

Inside and Outside Information

Daniel Quigley

University of Oxford

Ansgar Walther

Imperial College London and CEPR*

July 2022

Abstract

We study an economy with financial frictions in which a regulator designs a test that reveals *outside information* about a firm's quality to investors. The firm can also disclose verifiable *inside information* about its quality. We show that the regulator optimally aims for “public speech and private silence”, which is achieved with tests that give insiders an incentive to stay quiet. We fully characterize optimal tests by developing tools for Bayesian persuasion with incentive constraints, and use these results to derive novel guidance for the design of bank stress tests, as well as socially optimal corporate credit ratings.

*Quigley: Exeter College and Department of Economics, University of Oxford, Oxford OX1 3DP, UK. Email: daniel.quigley@economics.ox.ac.uk. Walther: Imperial College Business School, London SW7 2AZ, UK, and CEPR. Email: a.walther@imperial.ac.uk. We are grateful to our discussants Pierre Chaigneau, Nadya Malenko, Tim MacQuade, and Alessio Piccolo for their feedback. We also thank Heski Bar-Isaac, James Best, Vince Crawford, Peter Eso, Nicolas Inostroza, Ian Jewitt, Emir Kamenica, Alex Kohlhas, Paul Klemperer, Meg Meyer, Ines Moreno de Barreda, David Myatt, Larry Samuelson, Joel Shapiro, Anton Tsoy, Victoria Vanasco, Adrien Vigier, Selma Walther, Lucy White, Peyton Young, and audiences at various seminars and conferences, for their comments. A previous version of this paper was circulated under the title “Crowding out disclosure”.

1 Introduction

During the financial crisis of 2007-9, regulators such as the Federal Reserve used stress tests to alleviate investors’ uncertainty about the banking system. These stress tests are a form of *outside* information about firms that is released to investors. In addition, banks and other firms have wide discretion over their public corporate disclosures, which reveal *inside information* to investors. This potential interplay between inside and outside information appears frequently in financial markets. For instance, credit ratings agencies provide outside information about borrowers, who can also make their own inside disclosures.

In this paper, we analyze how outside information should be designed in an economy with financial frictions. Unlike the growing literature on this question (e.g., Goldstein and Leitner, 2018; Orlov et al., 2021; Inostroza and Pavan, 2022) we take into account that inside disclosures may respond endogenously to the design of outside information. We show that accounting for these responses changes the optimal design of outside information both qualitatively and quantitatively.

To motivate our analysis, imagine the Fed is designing a stress test to reveal outside information about the unobserved quality θ of a bank’s assets. Investors will run on the bank if their expectation of asset quality, after observing the test result, drops below a critical threshold t . Suppose the Fed wants to minimize the probability of runs. Using tools from the literature on Bayesian persuasion, it is easy to show that—in the absence of inside information—the optimal policy is a simple pass/fail test, which splits banks around a quality threshold θ^* . Banks with quality $\theta < \theta^*$ publicly “fail”, resulting in a run with probability one. Banks with assets of quality $\theta \geq \theta^*$ “pass” and avoid a run: The failure threshold θ^* is chosen so that the pass grade is just credible enough to save the bank, that is $E[\theta|\theta \geq \theta^*] = t$. Intuitively, this test achieves higher welfare than, say, full disclosure, because it pools strong types $\theta \geq t$ with as many weak types as possible, thus avoiding runs on vulnerable banks.¹

Suppose now that banks can verifiably disclose θ at some cost in anticipation of the test result, and consider the incentives of strong banks. If they do not disclose, they will avoid a run but investors will regard their assets as worth only t . As long as banks care about market perceptions beyond the immediate need to avoid a run, disclosure may be a best response for the strong, so long as disclosure costs are not too large. Once strong types disclose, investors’ expectations conditional on seeing a “pass” result without inside disclosure drop below the threshold t , because the strongest types of bank are no longer pooled into the

¹In richer models of bank runs (e.g., Inostroza, 2019), multiple failing grades can be optimal, but the unique passing grade is still defined and motivated in the same way.

passing grade. Thus, all non-disclosing banks now face a run. This is an instance of the classical “unraveling” of inside information, as first described by Grossman and Hart (1980). Indeed, whenever the simple pass/fail test incentivizes any inside disclosures among passing banks, our results show it cannot be optimal. In this sense, naive application of pass/fail tests do not tackle financial instability in an efficient way.

We study a class of models with financial frictions, deriving tests that optimally account for firms’ incentives to make disclosures. As we demonstrate in a series of examples, our general model nests models of bank runs along the lines of our example above (e.g., Diamond and Dybvig, 1983; Morris and Shin, 2000), a model of corporate financing with contracting frictions as in Holmstrom and Tirole (1997), in which outside information can be interpreted as the release of credit ratings, and richer treatments of financial instability such as Inostroza (2019). The important common feature of these applications is that financial constraints begin to bind whenever investors’ expectations about an indicator θ of underlying quality drop below a critical threshold. Moreover, welfare is typically a convex, increasing function of investors’ expectations when financial constraints are binding.

In our model, a regulator can commit to an arbitrary test that reveals outside information about θ , as in the literature on Bayesian persuasion (e.g., Kamenica and Gentzkow, 2011; Dworzak and Martini, 2019). After learning θ , firms / banks can either stay quiet or make additional inside disclosures in anticipation of the test result. We model inside disclosures as the firm choosing from a space of possible messages to investors. This space includes the possibility of costly verifiable messages, which credibly narrow down the set of possible types θ .

A key result in our characterization of optimal tests is that it is always optimal for the regulator to encourage *public speech and private silence*: without loss, the regulator can focus on designing tests that sustain “opaque” equilibria, in which firms make no inside disclosures. This result is reminiscent of a revelation principle, but requires a slightly different proof. In addition to arguing that the regulator can replicate any disclosures that firms would otherwise make, we must characterize the conditions under which opacity is a (perfect Bayesian) equilibrium of the messaging game between firms and investors. This result contributes an interesting economic intuition in its own right, but also makes the design problem substantially more tractable. Indeed, the regulator’s problem now boils down to a Bayesian persuasion problem subject to a series of incentive compatibility constraints, which ensure that no type of firm wants to deviate from silence.

While this finding helps to simplify the search for an optimal test, the regulator’s problem is still non-trivial and cannot be solved with existing methods for Bayesian persuasion. We show how existing techniques can be extended to our model with incentive compatibility con-

straints. After outlining this methodological contribution, we provide a full characterization of optimal tests.

We show that optimal tests still have a threshold property, in the following sense: all firms with $\theta < \theta^*$ fail the test, meaning that they face binding financial frictions with probability one; all remaining firms pass. However, a variety of passing grades are employed: the regulator designs a standard “pass” grade that lifts firms just above the threshold where financial constraints bind, but also a range of novel grades that pass strong firms “with flying colors” and raise investors’ expectations discretely above the critical threshold. These novel grades are released to preserve strong firms’ incentives to stay quiet. The regulator trades off the cost of releasing these grades, which force her to raise the failure threshold θ^* relative to the benchmark without inside disclosures, against firms’ incentive compatibility constraint. We show that the optimal way to resolve this trade-off is to assign a stochastic grade, by which strong types are revealed with just enough probability to discourage inside disclosures.

While there are typically multiple implementations of optimal tests in our model, we prove that the key economic properties we have highlighted are necessary features of *any* optimal test: all optimal tests have the threshold property described above and, moreover, assign stochastic grades to types of firms with a binding incentive constraint. That is, in an optimally designed test, strong firms should be confident of passing the test, but unsure about the precise grade they will receive.

We supplement this characterization with a set of comparative statics. We show that the incentive constraints associated with private silence become tighter from the regulator’s perspective if i) inside disclosures become less costly, ii) disclosure messages become “more verifiable” in a natural sense, or iii) the payoffs to firms’ managers become more sensitive to investors’ beliefs. In each of these cases, the regulator is more constrained, which means that she must ultimately raise the threshold θ^* below which firms fail. We also show that, in each of these scenarios, the regulator must make the optimal tests more informative, in that investors’ posterior expectations undergo a mean-preserving spread. Finally, we show that, under some mild regularity conditions, the failure threshold also increases when investors’ prior beliefs become more pessimistic.

Before we move on to the financial applications of our model, two further features of our general results are worth highlighting. First, the idea of optimal private silence, which we have emphasized as part of our solution, extends even further. Concretely, we can also consider a broader design problem in which the regulator can privately communicate with firms, showing them a preview of the test result before firms decide on inside disclosures. We show that it is always optimal for the regulator not to send a preview – another dimension of private silence. In terms of our model environment, this result also means that the timing of

the game we assume, namely, that firms have to prepare disclosures before observing the test result, is an optimal arrangement from the policy-maker’s perspective. It is always beneficial for the regulator to have the “last word” in this sense.

Second, in an instructive special case of our model, private silence is the *only* motivation for releasing outside information. Consider the case where investors’ prior beliefs are benign, in the sense that financial constraints would not bind under the prior. In this scenario, strong firms can still have a strict incentive to make disclosures in order to separate themselves in investors’ expectations. The unraveling argument again applies, and a positive mass of firms can face binding frictions in equilibrium. We show that a regulator can do better by designing a test that fails a small proportion of weak firms but provides strong firms with an incentive to stay quiet. Therefore, while the regulator would prefer to release no information at all, and to preserve prior beliefs, she optimally decides to release some information that is geared towards preventing unraveling. This example strongly conveys the idea that inside disclosures motivate the release of more accurate outside information.

We further discuss the applied insights that arise from our theoretical analysis. In the headline application to bank runs and stress tests, we provide a full characterization of optimal, granular stress tests. In practice, regulators in the crisis of 2007-9 released information about banks’ losses in adverse macroeconomic scenarios. Details of these losses were published even for banks that “passed” the test in the sense that the regulator did not find a capital shortfall. One distinctive feature of optimal stress tests in our model is that they maintain some constructive ambiguity, assigning stochastic grades to strong banks. By contrast, weak banks face certainty about the grades they will receive, and never enjoy more surplus than is needed to marginally avoid a bank run. Our results also imply that it is strictly suboptimal to give banks a preview of stress test results, because this only serves to increase their incentives to make private disclosures to investors.

