\theta \notin [(\theta' + \theta'')/2, (\theta' + \theta'')/2 + \eta]$, and (b) $\pi^\eta(\theta) = 1$ for all $\theta \in [(\theta' + \theta'')/2, (\theta' + \theta'')/2 + \eta]$. In the online Supplement, we show that $U_\sigma^{\Gamma^\eta}(x_\sigma^*(\theta), 1|x_\sigma^*(\theta)) \geq 0$ for all $\theta \in [0, 1]$.⁴² Hence $\Gamma^\eta \in \mathbb{G}_\sigma$. That, when agents play according to G_σ , the probability of default is strictly smaller under Γ^η than under Γ follows directly from the fact that the set of fundamentals that receive a pass grade under Γ^η is a strict superset of the set of fundamentals that receive a pass grade under Γ .

Next, consider the more interesting case in which $\min_{\theta \in [\theta'', 1]} U_\sigma^\Gamma(x_\sigma^*(\theta), 1|x_\sigma^*(\theta)) = 0$. For any $\sigma \in (0, \hat{\sigma})$, and any policy $\Gamma \in \mathbb{G}_\sigma$, let $\Theta_\sigma^\Gamma \equiv \arg \min_{\theta \in [0, 1]} U_\sigma^\Gamma(x_\sigma^*(\theta), 0|x_\sigma^*(\theta))$ and then let $\theta_\sigma^\# \equiv \inf \{\theta \in \Theta_\sigma^\Gamma : \theta \geq \theta''\}$. In the online Supplement, we show that necessarily $\theta_\sigma^\# > \theta''$ (the arguments are the same as those establishing that necessarily $U_\sigma^\Gamma(x_\sigma^*(\theta), 1|x_\sigma^*(\theta)) > 0$ for all (θ', θ'')). Also

⁴²Intuitively, because, under the original policy Γ , $U_\sigma^\Gamma(x_\sigma^*(\theta), 1|x_\sigma^*(\theta)) > 0$ for all $\theta > \theta'$, as long as η is small, under the new policy Γ^η , $U_\sigma^{\Gamma^\eta}(x_\sigma^*(\theta), 1|x_\sigma^*(\theta)) \geq 0$ for all $\theta > \theta'$. The difficult step is to show that $U_\sigma^{\Gamma^\eta}(x_\sigma^*(\theta), 1|x_\sigma^*(\theta)) > 0$ also for all $\theta \leq \theta'$. Because the new policy Γ^η assigns a pass grade to more fundamentals above θ' than the original policy Γ , and because default does not occur over any fundamental above θ' when agents follow a cut-off strategy with cutoff $x_\sigma^*(\theta)$, $\theta \leq \theta'$, the expected payoff from pledging of any marginal with signal $x_\sigma^*(\theta)$ is higher under Γ^η than under Γ , for any $\theta \leq \theta'$.

observe that, by definition of the partition D^Γ , $\pi(\theta) = 1$ in a right neighborhood of θ'' . Then let $d^\Gamma(\theta'') = (\theta'', \theta''']$ denote the interval to the immediate right $(\theta', \theta'']$ in the partition induced by the policy Γ and let $\hat{\theta} = \min \{ \theta''', \theta_\sigma^\# \}$.

Now, pick $\xi > 0$ small and let $\delta(\xi)$ be implicitly defined by

$$F((\theta' + \theta'')/2 + \xi) - F((\theta' + \theta'')/2) = F((\theta'' + \hat{\theta})/2 + \delta(\xi)) - F((\theta'' + \hat{\theta})/2). \quad (15)$$

Consider the policy $\Gamma^\xi = (\{0, 1\}, \pi^\xi)$ defined by (a) $\pi^\xi(\theta) = \pi(\theta)$ for all $\theta \notin [(\theta' + \theta'')/2, (\theta' + \theta'')/2 + \xi] \cup [(\theta'' + \hat{\theta})/2, (\theta'' + \hat{\theta})/2 + \delta(\xi)]$, (b) $\pi^\xi(\theta) = 1$ for all $\theta \in [(\theta' + \theta'')/2, (\theta' + \theta'')/2 + \xi]$, and (c) $\pi^\xi(\theta) = 0$ for all $\theta \in [(\theta'' + \hat{\theta})/2, (\theta'' + \hat{\theta})/2 + \delta(\xi)]$. In the online Supplement, we establish that, when $\xi > 0$ is small, such a policy is such that $\min_{\theta \in [0, 1]} U^{\Gamma^\xi}(x_\sigma^*(\theta), 1|x_\sigma^*(\theta)) > 0$ and hence $\Gamma^\xi \in \mathbb{G}_\sigma$.⁴³

By construction, $M(\Gamma^\xi) < M(\Gamma)$. Furthermore, when agents play according to G_σ , the probability of default under Γ^ξ is the same as under Γ . That $\min_{\theta \in [0, 1]} U^{\Gamma^\xi}(x_\sigma^*(\theta), 1|x_\sigma^*(\theta)) > 0$, however, implies that, starting from Γ^ξ , one can construct another policy $\Gamma^\# \in \mathbb{G}_\sigma$, sufficiently close to Γ^ξ in the L_1 norm, such that (1) $M(\Gamma^\#) < M(\Gamma)$ and (2), when the agents play according to G_σ , the probability of default under $\Gamma^\#$ is strictly smaller than under Γ . ■

Step 3. Steps 1 and 2 can be iterated to establish the following conclusions. Let $\sigma(\varepsilon)$ and $\sigma^\#(\varepsilon)$ be the thresholds from Lemmas 2 and 3, respectively. There exists a function $\bar{\sigma} : (0, \theta^{MS}) \rightarrow \mathbb{R}_{++}$, with $\bar{\sigma}(\varepsilon) \leq \min\{\sigma(\varepsilon), \sigma^\#(\varepsilon)\}$ for all $\varepsilon \in (0, \theta^{MS})$ and with $\bar{\sigma}(\varepsilon) \rightarrow 0^+$ as $\varepsilon \rightarrow 0^+$, such that the following is true: For any $\varepsilon \in (0, \theta^{MS})$, any $\sigma \in (0, \bar{\sigma}(\varepsilon)]$, and any policy $\Gamma = (\{0, 1\}, \pi) \in \mathbb{G}_\sigma$ with $M(\Gamma) > \varepsilon$, there exists another policy $\Gamma' = (\{0, 1\}, \pi') \in \mathbb{G}_\sigma$ with $M(\Gamma') \leq \varepsilon$ such that, when the agents play as in the auxiliary game G_σ , the probability of default under Γ' is strictly smaller than under Γ .⁴⁴

Starting from Γ' , one can then construct a nearby policy $\Gamma^* = (\{0, 1\}, \pi^*) \in \mathbb{G}_\sigma$ such that the probability of default under Γ^* is arbitrarily close to that under Γ' (and hence strictly smaller than under Γ) and such that $U_\sigma^{\Gamma^*}(x, 1|x) > 0$ for all x . The last property implies that Γ^* satisfies the perfect-coordination property.