Our comparative statics provide additional guidance for regulators in practice. For instance, we demonstrate that optimal tests must be more informative (and more granular) when bank managers have high-powered incentives that depend on current market perceptions, and when it is relatively easy for the bank to make verifiable disclosures. Moreover, the special case described above suggests that at the “eve of a crisis”, when investors have an intermediate degree of pessimism, it can be valuable to run stress tests with the sole purpose of disincentivizing inside disclosures.

One can also use our framework to study the issue of regulatory time-inconsistency in stress testing, which has been highlighted by Bouvard et al. (2015) and Parlasca (2021) among others. Indeed, we can re-interpret the agent making inside disclosures in our model as the regulator’s future self, who might want to make disclosures after the true θ is realized,

in addition to the test that her past self has committed to. We argue that financial regulators’ objective function typically places less weight on short-term market expectations than bank managers. Under this assumption, the optimal test we characterize is time-consistent, at least in the sense that the regulator is not tempted to make further disclosures after learning θ .

In the context of credit ratings, our results are the first to show how a broad spectrum of optimal ratings emerges endogenously as a result of financial constraints. Our characterization of optimal tests can be interpreted as an ideal set of credit ratings, which involve “junk” grades that effectively fail the firm and lead to binding financing constraints, bunching of firms at “investment grade” ratings that just avoid financial constraints, and “premium” grades that transfer additional surplus to the firm in order to prevent excessive corporate disclosures. The mechanism behind this design is related to Daley et al. (2020), who show that accurate ratings can crowd out intermediaries’ incentives to acquire costly signals of credit quality. However, our characterization is novel in that it emphasizes the granularity of optimally *designed* ratings, while the typical signal structure in the existing literature is an exogenously structured, binary test. Moreover, a range of recent work emphasizes additional incentive problems arising from the interaction of firms, ratings agencies and investors (e.g., Skreta and Veldkamp, 2009; Bolton et al., 2012; Manso, 2013). We view our result as an interesting benchmark describing the type of ratings that maximize overall surplus.

In addition to the work we have cited above, this paper relates to the broader literature on verifiable disclosures (e.g., Milgrom, 1981; Verrecchia, 1983; Shin, 2003; Acharya et al., 2011; Bond and Zeng, 2021). We focus on an information designer’s goal of discouraging or “crowding out” private disclosures. Similar crowding-out effects have been studied in a few related papers, including Goldstein and Yang (2017), Einhorn (2018) and Frenkel et al. (2018) in models with verifiable disclosures, as well as Feltovich et al. (2002) and Daley and Green (2014) in the context of signaling games and Gigler and Hemmer (1998) in a cheap talk environment, but these papers do not contain an analysis of optimal information design.

In the literature on Bayesian persuasion and information design, Gentzkow and Kamenica (2017) and Li and Norman (2021) study multi-sender persuasion. As they assume all senders have ex ante commitment power, these models cannot address the novel incentive constraints introduced by interim disclosures. Our work is also connected to Guo and Shmaya (2019), who derive the optimal design of (outside) information subject to the constraints implied by the receiver’s private information, but our methodology and solution take a very different shape. Best and Quigley (2020) identify the conditions under which commitment in Bayesian persuasion can be implemented in repeated games; by contrast, we emphasize optimal information design when full commitment is not feasible. In the accounting literature, Friedman

et al. (2015) study the interaction between persuasion and disclosure in a model with binary signals, which precludes most of the effects we emphasize, while DeMarzo et al. (2019) and Bertomeu et al. (2021) study the incentives for firms to block the release of outside information in a setting without additional inside disclosures. Relative to these papers, we contribute a novel characterization of optimal, granular outside information that anticipates the associated incentives to disclose in a general framework.

The remainder of the paper is structured as follows: Section 2 describes our general model environment and provides a series of examples that highlight potential applications in finance. Section 3 characterizes optimal tests – it is designed so that the reader interested in applications can skip the technical Subsection 3.2, which develops our solution method. Section 4 discusses the applied insights arising from our results, and Section 5 concludes.

2 Model

We begin by describing a general information design problem in the presence of inside disclosures. We then present a series of examples to highlight the applications of our framework to banking and corporate finance.

2.1 General Environment

A firm and a regulator interact with a population of (risk-neutral) investors. The firm has assets with fundamental value $\theta \in [\underline{\theta}, \bar{\theta}] \equiv \Theta$, where $\bar{\theta} > \underline{\theta} \geq 0$. All agents have a prior belief that θ is drawn from a continuous cumulative distribution $F(\theta)$.

The regulator designs a *test* which reveals a public signal s of θ . A test is defined by a set of possible realizations $s \in S$ and a family $\{G(s | \theta)\}_{s \in S, \theta \in \Theta}$ of conditional cumulative distributions of test results. The firm, after privately observing θ , sends investors a message $m \in M(\theta)$. Finally, investors observe m , which we refer to as *inside information*, and the realization s of the test, which we call *outside information*. Investors' beliefs about θ now determine the posterior expected value of the firm's assets, which we define as

$$z = E_G[\theta | s, m]. \quad (1)$$

The regulator's utility is $w(z, \theta)$ and the firm's utility is $u(z, \theta) - c(m)$. Here, $c(m) \geq 0$ measures the cost of any inside disclosures.²

²That disclosure costs depend only on m is for notational ease only; as will become clear, the results extend immediately when costs are allowed to be continuous in m and θ .

Equilibrium Definition: For a given test G ,³ a perfect Bayesian equilibrium is defined by a messaging strategy for the firm and an expected value z for all realizations $\{s, m\}$. In equilibrium, the firm maximizes its expected utility after observing θ , and investors compute the expectation in Equation (1) using Bayes' rule on the equilibrium path. If investors observe a realization $\{s, m\}$ off the equilibrium path, they place zero probability on any type θ who cannot send m (i.e., $m \notin M(\theta)$), so that messages are potentially *verifiable*. The regulator's problem is to choose the test G to maximize her expected utility subject to the constraint that m , and hence z , will be determined endogenously in equilibrium. As is common in the information and mechanism design literatures, we assume that if a test induces multiple PBE, the regulator can select her preferred one.

Preferences: We assume the following functional form for the regulator's preferences:

$$w(z, \theta) = \theta(1 - \lambda(z)), \quad (2)$$

where $\lambda(z) \geq 0$ is a decreasing, concave function on the interval $[0, t]$, and equal to zero on the interval $[t, 1]$. As we demonstrate in more detail below, this functional form naturally measures social surplus in economies with *financial frictions*, which begin to bind when market valuations fall below a critical threshold t . The function $\lambda(z)$ is interpreted as the fraction of assets lost due to financial frictions. We assume throughout the paper that we are not in the trivial case where $\lambda(z) \equiv 0$.

We further assume that the firm's preferences are given by

$$u(z, \theta) = [\phi z + (1 - \phi)\theta](1 - \lambda(z)). \quad (3)$$

This expression again accounts for the fraction $\lambda(z)$ of asset values lost due to financial frictions. The firm's utility is given by a weighted average of investors' expected value z of the remaining assets, and their true fundamental value θ , with $\phi > 0$ denoting the weight on current market beliefs. For example, this formulation measures the profits of a firm who sells a fraction ϕ of its assets to outside investors and retains remaining fraction until maturity.⁴ Alternatively, Equation (3) could be interpreted as the utility of a manager whose compensation depends on both current and future asset values.

For our analysis of test design, it is also useful to define the regulator's (interim) expected utility given investors' updated belief z about the firm's assets. Indeed, taking expectations

³With a small abuse of notation, we write G to concisely describe the regulator's choice of test, which consists of the signal space S and the conditional distributions $G(s|\theta)$ described above.

⁴Daley et al. (2020) have characterized the case in which ϕ is endogenous and serves as a signal of asset quality. By contrast, we treat ϕ as a constant and focus on direct disclosures m about asset quality.




Figure 1: **Regulator's Expected Utility.** Each panel shows the regulator's expected utility $v(z)$ as a function of investors' expectation $z = E[\theta|s, m]$, as defined in Equation (4). We have set the lowest type $\underline{\theta}$ equal to zero for illustrative purposes. Panel (a) shows piecewise linear payoffs, with $\lambda(z) = \bar{\lambda} * 1\{z < t\}$, where $\bar{\lambda}$ is a constant. Notice that $\lambda(z)$ has a jump at $z = t$, which makes $v(z)$ discontinuous. Our model of bank runs in Example 1 induces payoffs of this type with $\bar{\lambda} = 1$. Panel (b) shows payoffs for which $\lambda(z)$ is a strictly convex and continuous function on the interval $[0, t]$, and equal to zero on the interval $[t, \bar{\theta}]$. These payoffs are induced by our model of corporate investment in Example 2. Panel (c) shows a possible combination of the previous two cases, with both convex and discontinuous payoffs on the interval $[0, t]$, which arises in Example 3.

of (2), and noting that $E[\theta|z] = z$ by definition, we can define

$$v(z) \equiv E[u(z, \theta) | z] = z(1 - \lambda(z)). \quad (4)$$

We show below that, when designing information, the regulator focuses on maximizing the expected value of $v(z)$. Importantly, whenever financial frictions bind, this objective function is convex in market expectations z and, otherwise, it is linear. Figure 1 illustrates the shape of $v(z)$ in three different parametric cases, which also map to our applied examples below.

Message Spaces: The set $M(\theta)$ defines the inside disclosures that the firm can make after learning its type θ . We assume that this set always contains a costless *cheap talk* or *quiet* message denoted $m = \emptyset$, with $c(\emptyset) = 0$. Other messages m in our model can be *verifiable* disclosures, in the sense that not all types are able to send them. A common example is $M(\theta) = \{\emptyset, \theta\}$, in which case that each type can either credibly disclose their true θ , or stay quiet. We keep the message space more general because, in the applications to finance that interest us, there is some variation in firms' ability to provide full, verifiable disclosures. In addition, it allows us to develop comparative statics of optimal tests with respect to the verifiability of disclosures. We make the technical assumptions that each $M(\theta)$

is compact and continuous in θ .⁵ Moreover, we assume the sets $M^{-1}(m) = \{\theta : m \in M(\theta)\}$ are continuous in m .

In our analysis we also make use of the following *worst case* disclosure payoff:⁶

$$\bar{u}(\theta) = \max_{m \in M(\theta)} \left\{ \min_{\theta' : m \in M(\theta')} u(\theta', \theta) - c(m) \right\}. \quad (5)$$

The payoff $\bar{u}(\theta)$ can be interpreted as follows. Suppose that the firm chooses some message m after observing its type θ . The payoff inside the curly braces in Equation (5) is the utility that the firm achieves, net of disclosure costs, if investors respond by forming the most skeptical possible beliefs. These skeptical beliefs attach probability 1 to the worst type θ' that has access to message m , implying that investors' expectation of θ becomes $z = \theta'$. For example, if m is the "quiet" message $m = \emptyset$, then skeptical beliefs yield $z = \underline{\theta}$; if type θ has a message m that no other type can send, then skeptical beliefs yield $z = \theta$. The payoff $\bar{u}(\theta)$ in Equation (5) is the highest utility that the firm can achieve across all messages if investors respond skeptically. We assume that $\bar{u}(\theta) < u(\theta, \theta)$ for all $\theta > \underline{\theta}$. In economic terms, this assumption means that any message certifying that the type is at least θ must incur strictly positive costs, except (trivially) for the very worst type.

Remarks on Commitment and Timing Assumptions: We make two further observations on our model environment. First, the regulator commits to a test G and cannot change the associated conditional distributions after she learns the realization of θ . This assumption is standard in the growing literature on information design and Bayesian persuasion. That said, our framework lends itself to exploring the limits of such commitments. Indeed, one can use the same environment to study equilibria in which the regulator's future self can send further verifiable messages to investors after learning θ . Formally, this interpretation applies as soon as we set $\phi = 0$ in Equation (3), so that the firm's and the regulator's preferences are identical. We discuss the implications of this type of limited commitment in detail when we revisit our application to bank stress tests in Section 4.1.

Second, the timing of our model assumes the firm's disclosure opportunity arises after learning θ , but *before* observing the test result s . In Appendix C, we generalize the regulator's design problem so that she can also provide the firm with a preview of s , or a noisy signal of this result, prior to choosing its disclosure m . In this extension, the timing of disclosures

⁵Throughout, the space $M = \cup_{\theta} M(\theta)$ is always endowed with some underlying topology. By continuity, we mean that the correspondence $M(\theta)$ is both upper and lower semicontinuous in the usual sense.

⁶Under our assumptions on M , and if the cost function $c(m)$ is continuous, the Theorem of the Maximum applies, so that (5) is well-defined and is moreover an upper semi-continuous function of θ . Details are available on request.

itself becomes a regulatory choice variable. We return to this exercise after stating some of our characterizations of optimal tests in the next section.

2.2 Applied Examples

We now illustrate the applied relevance of our model in three examples, each of which maps our general framework to a specific model of financial frictions. Moreover, after we characterize optimal tests, we return to these examples to extract practical insights from our theoretical results.

Example 1. [Bank runs – Morris and Shin (2000)] In this example, we interpret the firm as a bank who interacts with a population of depositors. Each depositor can withdraw a fixed claim of \$1 immediately, or roll over their claim until the bank’s assets mature and yield a net return θ . We adapt the model of bank runs in Morris and Shin (2000) to a setting with additional inside and outside information θ . In this setting, which we describe in detail in Appendix D, depositors observe any disclosure m by the bank and the test result s , which can be interpreted as a regulatory stress test. For simplicity, we assume that the worst type of bank still has positive net present value $\underline{\theta} \geq 0$, so that the regulator always has an incentive to prevent runs. This means that all banks in the model are illiquid, yet solvent.⁷ After observing $\{m, s\}$, every depositor decides whether to withdraw early. We show that there is a run, meaning that all depositors withdraw early, if and only if

$$z < \frac{\kappa}{2} \equiv t \tag{6}$$

where z is depositors’ expectation of θ , defined as in Equation (1), and κ is a parameter measuring the illiquidity of the bank’s assets. A regulator wants to maximize depositors’ joint utility, which is given by θ if there is no bank run, and by a constant (normalized to zero) if there is a bank run. Thus, the regulator’s objective function is described by Equation (2), and the function $\lambda(z)$ capturing financial frictions in this example is a step function:

$$\lambda(z) = \begin{cases} 1, & z < t, \\ 0, & z \geq t. \end{cases} \tag{7}$$

A similar specification arises in Bouvard et al. (2015) and Goldstein and Leitner (2018).⁸ The bank’s managers choose inside disclosures to maximize the utility function in Equation

⁷It is easy to show, in addition to our characterizations below, that the optimal test design for insolvent banks would be to fail them, i.e., to reveal their true type, with probability one.

⁸In Goldstein-Leitner, the loss implied by a run is a constant in their model, while in our model it is a fraction of assets. With minor notational changes, all our results below also go through in their specification.

(3). As in the general model, $\phi > 0$ in this expression can reflect either a bank that sells a fraction of the its assets to outside investors, or managers who care about short-term results in the sense of their bank's current market valuation.

Example 2. [Corporate investment – Holmstrom and Tirole (1997)] Consider a firm who has an endowment of $A < 1$ units of cash and is able to make investments $I \in [0, 1]$ in a productive technology, which yields $(1 + \theta)I$ at a future date. We again assume that all firms have positive net present value, with $\underline{\theta} \geq 0$. The firm sells a security to risk-neutral investors, which entitles them to a fraction ϕ of the firm's future returns. Investors observe any disclosure m by the firm and the result of a credit rating s before buying the security. Investor's willingness to pay for the firm's security is then $\phi(1 + z)I$, where z is defined in Equation (1). The firm's budget constraint for investment is

$$A + \phi(1 + z)I \geq I.$$

Define the firm's maximum investment capacity $I(z)$ as:

$$I(z) = \min \left\{ 1, \frac{A}{1 - \phi(1 + z)} \right\}.$$

If $I(z) < 1$, then financial constraints bind, in the sense that the firm cannot invest up to its technological capacity. This occurs when market expectations satisfy

$$z < \frac{1 - A - \phi}{\phi} \equiv t.$$

Whenever this inequality holds, the fraction of valuable investments that is lost is $\lambda(z) = 1 - I(z)$, a concave function of z . Notice that financial constraints are never binding when $\phi = 1$ or if $A \geq 1$. Thus, frictions in this model are driven by the firm's need to attract outside financing while also retaining at least some of its assets. This requirement is a common motivation for frictions in corporate finance, for example, when retention is necessary to maintain appropriate managerial incentives as in Holmstrom and Tirole (1997). With these definitions, we obtain a special case of our general model. The regulator, who designs credit ratings to maximize the utility function in Equation (2), can be interpreted as a ratings agency who wants to maximize the joint surplus of firms and investors.

Example 3. [Cash flow and Liquidity Constraints – Inostroza (2019)] Inostroza (2019) studies the interaction between firms' cash flows and liquidity constraints in an environment motivated by bank runs. We discuss the key qualitative effects in his model, which generates objective functions that are a special case of our general framework. In Inostroza's

model, a bank raises cash by issuing securities against an investment project with unknown value θ and, at a future date, faces the threat of investor runs that can reduce the value of its initial investment. This generates a feedback loop: Any increase in investors' expectation z of θ raises their willingness to provide the bank with cash. This cash injection reduces the probability of a run, which further raises investors' willingness to pay. As a result, social surplus is strictly convex in z up to a threshold $z = t$, above which liquidity constraints stop binding and surplus becomes linear in z . More broadly, this intuition highlights an interaction between cash flow and liquidity constraints in corporate financing. One can view this environment as a third possible micro-foundation of our model.

3 Characterization of Optimal Tests

In this section, we derive necessary and sufficient conditions for an optimal test G . We present our analysis in three parts. In Section 3.1, we simplify the regulator's problem by showing that, without loss, she can focus on tests that discourage all inside disclosures. This allows us to develop the idea that the regulator optimally aims for *public speech and private silence*, in the sense that the firm stays quiet and investors rely only on outside regulatory information. The regulator's desire to implement private silence introduces an incentive compatibility constraint into the regulator's problem, which ensures that firms have no incentive to make inside disclosures. In Section 3.2, we extend existing methods from the literature on Bayesian persuasion in order to solve the constrained version of the regulator's problem. This subsection highlights an auxiliary methodological contribution of our work, but readers who are interested mostly in the applied implications can skip it. In Section 3.3, we use our method to fully characterize optimal tests and discuss their economic implications.

3.1 Public Speech and Private Silence

The regulator's problem in our model is complex because she designs the test G while anticipating that firms will choose their disclosure strategies endogenously in perfect Bayesian equilibrium. In order to make this problem tractable, we show that the regulator can focus on tests that discourage all inside disclosures in equilibrium. In other words, it is always optimal for the regulator to aim for public speech and private silence:

Lemma 1. [*Public Speech and Private Silence*] *For any test G , there exists an alternative test \hat{G} such that*

1. For any equilibrium under G , there is an equilibrium under \hat{G} in which the regulator obtains the same expected utility, and
2. The firm makes no verifiable disclosures in the equilibrium under \hat{G} , choosing $m = \emptyset$ with probability 1.

Lemma 1 shows that, given any test G and any associated equilibrium disclosures, the regulator can always do equally well with a test that induces the firm to stay quiet after learning θ . The first step of the argument is similar to a classical revelation principle: Since the regulator has full flexibility in designing the set of test results S and the conditional distributions $G(s|\theta)$, she can always construct a modified test \hat{G} , which replicates whatever disclosure the firm would have made in any equilibrium under G . In principle, one might worry that the new test may introduce new opportunities for the firm to affect investors' beliefs via (partially verifiable) disclosures. The second step, which goes slightly beyond the standard argument, shows that \hat{G} in fact does induce a perfect Bayesian equilibrium without any verifiable disclosures. This step involves constructing a set of beliefs for investors, including off the equilibrium path, that sustains silence as an optimal strategy for the firm regardless of its private information θ .