Step 4. The proof is completed by showing how to construct the function \mathcal{E} in the theorem

⁴³Heuristically, the result follows from the following observations. (a) For any $\theta < \theta'$, $U_\sigma^{\Gamma^\xi}(x_\sigma^*(\theta), 1|x_\sigma^*(\theta)) > U_\sigma^\Gamma(x_\sigma^*(\theta), 1|x_\sigma^*(\theta))$. This is because, when σ is small and $\theta < \theta'$, $x_\sigma^*(\theta) < (\theta' + \theta'')/2$. Because the new policy Γ^ξ assigns a pass grade to fundamentals above θ closer to $x_\sigma^*(\theta)$ than the original policy Γ , and because default does not occur over such fundamentals when agents follow a cut-off strategy with cutoff $x_\sigma^*(\theta)$, the expected payoff from pledging at $x_\sigma^*(\theta)$ is higher under Γ^ξ than under Γ . (b) Because $U_\sigma^\Gamma(x_\sigma^*(\theta), 1|x_\sigma^*(\theta)) > 0$ for all $\theta \in (\theta', \hat{\theta})$, as long as ξ is small, $U_\sigma^{\Gamma^\xi}(x_\sigma^*(\theta), 1|x_\sigma^*(\theta)) \geq 0$ for all $\theta \in (\theta', \hat{\theta})$. (c) For any $\theta > \hat{\theta}$, when σ and ξ are small, $x_\sigma^*(\theta) > (\theta'' + \hat{\theta})/2 + \delta(\xi)$. Because the new policy Γ^ξ assigns a pass grade to fundamentals below θ farther away from $x_\sigma^*(\theta)$ than the original policy Γ , and because default occurs over such fundamentals when agents follow a cut-off strategy with cutoff $x_\sigma^*(\theta)$, the expected payoff from pledging at $x_\sigma^*(\theta)$ is strictly higher under Γ^ξ than under Γ .

⁴⁴As we formally establish in the online Supplement, the thresholds $\sigma(\varepsilon)$ and $\sigma^\#(\varepsilon)$ identified in Steps 1 and 2 above are invariant to the initial policy Γ . The same steps used above to identify the policy $\Gamma^\#$ with mesh $M(\Gamma^\#) < M(\Gamma)$ can then be iterated till one arrives at a policy Γ' with mesh $M(\Gamma') \leq \varepsilon$.

relating the noise σ in the agents' exogenous private information to the bound $\mathcal{E}(\sigma)$ on the mesh of the policies. Let (ε_n) be a non-increasing sequence satisfying $\lim_{n \rightarrow \infty} \varepsilon_n = 0$. For each $n \in \mathbb{N}$, then let $\sigma_n = \bar{\sigma}(\varepsilon_n)$, with the function $\bar{\sigma}(\cdot)$ as defined in Step 3. The results in Steps 1-3 above imply that, given $(\varepsilon_n, \sigma_n)$, there exist strictly decreasing subsequences $(\tilde{\varepsilon}_n)$ and $(\tilde{\sigma}_n)$ satisfying $\lim_{n \rightarrow \infty} \tilde{\varepsilon}_n = \lim_{n \rightarrow \infty} \tilde{\sigma}_n = 0$ such that, for any $n \in \mathbb{N}$, the conclusions in Step 3 hold for $\varepsilon = \tilde{\varepsilon}_n$ and $\bar{\sigma}(\varepsilon_n) = \tilde{\sigma}_n$. Then let $\bar{\sigma} = \tilde{\sigma}_0 > 0$ and $\mathcal{E} : (0, \bar{\sigma}] \rightarrow \mathbb{R}_+$ be the function defined by $\mathcal{E}(\sigma) = \varepsilon_n$ for all $\sigma \in (\sigma_{n+1}, \sigma_n]$. The result in the theorem then follows from Steps 1-3, by letting $\mathcal{E}(\cdot)$ be the function constructed above.

This completes the proof of the theorem. Q.E.D.

Proof of Theorem 4. The proof is omitted as it follows from adapting the steps used to establish Theorems 1 and 2 in the baseline model, using the arguments indicated in the main text (see the discussion after Theorem 4).

Proof of Theorem 5. Without loss of generality, assume the policy $\Gamma = (\mathcal{S}, \pi)$ (a) is a (possibly stochastic) “pass/fail” policy (i.e., $\mathcal{S} = \{0, 1\}$, with $\pi(1|\theta)$ denoting the probability that signal $s = 1$ is disclosed when the fundamentals are θ), (b) is such that $\pi(1|\theta) = 0$ for all $\theta \leq \underline{\theta}$, and $\pi(1|\theta) = 1$ for all $\theta > \bar{\theta}$, and (c) satisfies the perfect coordination property. Note that Theorem 4 implies that, if Γ does not satisfy these properties, there exists another policy Γ' that satisfies 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 .

Suppose that Γ is such that there exists no θ^* such that $\pi(1|\theta) = 0$ for F -almost all $\theta \leq \theta^*$ and $\pi(1|\theta) = 1$ for F -almost all $\theta > \theta^*$.⁴⁵ We establish the result by showing that there exists a deterministic, pass/fail, monotone policy Γ^* satisfying the perfect coordination property that yields the policy maker a payoff strictly higher than Γ .

Recall that, for the policy Γ to satisfy the perfect coordination property, it must be that, when the policy discloses the signal $s = 1$, $U^\Gamma(x, 1|x) > 0$ for all x , where $U^\Gamma(x, 1|x)$ is the expected payoff of an agent with signal x who hears that $s = 1$ and who expects all other agents to follow a cut-off policy with cut-off x (in case the payoffs and the default outcome depend also on variables z only imperfectly correlated with θ , the expectation is over both θ and z).

Now let \mathbb{G}_σ denote the set of policies $\Gamma' = (\mathcal{S}, \pi')$ that, in addition to properties (a) and (b) above, are such that $U^{\Gamma'}(x, 1|x) \geq 0$, all x . Recall from the discussion in the proof of Theorem 3 that some policies Γ' in \mathbb{G}_σ need not satisfy the perfect coordination property (namely, those for which there exists x such that $U^{\Gamma'}(x, 1|x) = 0$). Let $\arg \max_{\Gamma \in \mathbb{G}_\sigma} \{U^P[\Gamma]\}$ be the set of policies that maximize

⁴⁵Clearly, if the policy $\Gamma = (\{0, 1\}, \pi)$ is such that there exists $\theta^* \in [\underline{\theta}, \bar{\theta}]$ such that $\pi(1|\theta) = 0$ for F -almost all $\theta \leq \theta^*$ and $\pi(1|\theta) = 1$ for F -almost all $\theta \geq \theta^*$, then the deterministic, binary, monotone policy $\Gamma^* = (\{0, 1\}, \pi^*)$ such that $\pi^*(1|\theta) = 0$ for all $\theta \leq \theta^*$ and $\pi^*(1|\theta) = 1$ for all $\theta > \theta^*$ also satisfies the perfect coordination property and yields the policy maker the same payoff as Γ , in which case the result trivially holds.

the policy maker's payoff over the set \mathbb{G}_σ .⁴⁶

Step 1 below shows that any $\tilde{\Gamma} \in \arg \max_{\Gamma \in \mathbb{G}_\sigma} \{U^P[\Gamma]\}$ is such that there exists θ^* such that $\tilde{\pi}(1|\theta) = 0$ for F -almost all $\theta \leq \theta^*$ and $\tilde{\pi}(1|\theta) = 1$ for F -almost all $\theta > \theta^*$. We establish this result by showing that, for any policy $\Gamma' \in \mathbb{G}_\sigma$ for which there exists no θ^* such that $\pi'(1|\theta) = 0$ for F -almost all $\theta \leq \theta^*$ and $\pi'(1|\theta) = 1$ for F -almost all $\theta > \theta^*$, there exists another policy $\Gamma'' \in \mathbb{G}_\sigma$ that strictly improves over Γ' . Step 2 then shows that the policy maker's payoff under any $\tilde{\Gamma}$ satisfying the above properties can be approximated arbitrarily well by a deterministic threshold policy $\Gamma^* \in \mathbb{G}_\sigma$ that satisfies the perfect coordination property (i.e., such that $U^{\Gamma^*}(x, 1|x) > 0$, all x).

Step 1. Take any policy $\Gamma' \in \mathbb{G}_\sigma$ for which there exists no θ^* such that $\pi'(1|\theta) = 0$ for F -almost all $\theta \leq \theta^*$ and $\pi'(1|\theta) = 1$ for F -almost all $\theta > \theta^*$. Let $X^{\Gamma'} \equiv \{x : U^{\Gamma'}(x, 1|x) = 0\}$. Clearly, if $X^{\Gamma'} = \emptyset$, there exists another policy $\Gamma'' \in \mathbb{G}_\sigma$ that yields the policy maker a payoff strictly higher than Γ' .⁴⁷ Hence, assume $X^{\Gamma'} \neq \emptyset$, and let $\bar{x} \equiv \sup X^{\Gamma'}$,

$$\begin{aligned}\theta_0 &\equiv \inf\{\theta : \exists \delta > 0 \text{ s.t. } \pi'(1|\theta') > 0 \text{ for } F\text{-almost all } \theta' \in [\theta, \theta + \delta)\}, \\ \theta_H &\equiv \sup\{\theta : \exists \delta > 0 \text{ s.t. } \pi'(1|\theta') < 1 \text{ for } F\text{-almost all } \theta' \in [\theta - \delta, \theta)\}.\end{aligned}$$

That, under Γ' , there exists no θ^* such that $\pi'(1|\theta) = 0$ for F -almost all $\theta \leq \theta^*$ and $\pi'(1|\theta) = 1$ for F -almost all $\theta > \theta^*$ implies that $\theta_0 < \theta_H$. Furthermore, $u(\theta_0, 1 - P(\bar{x}|\theta_0)) < 0$.⁴⁸ Now suppose that $u(\theta, 1 - P(\bar{x}|\theta)) \geq 0$ for all $\theta \in \Theta^+ \equiv \{\theta \geq \theta_0 : \exists \delta > 0 \text{ s.t. } \pi'(1|\theta') < 1 \text{ for } F\text{-almost all } \theta' \in [\theta - \delta, \theta)\}$. Then consider the policy $\Gamma'' = (\{0, 1\}, \pi'')$ defined by $\pi''(1|\theta) = \pi'(1|\theta)$ for all $\theta \notin \Theta^+$, and $\pi''(1|\theta) = \min\{\pi'(1|\theta) + \epsilon, 1\}$ for all $\theta \in \Theta^+$, for some $\epsilon > 0$ small. The above property, along with the monotonicity of $u(\theta, 1 - P(x|\theta))$ in x , implies that, for any $x \in X^{\Gamma'}$,

$$\int_{\theta_0}^{+\infty} u(\theta, 1 - P(x|\theta))\pi''(1|\theta)p(x|\theta)dF(\theta) > \int_{\theta_0}^{+\infty} u(\theta, 1 - P(x|\theta))\pi'(1|\theta)p(x|\theta)dF(\theta) = 0,$$

which in turn implies that $U^{\Gamma''}(x, 1|x) > 0$ for all $x \in X^{\Gamma'}$. Furthermore, because $U^{\Gamma'}(x, 1|x) > 0$ for all $x \notin X^{\Gamma'}$, there exists $\epsilon > 0$ small enough such that $U^{\Gamma''}(x, 1|x) > 0$ for all x . Hence $\Gamma'' \in \mathbb{G}_\sigma$. Furthermore, because $U^P(\theta, 1) > U^P(\theta, 0)$ for all $\theta \in \Theta^+$, the policy maker's payoff under Γ'' is strictly higher than under Γ' .

Next, assume Γ' is such that $u(\theta, 1 - P(\bar{x}|\theta)) < 0$ for some $\theta_1 \in \Theta^+$. Then fix $\epsilon > 0$ small and let $\delta(\epsilon)$ be implicitly defined by

$$\int_{\theta_0}^{\theta_0 + \epsilon} u(\theta, 1 - P(\bar{x}|\theta))\pi'(1|\theta)p(\bar{x}|\theta)dF(\theta) = \int_{\theta_1 - \delta(\epsilon)}^{\theta_1} u(\theta, 1 - P(\bar{x}|\theta))(1 - \pi'(1|\theta))p(\bar{x}|\theta)dF(\theta). \quad (16)$$

⁴⁶The arguments below also imply that $\arg \max_{\Gamma \in \mathbb{G}_\sigma} \{U^P[\Gamma]\} \neq \emptyset$.