In the remainder of the paper, we therefore focus on tests that induce silence. An interesting property of these tests is that disclosure costs are zero on the equilibrium path, since the firm always sends the “quiet” message $m = \emptyset$ and the associated costs are $c(\emptyset) = 0$.⁹

Given Lemma 1, we can reframe the regulator's problem in a much more tractable way. First, we can simplify the regulator's expected payoff using the fact that she focuses on tests that induce private silence. Let G be such a test. Given that firms stay quiet and so are uninformative, investors' expectations under G are therefore given by the random variable $z = E_G[\theta|s]$. As the regulator is the sole source of information in any such test, from now on we will refer to z as the *grade* that the firm achieves under the test, and G its distribution.

Taking the conditional expectation $E_G[w(z, \theta)|z]$ of the regulator's payoff in Equation (2), and then applying the law of iterated expectations, we find that the regulator's unconditional expected utility under G is given by

$$\begin{aligned} E_G[w(z, \theta)] &= E_G[z(1 - \lambda(z))] \\ &= E_G[v(z)]. \end{aligned}$$

Therefore, the regulator's problem is to design a test that keep firms quiet but also maximizes

⁹Notice that Lemma 1 is strengthened if the regulator's utility function in Equation (2) includes the firm's costs of disclosure $c(m)$. Therefore, our results below apply regardless of whether the costs of disclosure are private or social costs.

the expected value of $v(z)$. Notice that, under our general assumptions and in all our examples, this function is convex on the interval $[0, t]$ and linear on the interval $[t, 1]$, with a potential upward jump at the point $z = t$ (e.g., in the bank run application – see Figure 1, panel (a)). The following result builds on this derivation to provide a useful formulation of the regulator’s problem:

Lemma 2. [*Regulator’s Problem*] *The regulator finds an optimal test G by solving the following problem:*

$$V = \max_G \quad E_G[v(z)] \quad \text{subject to} \\ E_G[u(z, \theta) | \theta] \geq \bar{u}(\theta), \text{ for all } \theta, \quad (8)$$

where $\bar{u}(\theta)$ is the worst-case disclosure payoff, defined in Equation (5).

Lemma 2 employs the above derivation of the regulator’s objective function. Furthermore, it contains a tractable characterization of the incentive constraint determining “private silence”. Indeed, in order to satisfy this constraint, the regulator can do no worse than to focus on equilibria in which investors form the most skeptical possible beliefs when they see inside disclosures. In such equilibria, the highest payoff that the firm with type θ can earn by deviating from silence is exactly the worst case disclosure payoff $\bar{u}(\theta)$, defined in (5). The incentive compatibility constraint now reduces to the requirement that the firm’s expected payoff $E_G[u(z, \theta) | \theta]$ if it stays quiet, is at least equal to $\bar{u}(\theta)$. Were a test G to violate this inequality for some θ , then that type would have access to some verifiable message that allows it to do better than waiting for the test result, even if investors react skeptically.

Notice that, in general, our model features a potential feedback mechanism whereby firms’ outside options depend on the test that the regulator designs. This effect arises because the regulator’s test influences equilibrium play and, hence, the inferences drawn by investors following different disclosures. However, Lemmas 1 and 2 demonstrate that one can focus on a class of tests in which these feedback effects are absent. For any test that induces silence, private disclosures are off the equilibrium path. In the regulator’s preferred constellation of off-path beliefs, investors act skeptically following any disclosure, attaching probability one to the worst type that can make this disclosure. Hence, within the class of silence-inducing signals on which we focus, the value of firms’ outside option (i.e., of deviating to verifiable disclosure) is independent of the regulator’s information design.

The regulator’s problem in Lemma 2 further highlights the role of timing in our model. The constraint in problem (8) captures the idea that the firm has to decide on disclosures before observing the result s of the regulator’s test G . The problem would take a different shape if the firm was able to wait until it observes the realization of s and then make further

disclosures. This contrast is worth discussing in more detail. On one hand, preemptive disclosures are empirically plausible in some of the applications that motivate our paper. Financial markets react quickly to regulatory releases, leaving managers little time to prepare, audit and circulate verifiable inside information after the fact.¹⁰

On the other hand, it is perhaps more interesting to go further, by asking whether preemptive disclosures are part of an optimal institutional design. We address this question in Appendix C. As an additional feature of the design problem, we allow the regulator to show an arbitrary, potentially noisy preview of s to the firm before it makes its disclosure m . We show that the regulator can never gain by showing any informative preview. In this sense, the timing we have adopted can be viewed as part of an optimal design.

One of our core insights is that information design should proceed by ensuring public speech and private silence, as defined in Lemma 1. The discussion of regulatory previews, in the previous paragraph and Appendix C, also implies that this insight extends further. In fact, the regulator wishes to enforce private silence not only in the sense that the firm stays quiet when facing investors, but also in the sense that no information is privately communicated to the firm in advance of the regulatory test result. Practically, this means that regulators should aim to choose the timing and institutional detail of disclosures so as to have the “last word” before investors respond.

3.2 Bayesian Persuasion with Private Silence

While we have considerably simplified the regulator’s problem in Lemmas 1 and 2, the solution is still not trivial. A common approach to information design is to boil the problem down to a choice of cumulative distribution function $G(z)$ for grades/posterior expectations z . The regulator can then construct a test that induces this distribution of posteriors, so long as the prior distribution F is a mean-preserving spread of the chosen distribution of the regulator’s chosen distribution G (e.g., Dworzak and Martini, 2019). Unfortunately, this method cannot be applied to our model, where the regulator needs to consider the incentive compatibility constraint in (8). The reason for this discrepancy is that our incentive compatibility constraint is driven by the firm’s expected utility after learning the true realization θ of its type, which depends on firms’ higher-order expectations of investors’ beliefs, and is generally different from investor’s expectations z . Hence, the distribution of z is no longer sufficient to describe the regulator’s problem.

We extend the Dworzak-Martini method to our constrained problem. Central to our

¹⁰Daley et al. (2020) further discuss the relevance of preemptive signaling in the context of credit ratings. Technically, this timing is can also be motivated in a formally equivalent game, where outside information, such as proprietary research, is privately observed by investors.

approach is another change of variables. Instead of z , we study the distribution of the *expected grade* $x = E_G[z|\theta]$ that type θ achieves, not yet knowing the realization of z . We let the regulator choose the cumulative distribution of x instead of z . This change preserves the basic simplification in Dworzak and Martini (2019): Any feasible test necessarily induces a $H(x)$ which is a mean preserving contraction of the prior F .

Key to our solution is the following result, which states a relaxed version of the regulator’s problem:

Lemma 3. [*Relaxed Problem*] *Consider the following problem, in which the regulator chooses the distribution $H(x)$ of expected grades $x = E_G[z|\theta]$:*

$$\begin{aligned} \bar{V} = \max_H \quad & \int v(x)dH(x) \text{ subject to} \\ & F \text{ is a mean-preserving spread of } H, \\ & H(x) \leq F \circ \delta(x) \quad , \text{ for all } x \geq t \end{aligned} \quad (9)$$

where $\delta(x) = \{\theta : \bar{u}(\theta) \geq u(x, \theta)\}$ is the set of types that prefer their outside option to obtaining expected grade x , and where $F \circ \delta(x)$ denotes the probability mass of these types under the prior F .¹¹ The maximized value of this problem is weakly higher than the maximized value in Problem (8), so that $\bar{V} \geq V$.

Lemma 3 describes a new maximization problem, in which the regulator directly chooses a distribution $H(x)$ over expected grades, subject to two sets of constraint. We show in the Appendix that (9) is a relaxed version of (8), in that it always achieves a higher value.

The formulation of problem (9) introduces a second constraint in addition to the requirement that F is a mean-preserving spread of H . To understand the second constraint, recall that $\bar{u}(\theta)$ is the value of deviating to disclosure for type θ of the firm. The constraint compares two measures: First, the mass of firms who prefer their outside option to an expected grade x is $1 - F \circ \delta(x)$. Second, for any candidate distribution $H(x)$ of expected grades, we can write the mass of types who obtain an expected grade above x as $1 - H(x)$. The second constraint in problem (9) replaces the incentive constraints with the requirement that the second quantity must be greater than the first for all $x \geq t$. The key advantage of (9) is that some of the results in Dworzak and Martini (2019) extend to this type of “first order stochastic dominance” constraint (see Appendix A.1.2).

Two additional technical points are worth considering. First, the objective function in problem (9) focuses on $x = E[z|\theta]$, which is not always a sufficient statistic for expected

¹¹More rigorously, $F \circ \delta(x) = \int 1_{\theta \in \delta(x)} dF(\theta)$. The integral is well-defined because the functions $\bar{u}(\theta)$ and $u(x, \theta)$ are measurable.

utility because the firm's objective function in our model is not linear in z when $z < t$. Second, we have imposed the second constraint in problem (9) only on outside options above the threshold t , effectively ignoring incentive constraints for firms with outside options below this threshold. The proof of Lemma 3 shows that there is no loss of generality associated with these shortcuts. Indeed, for any feasible test G in problem (8), there is a test G' which achieves weakly higher value, and for which i) types $\theta \geq t$ get grades $z \geq t$ with probability 1, confining these types to the linear part of the objective function, and ii) types $\theta < t$ do not have binding incentive constraints. We then show that the distribution of expected grades implied by G' also yields value V in problem (9), which implies that this is indeed a relaxed version of the regulator's problem.

We now turn to the solution of the regulator's problem, which leverages Lemma 3 to find optimal tests. Our strategy is to find a solution to the relaxed problem (9) that is also feasible in the regulator's original problem (8), and must therefore be optimal.