⁴⁷To see this, note that, because there exists no θ^* such that $\pi'(1|\theta) = 0$ for F -almost all $\theta \leq \theta^*$ and $\pi'(1|\theta) = 1$ for F -almost all $\theta > \theta^*$, if $X^{\Gamma'} = \emptyset$, there must exist a set $(\theta', \theta'') \subseteq [\underline{\theta}, \bar{\theta}]$ of F -positive probability over which $\pi'(1|\theta) < 1$. The policy Γ'' can then be obtained from Γ' by increasing $\pi'(1|\theta)$ over such a set. Provided the increase is small, the new policy is such that $U^{\Gamma''}(x, 1|x) \geq 0$ for all x , and hence $\Gamma'' \in \mathbb{G}_\sigma$. Because $U^P(\theta, 1) > U^P(\theta, 0)$ over $[\underline{\theta}, \bar{\theta}]$, the new policy improves over the original one.

⁴⁸That $u(\theta_0, 1 - P(\bar{x}|\theta_0)) < 0$ follows from the fact that $\int_{\theta_0}^{+\infty} u(\theta, 1 - P(\bar{x}|\theta))\pi'(1|\theta)p(\bar{x}|\theta)dF(\theta) = 0$ and the monotonicity of $u(\theta, 1 - P(\bar{x}|\theta))$ in θ .

Consider the policy $\Gamma^\epsilon = (\{0, 1\}, \pi^\epsilon)$ defined by the following properties: (a) $\pi^\epsilon(1|\theta) = \pi'(1|\theta)$ 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) $\pi^\epsilon(1|\theta) = 0$ for all $\theta \in [\theta_0, \theta_0 + \epsilon]$; and (c) $\pi^\epsilon(1|\theta) = 1$ for all $\theta \in [\theta_1 - \delta(\epsilon), \theta_1]$. Note that, for ϵ strictly positive but small, such a policy is well defined. Also note that Condition (16) implies that $U^{\Gamma^\epsilon}(\bar{x}, 1|\bar{x}) = U^{\Gamma'}(\bar{x}, 1|\bar{x})$. We now establish that, for ϵ sufficiently small, under the new policy Γ^ϵ , $U^{\Gamma^\epsilon}(x, 1|x) > 0$ for any $x \in X^{\Gamma'}$, with $x \neq \bar{x}$. A necessary and sufficient condition for this to be the case is that, for any such x ,

$$\lim_{\epsilon \rightarrow 0^+} \frac{\partial}{\partial \epsilon} \{U^{\Gamma^\epsilon}(x, 1|x)p^{\Gamma^\epsilon}(1, x)\} > 0, \quad (17)$$

where $p^{\Gamma^\epsilon}(\bar{x}, 1) \equiv \int_{-\infty}^{+\infty} \pi^\epsilon(1|\theta)p(\bar{x}|\theta)dF(\theta)$. To see that (17) holds, use Condition (16) to observe that

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

Now use the fact that

$$\begin{aligned} U^{\Gamma^\epsilon}(x, 1|x)p^{\Gamma^\epsilon}(1, x) &= \int_{\theta_0+\epsilon}^{\theta_1-\delta(\epsilon)} u(\theta, 1 - P(x|\theta))\pi'(1|\theta)p(x|\theta)dF(\theta) \\ &\quad + \int_{\theta_1-\delta(\epsilon)}^{\theta_1} u(\theta, 1 - P(x|\theta))p(x|\theta)dF(\theta) \\ &\quad + \int_{\theta_1}^{+\infty} u(\theta, 1 - P(x|\theta))\pi'(1|\theta)p(x|\theta)dF(\theta) \end{aligned} \quad (18)$$

to observe that, for any $x \in X^{\Gamma'}, x \neq \bar{x}$,

$$\begin{aligned} \lim_{\epsilon \rightarrow 0^+} \frac{\partial}{\partial \epsilon} \{U^{\Gamma^\epsilon}(x, 1|x)p^{\Gamma^\epsilon}(1, x)\} &= -u(\theta_0, 1 - P(x|\theta_0))\pi'(1|\theta_0)p(x|\theta_0)f(\theta_0) + \\ &\quad + u(\theta_1, 1 - P(x|\theta_1))(1 - \pi'(1|\theta_1))p(x|\theta_1)f(\theta_1) \lim_{\epsilon \rightarrow 0^+} \delta'(\epsilon) \\ &= \pi'(1|\theta_0)f(\theta_0)p(\bar{x}|\theta_0)u(\theta_0, 1 - P(\bar{x}|\theta_0)) \left(\frac{p(x|\theta_1)u(\theta_1, 1 - P(x|\theta_1))}{p(\bar{x}|\theta_1)u(\theta_1, 1 - P(\bar{x}|\theta_1))} - \frac{p(x|\theta_0)u(\theta_0, 1 - P(x|\theta_0))}{p(\bar{x}|\theta_0)u(\theta_0, 1 - P(\bar{x}|\theta_0))} \right) > 0 \end{aligned} \quad (19)$$

where the inequality follows from the log-supermodularity of $p(x|\theta)$ and of $|u(\theta, 1 - P(x|\theta))|$, as per property (1) in Condition M, along with the fact that $u(\theta_0, 1 - P(x|\theta_0)) < 0$ for all $x \in X^{\Gamma'}$.⁴⁹

We conclude that, when ϵ is positive but small, under the new policy Γ^ϵ , for all $x \in X^{\Gamma'}, x \neq \bar{x}$, $U^{\Gamma^\epsilon}(x, 1|x) > 0$. That, under the same policy, $U^{\Gamma^\epsilon}(x, 1|x) \geq 0$ also for $x \notin X^{\Gamma'}$ follows from the fact that $U^{\Gamma^\epsilon}(x, 1|x)$ is continuous in (x, ϵ) . Together, the results above thus imply that the new policy $\Gamma^\epsilon \in \mathbb{G}_\sigma$.

We now show that, when property (2) in Condition M holds, the new policy yields the policy maker an expected payoff strictly higher than Γ' . To see this, observe that, for any $\epsilon \geq 0$, the policy maker's payoff under the policy Γ^ϵ is equal to

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

⁴⁹Clearly, because $u(\theta_0, 1 - P(x|\theta_0)), u(\theta_0, 1 - P(\bar{x}|\theta_0)), u(\theta_1, 1 - P(\bar{x}|\theta_1)) < 0$, if $u(\theta_1, 1 - P(x|\theta_1)) > 0$, the inequality in (19) holds irrespective of the assumed modularities.

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

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

Therefore, $\lim_{\epsilon \rightarrow 0^+} \frac{dU^P[\Gamma^\epsilon]}{d\epsilon} > 0$ if and only if

$$\frac{U^P(\theta_1, 1) - U^P(\theta_1, 0)}{U^P(\theta_0, 1) - U^P(\theta_0, 0)} > \frac{p(\bar{x}|\theta_1)u(\theta_1, 1 - P(\bar{x}|\theta_1))}{p(\bar{x}|\theta_0)u(\theta_0, 1 - P(\bar{x}|\theta_0))}.$$

Property (2) in Condition M guarantees this is the case.

We conclude that the policy $\Gamma^\epsilon \in \mathbb{G}_\sigma$ strictly improves upon Γ' . Furthermore, the construction of Γ^ϵ above also implies that (a) there exists a deterministic binary monotone policy $\Gamma'' \in \mathbb{G}_\sigma$ that strictly improves over Γ' and (b) $\arg \max_{\Gamma \in \mathbb{G}_\sigma} \{U^P[\Gamma]\} \neq \emptyset$. Hence any policy $\tilde{\Gamma} \in \arg \max_{\Gamma \in \mathbb{G}_\sigma} \{U^P[\Gamma]\}$ must be characterized by a threshold θ^* such that $\tilde{\pi}(1|\theta) = 0$ for F -almost all $\theta \leq \theta^*$ and $\tilde{\pi}(1|\theta) = 1$ for F -almost all $\theta > \theta^*$.

Step 2. Let θ^* be the cut-off corresponding to the policy $\tilde{\Gamma} \in \arg \max_{\Gamma \in \mathbb{G}_\sigma} \{U^P[\Gamma]\}$ and note that the latter must satisfy (4). The result in the theorem then follows from observing that, given $\tilde{\Gamma}$, there exists a nearby deterministic monotone policy $\Gamma^* \in \mathbb{G}_\sigma$ with cut-off $\theta^{**} = \theta^* + \varepsilon$, for $\varepsilon > 0$ but small, such that Γ^* satisfies the perfect coordination property (i.e., $U^{\Gamma^*}(x, 1|x) > 0$ all x) and yields the policy maker a payoff arbitrarily close to that under $\tilde{\Gamma}$ (and hence strictly higher than that under Γ'). Q.E.D.

Proof of Proposition 1. Given any threshold $\hat{\theta}$, and any signal x , let

$$\psi_h(x, \hat{\theta}, \sigma) = \int_{\Theta} u_h \left(\theta, 1 - \Phi \left(\frac{x - \theta}{\sigma} \right) \right) d\Lambda(\theta|x, 1),$$

denote the payoff of an agent with signal x , of precision σ^{-2} , who, after hearing that the bank passed the test, learns that $\theta > \hat{\theta}$, and who expects all other agents to buy security h when their signal exceeds x and short-sell it otherwise, with $h = D$ in case the bank finances itself with debt, and $h = E$ in the case the bank finances itself with equity (here $\Lambda(\cdot|x, 1)$ represents the posterior belief over θ for an agent with exogenous signal x who learns that the bank passed the test). We start with the following result:

Lemma 4. *Assume that (a) z is independent of θ , (b) $\rho_S = 1$, (c) the noise in the agents' signals is Gaussian (that is, $x_i = \theta + \sigma\epsilon_i$, with ϵ_i drawn from a standard Normal distribution, independently of (θ, z)), and (d) $l(\theta) = l$ for all θ , with $l \in \mathbb{R}_+$. Then, for any $\hat{\theta} \in (\underline{\theta}, \bar{\theta})$, any $x > \hat{\theta}$,*

$$\lim_{\sigma \rightarrow 0^+} \frac{\partial}{\partial \sigma} \psi_E(x, \hat{\theta}, \sigma) > 0 > \lim_{\sigma \rightarrow 0^+} \frac{\partial}{\partial \sigma} \psi_D(x, \hat{\theta}, \sigma).$$

Proof of Lemma 4. Note that, for any $x \geq \hat{\theta}$,

$$\begin{aligned} \frac{\partial}{\partial \sigma} \psi_h(x, \hat{\theta}, \sigma) &= \frac{\partial}{\partial \sigma} \int_{\hat{\theta}}^{\infty} u_h(\theta, 1 - \Phi\left(\frac{x-\theta}{\sigma}\right)) \frac{\phi\left(\frac{x-\theta}{\sigma}\right)}{\sigma \Phi\left(\frac{x-\hat{\theta}}{\sigma}\right)} d\theta = \frac{\partial}{\partial \sigma} \frac{\int_{1-\Phi\left(\frac{x-\hat{\theta}}{\sigma}\right)}^1 u_h(x-\sigma\Phi^{-1}(1-A), A) dA}{\Phi\left(\frac{x-\hat{\theta}}{\sigma}\right)} \\ &= \frac{\int_{1-\Phi\left(\frac{x-\hat{\theta}}{\sigma}\right)}^1 \frac{\partial u_h(x-\sigma\Phi^{-1}(1-A), A)}{\partial \theta} (-\Phi^{-1}(1-A)) dA}{\Phi\left(\frac{x-\hat{\theta}}{\sigma}\right)} + \frac{[\psi(x, \hat{\theta}, \sigma) - u_h(\hat{\theta}, 1 - \Phi\left(\frac{x-\hat{\theta}}{\sigma}\right))] \phi\left(\frac{x-\hat{\theta}}{\sigma}\right) \left(\frac{x-\hat{\theta}}{\sigma^2}\right)}{\Phi\left(\frac{x-\hat{\theta}}{\sigma}\right)} \end{aligned} \quad (20)$$

The second equality follows from the change in variables, $A = 1 - \Phi((x - \theta) / \sigma)$, whereas the second equality follows from the chain rule of differentiation.

The proof proceeds in two steps. Step 1 shows that, when $\sigma \rightarrow 0^+$, and $x > \hat{\theta}$, the second term in the right-hand side of the last equality in (20) vanishes. Step 2 shows that, when $\sigma \rightarrow 0^+$, and $x > \hat{\theta}$, the first term in the right-hand side of the last equality in (20) is positive for equity but negative for debt.

Step 1. Because $g_h(\cdot)$ and $b_h(\cdot)$ are bounded, for any σ , $\psi_h(x, \hat{\theta}, \sigma) - u_h(\hat{\theta}, 1 - \Phi\left(\frac{x-\hat{\theta}}{\sigma}\right))$ is also bounded. Next, by L'Hopital's rule, we have that, for any $x > \hat{\theta}$, $\lim_{\sigma \rightarrow 0^+} \phi\left(\frac{x-\hat{\theta}}{\sigma}\right) \left(\frac{x-\hat{\theta}}{\sigma^2}\right) = 0$, implying that the second term in the right-hand side of the last equality in (20) vanishes as $\sigma \rightarrow 0^+$.