3.3 Optimal Tests

In Lemma 2, we stated the regulator's problem of finding an optimal test subject to the incentive compatibility constraint that firms must stay quiet in equilibrium. In addition, Lemma 3 in the previous subsection developed an auxiliary problem that is useful in solving for optimal tests. In Appendix A.2, we show how existing characterizations of optimal tests can be extended to solve this auxiliary problem. The following proposition uses these intermediate results to fully characterize an optimal solution to the regulator's problem:

Proposition 1. [*Characterization of Optimal Tests*] *The following conditions define an optimal test G in the regulator's problem (8):*

1. *Every type $\theta < \theta^*$ is fully revealed with probability 1. We refer to θ^* as the failure threshold.*
2. *Every type $\theta \geq \theta^*$ is fully revealed with probability $\alpha(\theta) \in [0, 1]$ and receives grade $z = t$ with probability $1 - \alpha(\theta)$, where $\alpha(\theta)$ is the smallest positive value satisfying*

$$\alpha(\theta)\theta + (1 - \alpha(\theta))t \geq \bar{x}(\theta), \tag{10}$$

and where $\bar{x}(\theta) = \frac{1}{\phi} [\bar{u}(\theta) - (1 - \phi)\theta]$ denotes the expected grade that makes type θ indifferent between silence and disclosure.

3. The failure threshold θ^* is smallest value satisfying the inequality:

$$\int_{\theta^*}^{\bar{\theta}} (\theta - \max\{t, \bar{x}(\theta)\}) dF(\theta) \geq 0. \quad (11)$$

Proposition 1 defines an optimal test for the regulator’s problem (8), which is illustrated in Figure 2. A useful point to keep in mind when interpreting our results is that the regulator is effectively trying to minimize the expected losses associated with financial frictions, which are given by $z\lambda(z)$ whenever a firm obtains a grade $z < t$.¹² We refer to any such grade as a “failing” grade, because it implies that the firm faces binding financial frictions after investors learn the test result.

First, we show that there is always an optimal test that fully reveals all firms with type $\theta < \theta^*$. These types fail the test with probability 1, which is why we refer to θ^* as the *failure threshold* in the rest of the paper. Below, we explain why the threshold property is optimal.

Second, we show that the optimal test assigns a binary lottery to the remaining types $\theta \geq \theta^*$. Each of these types receives either a passing grade $z = t$, which takes the firm exactly to the threshold that eliminates financial frictions, or a grade that fully reveals the true realization of θ . The intuition is as follows: One can partition the firm into “weak” types $\theta < t$, for whom financial frictions would bind under full disclosure, and “strong” types with $\theta \geq t$. An ideal scenario for the regulator would be a pooling scheme, in which she passes strong types with probability 1, and also assigns a passing grade to as many weak types as possible. However, this pooling scheme would also give some strong types an incentive to deviate towards verifiable disclosure. To avoid this issue, the regulator assigns such types a grade that passes “with flying colors” and reveals their true quality with some probability $\alpha(\theta)$. This probability is chosen as the smallest value that satisfies type θ ’s incentive compatibility constraint, which in this case reduces to Equation (10). As we elaborate below, this is the most efficient way to satisfy the incentive constraints. Notice that the regulator sets $\alpha(\theta) = 0$ for weak types $\theta \in [\theta^*, t]$ who pass the test. This is because these types can only benefit from the regulator’s pooling scheme, so they need not be given an additional incentive to stay quiet.

The third part of Proposition 1 defines how the failure threshold θ^* is chosen. The regulator would like θ^* to be as small as possible in order to minimize financial frictions.

¹²To see this more directly, notice that the objective function of the regulator in problem (8) can be written as

$$E_G[v(z)] = E_G[z(1 - \lambda(z))] = \mu_0 - E_G[z\lambda(z)],$$

where $\mu_0 = \int \theta dF(\theta)$ is the prior mean of θ , which is independent of the regulator’s chosen test G . Hence, the problem can be written equivalently as minimizing $E_G[\theta\lambda(z)]$, subject to the same constraints as in (8).

However, she is constrained by the requirement that the pass grade must credibly induce investors expectations z to lie above the threshold t at which financial constraints start to bind. Using Bayes' rule, investors' posterior probability on type θ upon observing the passing grade can be written (informally) as

$$p(\theta|\text{pass}) = \frac{(1 - \alpha(\theta)) dF(\theta)}{\int_{\theta^*}^{\bar{\theta}} (1 - \alpha(\theta')) dF(\theta')}$$

The credibility requirement is that the passing grade actually induces a posterior mean $z = \int \theta p(\theta|\text{pass}) d\theta \geq t$ or, equivalently,

$$\int_{\theta^*}^{\bar{\theta}} (\theta - t) (1 - \alpha(\theta)) dF(\theta) \geq 0 \tag{12}$$

Substituting the revelation probabilities $\alpha(\theta)$ implied by Equation (10) now also yields the inequality in Equation (11). The regulator optimally chooses the lowest threshold θ^* that is consistent with credibility. From this characterization, it also follows that the regulator can pass all types, choosing $\theta^* = \underline{\theta}$, if and only if investors' prior beliefs satisfy

$$E[\theta] \geq E[\max\{t, \bar{x}(\theta)\}]. \tag{13}$$

Intuitively, when investors are sufficiently optimistic a priori in the sense of Equation (13), the regulator can design a test which ensures that financial frictions are never binding. Conversely, if investors' prior belief $E[\theta]$ is below the right-hand side of (13), then the regulator is forced to choose $\theta^* > \underline{\theta}$ and faces binding financial frictions for a positive mass of weak types.

Notice that, whenever financial frictions bind, the regulator optimally to choose a failure threshold that satisfies the credibility constraint in Equation (12) with equality, and therefore induces a passing grade that satisfies $z = t$. Intuitively, it is always wasteful to assign a higher passing grade $z > t$ to weak types. Such a policy only transfers additional surplus to weak types without alleviating financial frictions. Therefore, as long as there are still weak types that face binding financial frictions, it is always better to pool more of these types into the passing grades until investors expectations are reduced to exactly $z = t$.

This argument clarifies why the regulator satisfies strong types' incentive constraints using the *stochastic* grades we described above, which randomize between full revelation and the passing grade $z = t$. An alternative design would be to assign strong types to higher passing grades $z > t$ with probability one, potentially pooling some weak types into the same grade. However, this policy again involves an inefficient transfer of surplus to the weak

types. It is more efficient to pool strong and weak types on the grade $z = t$, and assign higher grades randomly to strong types in order to satisfy their incentive constraints. Redistributing as much surplus as possible to the strong encourages private silence, introducing scope for public speech to protect more firms from financial frictions.

In Figure 2, panel (a), we illustrate these properties of optimal tests. In the figure, type $\theta_0 < \theta^*$ is an example of a failing type. The thick arrow from this type to itself denotes that type θ_0 obtains a grade equal to its true value (i.e., is fully revealed) with probability one. Type θ_1 is a passing weak type, who obtains the passing grade $z = t$ with probability one, again represented by a thick arrow. Type θ_2 is a strong type who does not have a binding incentive compatibility constraint, so that the regulator is able to set $\alpha(\theta_2) = 0$ and assign the passing grade with probability one. From Equation (10) and the figure, one can verify that setting $\alpha(\theta_2) = 0$ is feasible because $\bar{x}(\theta_2) < t$. Finally, type θ_3 is a strong type with a binding incentive compatibility constraint. The regulator assigns a stochastic grade, which is represented in the figure by thin arrows that are labeled with the associated probabilities. Type θ_3 is fully revealed with probability $\alpha(\theta_3) > 0$ and obtains the passing grade with the complementary probability.

Panel (b) of Figure 2 further illustrates the cumulative distribution $G(z)$ of grades that is induced by an optimal test with these properties. This distribution can be understood by considering four regions. First, all types below the failure threshold θ^* are fully revealed, so that the distribution of grades is identical to the prior distribution of types below this threshold. Second, there is bunching (i.e., a point mass) of types that are assigned the passing grade, implying an upward jump in $G(z)$ at the point $z = t$.¹³ Third, among types with a binding incentive constraint (i.e., $\theta > \bar{x}^{-1}(t)$), truthful revelation of grades that pass “with flying colors” occurs with probability $\alpha(\theta)$. Hence, the distribution of grades in this region is a “flatter” version of the prior distribution among these types.

As a final note on Proposition 1, we can use the intuition above to explain why a single failure threshold is optimal. Suppose that the regulator considers a marginal increase in the probability that a weak type $\theta < t$ achieves the passing grade t . The expected marginal cost of this change is that the passing grade becomes less credible. Indeed, we can write the marginal cost as $v(t - \theta) dF(\theta)$, where v stands for the shadow price associated with the credibility constraint (11). The marginal benefit, which arises because type θ obtains larger grades, is given by $(v(t) - v(\theta)) dF(\theta)$. The policy change increases the regulator’s

¹³Between the failure threshold and the passing grade t , the distribution of grades is flat, meaning that the regulator never reveals grades $z \in (\theta^*, t)$.