Step 2. For any $x > \hat{\theta}$, when $\sigma \rightarrow 0^+$, first term in the right-hand side of the last equality in (20) converges to $\int_0^1 \frac{\partial u_h(x, A)}{\partial \theta} (-\Phi^{-1}(1 - A)) dA$. Next, observe that, because z and θ are independent,

$$u_h(\theta, A) = \int_{\underline{z}}^{\hat{z}_h(A)} \hat{b}_h(\theta, A, z) dQ(z) + \int_{\hat{z}_h(A)}^{\bar{z}} \hat{g}_h(\theta, A, z) dQ(z),$$

where $\hat{z}_h(A)$ is a shortcut for $\hat{z}_h(\theta, A)$ and is independent of θ because $l(\theta)$ is invariant in θ .⁵⁰ This means that

$$\frac{\partial u_h(\theta, A)}{\partial \theta} = \int_{\underline{z}}^{\hat{z}_h(A)} \frac{\partial \hat{b}_h(\theta, A, z)}{\partial \theta} dQ(z) + \int_{\hat{z}_h(A)}^{\bar{z}} \frac{\partial \hat{g}_h(\theta, A, z)}{\partial \theta} dQ(z).$$

Using the change in variables $\omega = -\Phi^{-1}(1 - A)$, and the fact that, for any x , $\phi(x) = \phi(-x)$, we then have that $\int_0^1 (-\Phi^{-1}(1 - A)) dA = \int_{-\infty}^{\infty} \omega \phi(\omega) d\omega = 0$.

Equity. Under the assumption that $\rho_S = 1$, we have that

$$\frac{\partial \hat{b}_E(\theta, A, z)}{\partial \theta} = 0 \quad \text{and} \quad \frac{\partial \hat{g}_E(\theta, A, z)}{\partial \theta} = \frac{V'(\theta)}{1+q},$$

implying that $\frac{\partial u_E(\theta, A)}{\partial \theta} = \frac{V'(\theta)}{1+q} (1 - Q(\hat{z}_E(A)))$.

For any $x > \hat{\theta}$, we thus have that

$$\begin{aligned} \int_0^1 \frac{\partial u_E(x, A)}{\partial \theta} (-\Phi^{-1}(1 - A)) dA &= \frac{V'(x)}{1+q} \int_0^1 (1 - Q(\hat{z}_E(A))) (-\Phi^{-1}(1 - A)) dA \\ &= \frac{V'(x)}{1+q} \text{cov}(1 - Q(\hat{z}_E(A)), -\Phi^{-1}(1 - A)) > 0 \end{aligned}$$

⁵⁰Recall that $\hat{z}_h(\theta, A)$ is implicitly defined by the solution to $l(\theta) + \rho_S q p_h^*(A, z) = D$ whenever the equation has a solution, is equal to \underline{z} when $l(\theta) + \rho_S q p_h^*(A, \underline{z}) > D$, and is equal to \bar{z} when $l(\theta) + \rho_S q p_h^*(A, \bar{z}) < D$.

where the inequality follows from the fact that $\hat{z}_E(A)$ is decreasing in A .

Debt. Using the fact that

$$\frac{\partial \hat{b}_D(\theta, A, z)}{\partial \theta} = \frac{L_D}{qL_D + D} \gamma'(\theta) \quad \text{and} \quad \frac{\partial \hat{g}_D(\theta, A, z)}{\partial \theta} = 0,$$

we have that $\frac{\partial u_D(\theta, A)}{\partial \theta} = \frac{L_D}{qL_D + D} \gamma'(\theta) Q(\hat{z}_D(A))$, implying that, for any $x > \hat{\theta}$,

$$\begin{aligned} \int_0^1 \frac{\partial u_D(x, A)}{\partial \theta} (-\Phi^{-1}(1 - A)) dA &= \frac{L_D}{qL_D + D} \gamma'(\theta) \int_0^1 Q(\hat{z}_D(A)) (-\Phi^{-1}(1 - A)) dA \\ &= \frac{L_D}{qL_D + D} \gamma'(\theta) \text{cov}(Q(\hat{z}_D(A)), -\Phi^{-1}(1 - A)) < 0 \end{aligned}$$

where the inequality follows again from the fact that $\hat{z}_D(A)$ is decreasing in A . The lemma then follows from combining the results from steps 1 and 2. ■

Now observe that, for any precision of private information σ^{-2} , and any monotone pass/fail policy with threshold $\hat{\theta}$, after the signal $s = 1$ is disclosed, purchasing the bank's security h is the unique rationalizable action for all investors if, and only if, $\psi_h(x, \hat{\theta}, \sigma) > 0$ for all $x \in \mathbb{R}$ (the arguments are analogous to the ones in the proof of Theorem 2). From the discussion following Theorem 5, then observe that the threshold $\theta_h^*(\sigma)$ defining the optimal monotone policy when the precision of the investors' information is σ^{-2} is given by $\theta_h^*(\sigma) = \inf \left\{ \hat{\theta} : \psi_h(x, \hat{\theta}, \sigma) \geq 0 \text{ for all } x \in \mathbb{R} \right\}$.

Next, let $x_h^*(\sigma) \equiv \arg \min_{x \in \mathbb{R}} \psi_h(x, \theta_h^*(\sigma), \sigma)$ and note that $x_h^*(\sigma)$ is a solution to the equation $\psi_h(x_h^*(\sigma), \theta_h^*(\sigma), \sigma) = 0$. Also observe that $\lim_{\sigma \rightarrow 0^+} \theta_h^*(\sigma) = \theta_h^{MS}$, for $h = E, D$.

Now recall that, by definition of $\theta_h^\#$ along with the monotonicity of $u_h(\theta, 1/2)$ in θ , for any $\theta > \theta_h^\#$, $u_h(\theta, 1/2) > 0$. Hence, for any $\hat{\theta} > \theta_h^\#$, any $\sigma > 0$, and any $x \leq \hat{\theta}$,

$$\psi_h(x, \hat{\theta}, \sigma) = \int_{\hat{\theta}}^{\infty} u_h\left(\theta, 1 - \Phi\left(\frac{x - \theta}{\sigma}\right)\right) \frac{\phi\left(\frac{x - \theta}{\sigma}\right)}{\sigma \Phi\left(\frac{x - \hat{\theta}}{\sigma}\right)} d\theta > 0. \quad (21)$$

The assumption that $\theta_h^{MS} > \theta_h^\#$, $h = E, D$, along with the continuity of $\theta_h^*(\sigma)$ in σ , then imply that there exist $\hat{\sigma} > 0$, such that, for any $\sigma \in (0, \hat{\sigma})$, $\theta_h^*(\sigma) > \theta_h^\#$, for $h = E, D$. The result in (21) then implies that, for any such σ , $x_h^*(\sigma) > \theta_h^*(\sigma)$.