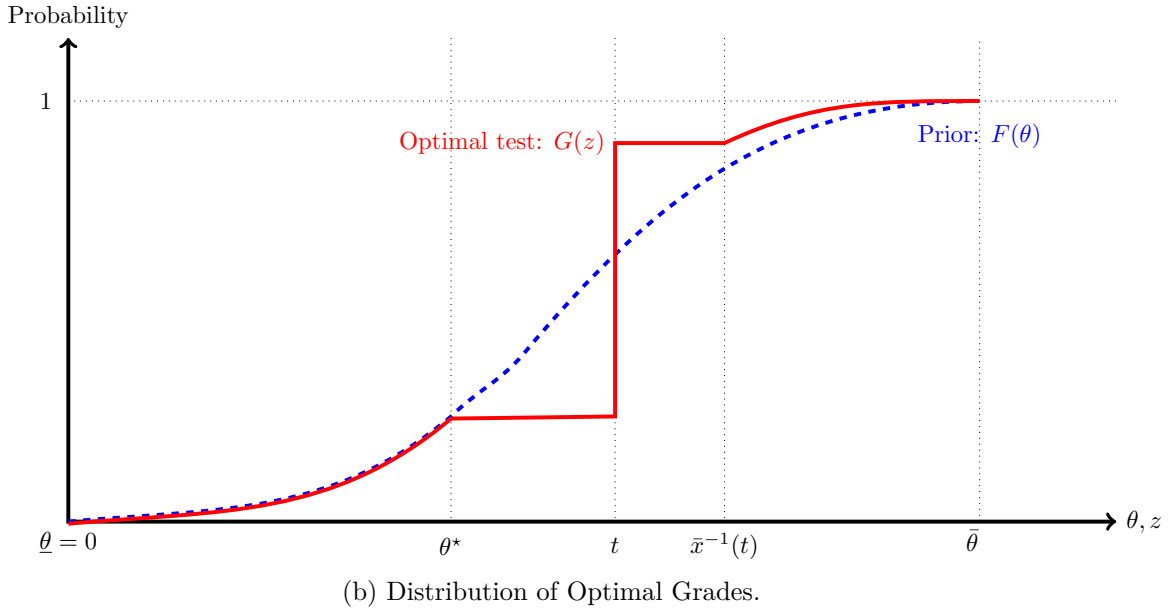
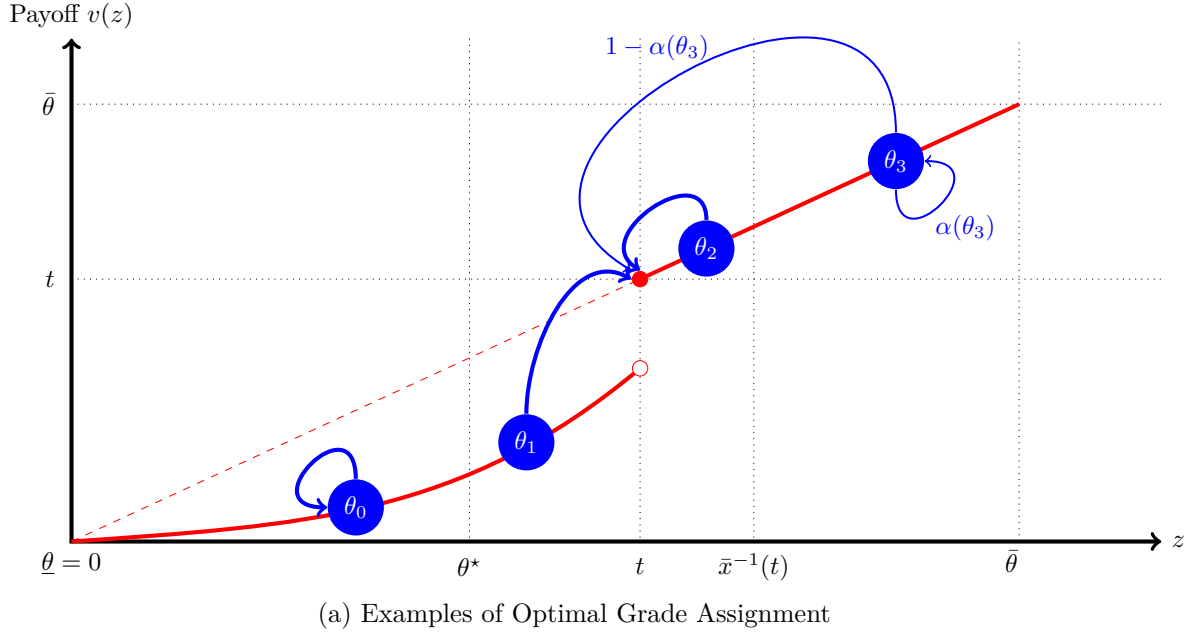


Figure 2: **Illustration of Optimal Tests.** Panel (a) plots the optimal assignment of types θ to grades z under the optimal test described in Proposition 1. Assignments are drawn for four example types $\theta \in \{\theta_0, \theta_1, \theta_2, \theta_3\}$, and the associated expected payoffs $v(z)$ are on the vertical axis. Thick arrows denotes a deterministic grade that is assigned with probability one to a given type. Thin arrows denotes a stochastic grade and are labeled with the associated probabilities. Panel (b) shows the implied distribution of optimal grades. The blue / dashed curve is the cumulative prior distribution of types θ . The red / solid curve is the cumulative distribution of grades $z = E[\theta|s]$ implied by the optimal test described in Proposition 1. The incentive constraint associated with private silence is binding for all types $\theta > \bar{x}^{-1}(t)$.

expected utility if marginal benefits exceed marginal costs, that is:

$$\frac{v(t) - v(\theta)}{t - \theta} \geq v.$$

The left-hand side of this inequality is the average gain obtained by promoting type θ to the passing grade t . Importantly, this is increasing in θ because $v(\cdot)$ is convex on the interval $[0, t]$. Therefore, the regulator always prefers to raise the probability of passing higher types until her credibility budget is used up, resulting in the single failure threshold θ^* .

3.4 Necessary Conditions for Optimal Tests

Proposition 1 fully describes an optimal test, but this implementation is not necessarily unique.¹⁴ In particular, there are potential indeterminacies due to the fact that the regulator's objective function is linear in investors' expectations z on some intervals. For example, different tests can induce the same distribution of surplus across passing types, $\theta \geq \theta^*$. Since the regulator's objective function $v(z)$ is linear for $z \geq t$, these alternatives are equally good for her — mean-preserving spreads / contractions among passing grades do not affect her expected utility so long as firms' incentives to stay quiet are preserved. Moreover, in applications such as Example 1 (bank runs – see Figure 1, panel (a)), the objective function is also linear for $z < t$, with a jump at $z = t$. In these cases, there are alternative optimal ways to assign failing grades to types $\theta < \theta^*$, for example, by pooling all firms below the threshold θ^* into a single failing grade instead of fully revealing them.

As a result, it becomes important to understand necessary conditions for optimality, so as to extract the economically important features that are shared by *all* optimal tests. The following result states two sets of necessary conditions.

Proposition 2. [Necessary Conditions for Optimal Tests] *All optimal tests G in the regulator's problem (8) satisfy the following conditions:*

1. *Every type $\theta < \theta^*$ obtains a grade $z < t$ with probability 1, and every type $\theta \geq \theta^*$ obtains a grade $z \geq t$ with probability 1, where the failure threshold θ^* is the same as in Proposition 1.*
2. *For almost every type θ of the firm such that $\bar{x}(\theta) > t$, the incentive compatibility constraint binds, and the firm achieves expected utility $\bar{u}(\theta)$.*
3. *For every type θ with a binding incentive compatibility constraint, the test result z is stochastic conditional on θ .*

¹⁴However, as we show in the appendix, the implied distribution of expected grades x is unique.

The first condition in Proposition 2 shows that all optimal tests share the stark separation around the failure threshold that we demonstrated in Proposition 1. All firm types below the threshold fail with probability one, and all types above pass with probability one, in any optimal test. Moreover, all optimal tests share the same threshold θ^* , which we characterized by Equation (11). Intuitively, the failure threshold captures the optimal trade-off between passing as many types as possible while also ensuring that the passing grade remains credible. The trade-off is not affected by any re-design of tests along linear regions of the regulator's objective and, therefore, we can find a unique optimal solution θ^* .

The second condition considers the incentive compatibility constraint in the regulator's problem, which ensures that no firm has an incentive to make additional disclosures. We show that the set of types θ for which this constraint is binding is the same at all optimal solutions. Indeed, exactly as in the particular solution that we define in Proposition 1, the constraint binds for all types such that $\bar{x}(\theta) > t$, who would have an incentive to deviate to disclosure if they were pooled to the passing grade $z = t$ with probability one.

An interesting third point is that, whenever the incentive constraint binds, the regulator assigns a stochastic grade z conditional on θ . This result contrasts with other models of information design with a continuum of types, in which optimal tests are often deterministic. The optimality of stochastic tests results from the nature of the regulator's trade-off. The regulator needs to provide strong types with high grades in order to encourage private silence, but also wants to pool them with weak types. As described in section 3.3, the most efficient way to manage this trade-off is to ensure weak types are pooled with stronger ones only on the grade $z = t$ that exactly avoids financial frictions. Randomization allows the regulator to balance these objectives, while deterministic tests do not.

In Appendix C, we derive an additional necessary condition for optimal tests. Recall that our results so far have relied on Lemma 1, which allows the regulator to focus on tests that induce private silence without loss of optimality. We further show that, even in the broader class of all possible tests, any optimal test *must* induce silence from all types $\theta \geq \theta^*$ above the failure threshold. In other words, the regulator's optimum cannot be achieved via a combination of a test and subsequent equilibrium disclosures. Private speech by the strong is *strictly* detrimental for the regulator, emphasizing the desirability of the *public speech and private silence* principle.

Our results in this subsection provide a tight characterization of optimal tests. Before moving on to applications of our general results to banking and corporate finance, we analyze some interesting comparative statics of optimal tests with respect to model primitives.

3.5 Comparative Statics

A useful observation is that the optimal test described in Proposition 1 depends only on two sets of parameters: The expected grades $\bar{x}(\theta)$ that make each type of firm indifferent between silence and disclosure, and the prior distribution of types $F(\theta)$. Knowledge of these two objects is enough to solve Equations (10) and (11), which determine the structure of the optimal test.

Interestingly optimal tests are thus independent of the nature of financial frictions, as measured by the function $\lambda(z)$: holding $\bar{x}(\theta)$ and $F(\theta)$ fixed, the same optimal test arises regardless of whether one studies bank runs in the tradition of Diamond and Dybvig (1983) (as in Example 1), or corporate financing frictions à la Holmstrom and Tirole (1997) (Example 2). Of course, the regulator's maximized utility does depend on frictions, since the total surplus lost to frictions is given by $\int_{\underline{\theta}}^{\theta^*} \theta \lambda(\theta) dF$. However, these differences do not affect optimal policy.