Next, for any $\hat{\theta} \in (\underline{\theta}, \bar{\theta})$, any $\sigma > 0$, let $\Psi_h(\hat{\theta}, \sigma) \equiv \inf_{x \in \mathbb{R}} \psi_h(x, \hat{\theta}, \sigma)$ and, for any $(\hat{\theta}, \sigma)$ such that $\arg \min_{x \in \mathbb{R}} \psi_h(x, \hat{\theta}, \sigma) \neq \emptyset$, let $x_h^{**}(\hat{\theta}, \sigma) \in \arg \min_{x \in \mathbb{R}} \psi_h(x, \hat{\theta}, \sigma)$. Note that, when $\tilde{\sigma} = \sigma$, $x_h^{**}(\theta_h^*(\sigma), \tilde{\sigma}) = x_h^*(\sigma)$.

Now observe that the arguments in the proof of Lemma 4 imply that there exists $\sigma^\# > 0$ small such that, for any $\sigma \in (0, \sigma^\#)$, any $x > \theta_h^*(\sigma)$, $\psi_h(x, \theta_h^*(\sigma), \sigma)$ is differentiable in σ with a derivative that is continuous in σ and uniformly bounded over $\{(\sigma, x) \in (0, \sigma^\#) \times \mathbb{R} : x > \theta_h^*(\sigma)\}$, $h = E, D$. Now let $\bar{\sigma} \equiv \min\{\sigma^\#, \hat{\sigma}\}$. The envelope theorem, along with the property above, then implies that, for any $\sigma, \tilde{\sigma} \in (0, \bar{\sigma})$, with $\tilde{\sigma} \geq \sigma$, and $h = E, D$,

$$\left. \frac{\partial}{\partial \sigma} \Psi_h(\hat{\theta}, \tilde{\sigma}) \right|_{\hat{\theta} = \theta_h^*(\sigma)} = \left. \frac{\partial}{\partial \sigma} \psi_h(x, \hat{\theta}, \tilde{\sigma}) \right|_{x = x_h^{**}(\theta_h^*(\sigma), \tilde{\sigma}), \hat{\theta} = \theta_h^*(\sigma)}$$

with $\text{sign} \left(\frac{\partial}{\partial \sigma} \psi_h \left(x, \hat{\theta}, \tilde{\sigma} \right) \Big|_{x=x_h^{**}(\theta_h^*(\sigma), \tilde{\sigma}), \hat{\theta}=\theta_h^*(\sigma)} \right) = \text{sign} \left(\lim_{\sigma \rightarrow 0^+} \frac{\partial}{\partial \sigma} \psi_h \left(x_h^*(\sigma), \theta_h^*(\sigma), \sigma \right) \right)$.

The above properties, along with Lemma (4) and the fundamental theorem of calculus, then imply that, for any $\sigma, \sigma' \in (0, \bar{\sigma})$, with $\sigma' > \sigma$,

$$\Psi_E(\theta_E^*(\sigma), \sigma') = \int_{\tilde{\sigma}=\sigma}^{\tilde{\sigma}=\sigma'} \frac{\partial}{\partial \sigma} \Psi_E(\hat{\theta}, \tilde{\sigma}) \Big|_{\hat{\theta}=\theta_E^*(\sigma)} d\tilde{\sigma} > 0 > \int_{\tilde{\sigma}=\sigma}^{\tilde{\sigma}=\sigma'} \frac{\partial}{\partial \sigma} \Psi_D(\hat{\theta}, \tilde{\sigma}) \Big|_{\hat{\theta}=\theta_D^*(\sigma)} d\tilde{\sigma} = \Psi_D(\theta_D^*(\sigma), \sigma').$$

The result in the proposition then follows from the above properties along with the monotonicity of $\Psi_h(\cdot, \sigma')$ in the truncation point θ_h^* , $h = E, D$. Q.E.D.

References

- Albagli, E., Hellwig, C., Tsyvinski, A., 2015. A theory of asset prices based on heterogeneous information.
- Allen, F., Gale, D., 1998. Optimal financial crises. *Journal of Finance* 53 (4), 1245–1284.
- Alonso, R., Camara, O., 2016a. Persuading voters. *American Economic Review* 106 (11), 3590–3605.
- Alonso, R., Camara, O., 2016b. Bayesian persuasion with heterogeneous priors. *Journal of Economic Theory* 165, 672–706.
- Alvarez, F., Barlevy, G., 2015. Mandatory disclosure and financial contagion. WP, NBER.
- 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.
- Arieli, I., Babichenko, Y., 2019. Private bayesian persuasion. *Journal of Economic Theory* 182, 185–217.
- Bardhi, A., Guo, Y., 2017. Modes of persuasion toward unanimous consent. *Theoretical Economics*, forthcoming.
- Basak, D., Zhou, Z., 2019. Diffusing coordination risk. *American Economic Review* forthcoming.
- Basak, D., Zhou, Z., 2020. Timely persuasion. Mimeo, Indian School of Business and Tsinghua University.
- Bergemann, D., Morris, S., 2016a. Bayes correlated equilibrium and the comparison of information structures in games. *Theoretical Economics* 11 (2), 487–522.
- Bergemann, D., Morris, S., 2016b. Information design, bayesian persuasion and bayes correlated equilibrium. *The American Economic Review* 106 (5), 586–591.
- Bergemann, D., Morris, S., 2019. Information design: A unified perspective. *Journal of Economic Literature* 57(1), 44–95.