In order to understand the comparative statics of optimal tests, we therefore focus on variations in $\bar{x}(\theta)$ and $F(\theta)$. We begin with an analysis of the firm's disclosure-equivalent grade $\bar{x}(\theta)$:

Proposition 3. [*Comparative Statics: Disclosure Incentives*] *If $\bar{x}(\theta)$ increases for all θ , then:*

1. *The failure threshold θ^* in any optimal test increases and the regulator's maximized expected utility decreases.*
2. *The optimal test described in Proposition 1 becomes more informative: The distribution of optimal grades $z = E_G[\theta|s]$ undergoes a mean-preserving spread.*

Proposition 3 shows what happens when $\bar{x}(\theta)$ increases for all θ , so that higher expected grades are needed to discourage inside disclosures. The first point in the Proposition is straightforward: The regulator would like to enforce private silence, so that an increase in firms' incentives to make private disclosures is problematic. As a result, the regulator is forced to assign grades strictly above the passing threshold t to strong types with higher probability. In turn, the passing grade becomes less credible, and the regulator must raise the failure threshold θ^* and pass fewer weak firms. This change is clearly detrimental for the regulator's expected utility.

The second point in Proposition 3 is more subtle and concerns the informativeness of the optimal test that we described in Proposition 1.¹⁵ Whenever outside options increase,

¹⁵Due to the potential multiplicity of optimal tests, it is not possible to say that *any* optimal test for high $\bar{x}(\theta)$ is more informative than *any* optimal test for low $\bar{x}(\theta)$.

the grade $z = E_G[\theta|s]$ that the regulator releases as a result of her test becomes more variable in the sense of mean-preserving spreads (second order stochastic dominance).¹⁶ In our setting, any decision maker whose utility depends only the expectation of θ would prefer to observe the optimal test corresponding to the higher disclosure payoff function $\bar{x}(\theta)$. This phenomenon is the result of two economic forces. First, when the disclosure-equivalent grade $\bar{x}(\theta)$ increases, the failure threshold rises and the regulator fails a larger mass of weak types $\theta < \theta^*$ by fully revealing them. This change provides more information about weak firms to an outside observer. Second, the regulator assigns grades that pass strong types “with flying colors” more frequently so as to incentivize their silence. An outside observer therefore sees more accurate information about both weak and strong types.

Proposition 3 is also useful because it gives us a direct sense of the impact of verifiable disclosures on optimal tests. Indeed, the “unconstrained” case, in which firms cannot make verifiable disclosures at all, is equivalent to $\bar{x}(\theta) \leq 0$ for all θ in our model. Proposition 3 immediately implies that, compared to the constrained solution that we have analyzed, unconstrained-optimal tests are always less informative and fail a smaller mass of firms: the regulator engages in more elaborate public speech so as to maintain private silence.

To obtain further insights, recall from Proposition 1 that the disclosure-equivalent grade $\bar{x}(\theta)$ depends on i) the disclosure technology available to firms, via the the worst-case disclosure payoff $\bar{u}(\theta)$, and on ii) the preference parameter ϕ , which measures the sensitivity of firms’ payoffs to investors’ beliefs. The following auxiliary result explores the relationship between disclosure incentives and model primitives:

Lemma 4. [*Sources of Incentives to Disclose*] *The firm’s disclosure-equivalent grade $\bar{x}(\theta)$ increases for all θ and, therefore, optimal tests become more informative, if one of the following applies:*

- Lower disclosure costs: *The cost $c(m)$ of every verifiable message m decreases; or*
- More verifiable messages: *Holding constant the set $M = \cup_{\theta} M(\theta)$ of all possible messages, the set $M^{-1}(m)$ of types that can send message m becomes smaller (in the set inclusion order) for all m ; or*
- Firms are more sensitive to investors’ beliefs: *The parameter ϕ in the firm’s utility function in Equation (3) increases.*

¹⁶We note that this ranking is weaker than the Blackwell (1953) criterion, which would require a mean-preserving spread in the entire perceived distribution of $\theta|s$ as opposed to its mean. As the regulator cares only about z , Blackwell’s criterion is less useful here: a given distribution over z can often be implemented by more than one test, which need not be Blackwell-ranked.

Lemma 4 first demonstrates that firm’s incentives to disclose become stronger if disclosure becomes less costly, which is intuitive. The second point in the Lemma shows that incentives to disclose also become stronger if the messages available to the firm are “more verifiable”. The sense in which we define verifiability is the size of the set of types θ that have access to each message. For example, in a cheap talk setting, every type can send every message so that $M^{-1}(m)$ is as large as possible. By contrast, a scenario in which $M^{-1}(m)$ becomes smaller captures more verifiability. In this scenario, investors’ skeptical beliefs, which are used to derive the worst-case disclosure payoff $\bar{u}(\theta)$ in Equation (5), become more forgiving from the firm’s perspective. The firm still expects to be judged as the worst type that can send any given m , but that worst type (weakly) increases when messages become more verifiable. Hence, incentives to disclose become stronger.

Combining these insights with Proposition 3 shows that either a decrease in disclosure costs or an increase in verifiability make the regulator worse off, forcing her to fail more firms and to release more informative tests overall. In the limit as disclosure costs become prohibitive or all messages become cheap talk, the regulator is not constrained by inside disclosures at all. The other extreme, where fully verifiable disclosure is costless, is the least favorable for the regulator.

The third point in the Lemma states that incentives to disclose become stronger when the sensitivity of firms’ payoffs to investors’ beliefs is large. This property is also intuitive. Verifiable disclosures in our model are driven by firms’ desire to override outside grades that do not reflect their true quality. The more firms care about investors’ beliefs, the stronger their incentive to make verifiable disclosures.

As a complementary result, we also present comparative statics of optimal tests with respect to investors’ prior beliefs $F(\theta)$.

Proposition 4. [*Comparative Statics: Prior Beliefs*] *Assume that the disclosure-equivalent grade $\bar{x}(\theta)$ has a single crossing with t . If the prior distribution $F(\theta)$ becomes more optimistic in the sense of the Monotone Likelihood Ratio Property (MLRP), then the failure threshold θ^* in any optimal test decreases and the regulator’s maximized utility increases.*

In information design problems related to ours, it tends to be optimal to release less information when investors’ prior beliefs are optimistic (e.g., Bouvard et al., 2015; Goldstein and Leitner, 2018). In models without inside disclosures, this property arises because it becomes easier to make passing grades credible. Proposition 4 confirms this intuition, but the reasoning in our case is more nuanced. In principle, investors’ optimism can make the regulator’s credibility constraint in Equation (11) either easier or more difficult to fulfill. On one hand, investors place larger weight on strong types a priori, so that credibility is easier

to achieve. On the other hand, because the probability weight on strong types increases, it becomes more costly to provide these types with incentives to stay quiet.

Proposition 4 provides conditions under which the former effect dominates, so that the standard intuition prevails. First, we assume that the disclosure-equivalent grade $\bar{x}(\theta)$ has at most one crossing with the threshold t below which financial frictions bind. For example, consider the case where the messages available to the firm are $M(\theta) = \{\theta, \emptyset\}$, so that the firm can either disclose the true θ or stay quiet. If the costs of full disclosure are a constant c , then we have $\bar{x}(\theta) = \theta - \frac{1}{\phi}c$, which is linear, and therefore has the single-crossing property. More generally, single crossing holds in this message space as long as the costs of disclosure do not rise too fast with the true type θ . Second, we consider optimism in the sense of MLRP, which places enough structure on prior beliefs to ensure that the integral defining the credibility constraint in Equation (11) increases with optimism. By contrast, without inside disclosure, a weaker notion of optimism such as first-order stochastic dominance would have been sufficient to make the regulator better off.

3.6 Private Silence as the Only Motivation for Public Speech: A Special Case

We close this section by highlighting a special case of the characterizations we have developed:

Corollary 1. *[Private Silence as the Only Motivation for Public Speech] Suppose that the investor’s prior beliefs satisfy*

$$t < E[\theta] < E[\max\{t, \bar{x}(\theta)\}]. \quad (14)$$

If firms cannot make verifiable disclosures, then the regulator’s optimal test conveys no information to investors. By contrast, if firms can make verifiable disclosures, then the regulator optimally releases an informative test and fails a strictly positive mass of firm types, as described in Proposition 1.

Corollary 1 considers parameter constellations where investors’ prior expectation $E[\theta]$ is in an intermediate range, as defined in Equation (14). In this case, a regulator who fully controls the informational environment does not need to provide any information to investors at all. Indeed, since $E[\theta] > t$, financial frictions never bind if investors judge all firms according to their prior expectation. The regulator can do no worse than to preserve this situation by sending a single, uninformative grade $z_0 \equiv E[\theta]$ as her test result. However, if firms are now given the opportunity to make verifiable inside disclosures, the regulator’s opaque test “unravels”. A set of strong firms have an incentive to deviate to disclosure, which means

that they are no longer pooled into the regulator’s test result z_0 . Hence, investors’ posterior belief given z_0 become more pessimistic, providing an incentive for further strong firms to disclose. This process makes the test result gradually less credible, and financial frictions will eventually bind for any firm that does not make a verifiable disclosure.¹⁷ An optimal response to this problem is for the regulator to release some information that encourages private silence. This example makes the key point of our paper quite starkly, since enforcing private silence is now the only motivation for public speech.

4 Applications

4.1 Bank Stress Tests in Times of Financial Crises

In Examples 1 and 3 above, we demonstrated that our framework can be applied to models of banking crises. In this context, our characterization of optimal tests highlights the preferred policy of a regulator who can release credible information to investors, and who is concerned about market expectations dropping below a threshold and triggering bank runs. We view this setting as informative for the problem of real-world financial regulators who have to design stress tests in times of financial crises.

For example, during the financial crisis of 2007-9, the Federal Reserve designed a novel stress testing exercise as part of its Supervisory Capital Assessment Program (SCAP). The Fed’s stated rationale for announcing this test was to provide credible information that would resolve uncertainty about potential losses in the US banking system.¹⁸ All large US bank holding companies were required to take part. For each of the 19 participating banks, the Fed calculated projected losses in hypothetical, adverse macroeconomic scenarios. The Fed also released granular breakdowns of these losses for the “more adverse” scenario in its test to the public.¹⁹ A number of studies have confirmed that the SCAP conveyed credible information to investors, which manifested itself in bank stock returns and trading volumes (Morgan et al., 2014; Flannery et al., 2017), CDS spreads and systematic risk (Sahin et al., 2020), as well as indicators of information asymmetry and uncertainty (Fernandes et al., 2020).²⁰

Since the financial crisis, regulators have also adopted a more regular stress testing regime,

¹⁷When disclosure costs are strictly positive, this process does not necessarily result in *all* types $\theta \geq t$ making a disclosure. However, it remains true that non-disclosing firms will face an effective grade $z < t$ with probability 1.