- 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.
- Brunnermeier, M. K., Pedersen, L. H., 2005. Predatory trading. *The Journal of Finance* 60 (4), 1825–1863.
- Calzolari, G., Pavan, A., 2006a. On the optimality of privacy in sequential contracting. *Journal of Economic theory* 130 (1), 168–204.
- Calzolari, G., Pavan, A., 2006b. Monopoly with resale. *The RAND Journal of Economics* 37 (2), 362–375.
- Chan, J., Gupta, S., Li, F., Wang, Y., 2019. Pivotal persuasion. *Journal of Economic theory* 180, 178–202.
- Che, Y.-K., Hörner, J., 2018. Recommender systems as mechanisms for social learning. *Quarterly Journal of Economics* 133 (2), 871–925.
- Cong, L. W., Grenadier, S., Hu, Y., et al., 2016. Dynamic coordination and intervention policy. Tech. rep.
- Copeland, A., Martin, A., Walker, M., 2014. Repo runs: Evidence from the tri-party repo market. *The Journal of Finance* 69 (6), 2343–2380.
- Corona, C., Nan, L., Gaoqing, Z., 2017. The coordination role of stress test disclosure in bank risk taking. WP, Carnegie Mellon University.
- Covitz, D., Liang, N., Suarez, G. A., 2013. The evolution of a financial crisis: Collapse of the asset-backed commercial paper market. *The Journal of Finance* 68 (3), 815–848.
- Denti, T., 2020. Unrestricted information acquisition. WP, Cornell University.
- Diamond, D. W., Dybvig, P. H., 1983. Bank runs, deposit insurance, and liquidity. *Journal of Political Economy* 91 (3).
- Doval, L., Ely, J., 2017. Sequential information design. WP, Northwestern University.
- Dworczak, P., 2019. Mechanism design with aftermarkets: Cutoff mechanisms. WP, Northwestern University.
- Dworczak, P., Pavan, A., 2019. Robust (bayesian) persuasion. WP, Northwestern University.
- Edmond, C., 2013. Information manipulation, coordination, and regime change. *Review of Economic Studies* 80 (4), 1422–1458.
- Ely, J. C., 2017. Beeps. *The American Economic Review* 107 (1), 31–53.
- Farhi, E., Tirole, J., 2012. Collective moral hazard, maturity mismatch, and systemic bailouts. *American Economic Review* 102, 60–93.
- Faria-e Castro, M., Martinez, J., Philippon, T., 2016. Runs versus lemons: information disclosure and fiscal capacity. *The Review of Economic Studies*, 060.
- Flannery, M., Hirtle, B., Kovner, A., 2017. Evaluating the information in the federal reserve stress tests. *Journal of Financial Intermediation* 29 (29), 1–18.
- Garcia, F., Panetti, E., 2017. A theory of government bailouts in a heterogeneous banking union. WP, Indiana University.

- Gentzkow, M., Kamenica, E., 2016. A Rothschild-Stiglitz approach to bayesian persuasion. *American Economic Review*, P&P 106 (5), 597–601.
- Gick, W., Pausch, T., 2012. Persuasion by stress testing: Optimal disclosure of supervisory information in the banking sector. WP, Harvard University and Deutsche Bundesbank.
- Gitmez, A., Molavi, P., 2019. Media capture: A bayesian persuasion approach. WP, MIT.
- Goldstein, I., Huang, C., 2016. Bayesian persuasion in coordination games. *The American Economic Review* 106 (5), 592–596.
- Goldstein, I., Leitner, Y., 2018. Stress tests and information disclosure. *Journal of Economic Theory* 177, 34–69.
- Goldstein, I., Pauzner, A., 2015. Demand–deposit contracts and the probability of bank runs. *Journal of Finance* 60 (3), 1293–1327.
- 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.
- Gorton, G., Metrick, A., 2012. Securitized banking and the run on repo. *Journal of Financial economics* 104 (3), 425–451.
- Guo, Y., Shmaya, E., 2019. The interval structure of optimal disclosure. *Econometrica* 87 (2), 653–675.
- Henry, J., Christoffer, K., 2013. A macro stress testing framework for assessing systemic risks in the banking sector. WP, European Central Bank.
- Homar, T., Kick, H., Salleo, C., 2016. Making sense of the **EU** wide stress test: a comparison with the srisk approach. WP, European Central Bank.
- Inostroza, N., 2019. Persuading multiple audiences: An information design approach to banking regulation. WP, University of Toronto.
- Kamenica, E., Gentzkow, M., 2011. Bayesian persuasion. *American Economic Review* 101, 2590–2615.
- Kolotilin, A., Mylovanov, T., Zapechelnuyk, A., Li, M., 2017. Persuasion of a privately informed receiver. *Econometrica* 85 (6), 1949–1964.
- Kyle, A. S., 1985. Continuous auctions and insider trading. *Econometrica: Journal of the Econometric Society*, 1315–1335.
- Laclau, M., Renou, L., 2017. Public persuasion. WP, PSE.
- Lerner, J., Tirole, J., 2006. A model of forum shopping. *American economic review* 96 (4), 1091–1113.
- Li, F., Yangbo, S., Zhao, M., 2019. Global manipulation by local obfuscation. WP, UNC.
- Mathevet, L., Perego, J., Taneva, I., 2019. On information design in games. *Journal of Political Economy*, forthcoming.
- Morgan, D., Persitani, S., Vanessa, S., 2014. The information value of the stress test. *Journal of Money, Credit and Banking*, Vol. 46, No. 7.
- Morris, S., Daisuke, O., Takahashi, S., 2019. Information design in binary-action supermodular games. WP, MIT.

- Morris, S., Shin, H. S., 2006. Global games: Theory and applications. *Advances in Economics and Econometrics*, 56.
- Morris, S., Yang, M., 2019. Coordination under continuous stochastic choice. WP, MIT.
- Myerson, R. B., 1986. Multistage games with communication. *Econometrica*, 323–358.
- Orlov, D., Zryumov, P., Skrzypacz, A., 2017. Design of macro-prudential stress tests. WP, Stanford University.
- Pérignon, C., Thesmar, D., Vuillemeys, G., 2018. Wholesale funding dry-ups. *The Journal of Finance* 73 (2), 575–617.
- Petrella, G., Resti, A., 2013. Supervisors as information producers: Do stress tests reduce bank opaqueness? *Journal of Banking and Finance* 37 (12), 5406–5420.
- Rayo, L., Segal, I., 2010. Optimal information disclosure. *Journal of Political Economy* 118 (5), 949–987.
- Rochet, J., Vives, X., 2004. Coordination failures and the lender of last resort: Was bagethor right after all? *Journal of the European Economic Association* 2 (6), 1116–1147.
- Shimoji, M., 2017. Bayesian persuasion in unlinked games. WP, University of York.
- Szkup, M., Trevino, I., 2015. Information acquisition in global games of regime change. *Journal of Economic Theory* 160, 387–428.
- Taneva, I. A., 2019. Information design. *American Economic Journal: Microeconomics* forthcoming.
- Williams, B., 2017. Stress tests and bank portfolio choice. WP, New York University.
- Yang, M., 2015. Coordination with flexible information acquisition. *Journal of Economic Theory* 158, 721–738.