¹⁸See <https://www.federalreserve.gov/newsevents/pressreleases/files/bcreg20090424a1.pdf> for a description of the framework and its rationale, which was published before the results were announced.

¹⁹For example, see Table 3 in <https://www.federalreserve.gov/newsevents/pressreleases/files/bcreg20090507a1.pdf>.

²⁰Petrella and Resti (2013) and Ahnert et al. (2018) present analogous results for the European stress tests in 2011.

which releases information about banks’ health in normal times. The stated goal of these tests is different, and focuses more on their ability to reinforce transparency and to mitigate moral hazard in the financial system. While our analysis can shed some light on the optimal design of tests in normal times (recall Corollary 1), we do not formally model moral hazard or market discipline. Hence, we view our results as primarily useful for the design of stress tests *in times of financial distress*, such as the SCAP and its international equivalents, which were driven mostly by the desire to avoid market distress.

In this subsection, we discuss three novel insights on bank stress tests that arise from our analysis.

Optimal granular stress tests: A unique feature of optimal tests in our model is that the regulator releases three types of grades: Failing grades $z < t$ lead to binding financial frictions, a passing grade $z = t$ pools weak and strong banks, and a range of grades $z > t$ pass strong firms “with flying colors” in order to discourage inside disclosures. Thus, optimal tests are quite granular. As long as any of the incentive compatibility constraints associated with inside disclosure are binding, any optimal test induces more than one passing grade. In addition, if there is strict convexity in the regulator’s objective function $v(z)$ when frictions bind at $z < t$ (e.g., in Example 3), then any optimal test also involves multiple failing grades.

Interestingly, the granularity of the optimal test in our model is somewhat consistent with how the 2009 SCAP was implemented in practice. Indeed, the Fed published detailed results about each bank’s losses in adverse scenarios. This included losses among banks such as Goldman Sachs and JP Morgan, who “passed” the test, in the sense that the Fed did not conclude that they had a capital shortfall. Similarly granular results were published by the European Banking Authority in 2011, when the European banking crisis was still ongoing.²¹ Our model provides a possible rationale for these practices, if one of the motives for public disclosure of stress test results is to prevent the “unraveling” of information via signals from strong banks. Shahhosseini (2016) provides some evidence consistent with the effects we highlight. She studies banks’ reactions to being included in stress tests, and highlights that stress-tested banks made fewer loan charge-offs and more frequently changed the classification of loan losses, which one could interpret as tested banks becoming more opaque.

Our results on the granularity of optimal tests are new to the existing literature on stress test design, which shows that simple pass/fail tests are optimal in many settings. This in-

²¹See, for example, <https://www.eba.europa.eu/risk-analysis-and-data/eu-wide-stress-testing/2011>, in particular the document entitled “*Summary of CT1 ratios under the adverse scenario.pdf*”.

cludes Faria-e-Castro et al. (2016), who study the interaction of stress tests with government bailouts, Inostroza and Pavan (2022), where a regulator can communicate privately with dispersed investors, and Orlov et al. (2021), who advocate for an initial binary test followed by capital injections and ultimately full disclosure. Inostroza (2019) emphasizes the merits of multiple failing grades, as in our Example 3, but in his model there is no rationale for granularity in passing grades. In Goldstein and Leitner (2018), strong firms have the outside option of retaining an asset rather than selling it to investors. However, the associated incentive constraints are specific to the nature of the outside option. Our model with endogenous disclosure generates outside options that are starker and more tractable. As a result, we can derive a full characterization and comparative statics of optimal tests in a model with an arbitrary distribution of firm types and signals. In this context, particularly novel features of our analysis in this context are that optimal grades involve strict separation around the failure threshold θ^* , and assign stochastic test results to all types with a binding incentive compatibility constraint. Our work also connects to Philippon and Skreta (2012), who show in a very different environment that regulatory policy in a crisis should take into account its effect on endogenous information and adverse selection in financial markets.

Determinants of Regulatory Trade-offs in Stress Test Design: Our analysis of comparative statics in Section 3.5 sheds light on the institutional details that affect the optimal design of bank stress tests. In common with existing literature, we show that regulators must release more accurate information, usually at the cost of failing more banks, when investors’ beliefs become more pessimistic in a financial crisis. In addition, we show that it is not enough to reassure investors in isolation. Instead, regulators need to consider the endogenous response of banks’ disclosures, make stress tests informative enough to prevent “unraveling” of information, and encourage a regime with *public speech and private silence*.

Our analysis of banks’ incentives and disclosure technologies in Proposition 3 and Lemma 4 highlights additional features that regulatory designs must take into account. First, the required informativeness of stress tests is increasing in banks’ ability to make verifiable disclosures. A sizable applied literature in accounting studies the extent to which banks’ disclosures have the potential to credibly reassure investors. An emerging consensus is that banks’ disclosures about opaque areas of their business model, in particular about derivatives and securitization activities, provide particularly valuable information to investors during crises (e.g., Barth and Landsman, 2010; Iren et al., 2014). By contrast, standard fair value accounting in normal times is considered less important. Sowerbutts et al. (2013) also document that banks made heightened disclosures during the crisis of 2008. Our analysis implies that the more detailed banks’ potential verifiable disclosures become, the more regulators should be

concerned about “unraveling” effects, and respond accordingly with more informative stress tests and a greater emphasis on multiple passing grades that incentivize private silence.

Moreover, Lemma 4 highlights that banks’ incentives to disclose may be driven by the incentives of their managers and, in particular, by the sensitivity of managers’ payoffs to investors’ expectations. This finding is particularly salient in banking, because senior bankers tend to have bonus-based compensation schemes. These patterns have been documented both by theoretical analyses of optimal compensation in banking (e.g., Thanassoulis, 2012; Acharya et al., 2016) and by the associated empirical evidence (e.g., Bannier et al., 2013; Célérier and Vallée, 2019). All else equal, our results suggest that regulators should be more concerned about inside disclosures when managers’ payoffs are more sensitive to current market expectations, for example because managers own their bank’s stock / options or because their bonuses are sensitive to annual stock price performance.

Perhaps surprisingly, once these institutional features have been taken into account, the optimal test described in Proposition 1 does not depend on the details of a particular model of financing frictions that regulators have in mind. Indeed, the optimal test has the same structure, and the same failure threshold, when financial frictions take the shape of a binary “run / don’t run” outcome (as in our Example 1), or when they are manifested in a smooth, convex function that captures the gradual deterioration of market conditions as investors become more pessimistic (as in Example 3). Therefore, our analysis shows that tests need not be “fine-tuned” to particular financial frictions. Rather, optimal designs depend crucially on market beliefs and insiders’ incentives.

A final point worth highlighting is our recommendation for stress testing at the “eve of a crisis”, when investors’ prior beliefs are relatively pessimistic but lie in the intermediate region characterized in Corollary 1. In these scenarios, there is no need for stress tests at all unless regulators expect inside disclosures by banks. Running a stress test in these situations is still beneficial because it ensures that information does not unravel, and presents information to investors in a way that minimizes the probability of future bank runs.

Credibility and Regulatory Commitments to Stress Test Accuracy: Several studies of stress tests have highlighted potential time inconsistencies in regulator’s behavior. For example, Bouvard et al. (2015) show that regulators may have an incentive to withhold private information about bank vulnerability ex post, in an attempt to signal strength to investors, while it would be optimal ex ante to commit to an informative stress test.

In our framework, we can address the question whether a regulator would change the nature of the optimal test after learning the true quality θ of banks’ assets. Concretely, notice that the regulators’ and the firms’ preferences, defined in Equations (2) and (3), are

exactly the same when we set the sensitivity parameter ϕ to zero. Hence, analyzing stress tests that induce private silence when $\phi = 0$ is equivalent to analyzing tests that would induce the *regulator* to refrain from further disclosures, and to obey her ex-ante information design, after she learns the realization of θ .

Using our characterization of optimal silence-inducing tests in Proposition 1, we can infer that incentive constraints associated with private silence are never binding in the case where $\phi = 0$. The economic intuition is as follows: Incentive constraints in our model when strong banks have an incentive to signal their quality to investors, whose beliefs drive bank managers' payoffs whenever $\phi > 0$. By contrast, under the regulators' preferences with $\phi = 0$, there is no gain from signaling quality as long as one achieves a grade $z \geq t$ that avoids financial frictions. For this reason, the optimal tests we characterize are time-consistent in the sense that the regulator would not want to make additional disclosures after learning the true type θ .

Two caveats to this point remain in the light of the recent literature. First, as well as time-inconsistency in stress tests during crises, there can be time-inconsistency in normal times, as pointed out by Parlasca (2021). Indeed, if one of the objectives of stress tests in normal times is to alleviate moral hazard, then the regulators have an incentive to threaten tough tests ex ante, but relax these tests ex post if there is a possibility of binding financial frictions. Second, in the model considered by Orlov et al. (2021), it is sometimes optimal to fail strong types of bank in order to make it easier for failing banks to raise capital after the test is released. This type of mechanism may not be time-consistent if the regulator prefers to pass strong banks instead after their strength has been realized.

4.2 Optimal Credit Ratings

In Example 2, we demonstrated that our setting nests a standard model of corporate financing in the spirit of Holmstrom and Tirole (1997). Financial frictions in this model arise when the firm's borrowing capacity becomes insufficient to make all available investments that have positive NPV. Releases of information about asset quality in this setting can be interpreted as credit ratings, whereby favorable ratings allow firms to borrow more and ensure that financial constraints do not bind. The regulator can be interpreted as an idealized ratings agency that maximizes the joint utility of investors and the firm.

This interpretation allows us to contribute insights to a growing literature on the design of credit ratings. Motivated by the crisis of 2008, this literature focuses to a large extent on imperfections in how credit ratings are produced and consumed. These issues include firms' incentives to shop for favorable ratings (Skreta and Veldkamp, 2009; Bolton et al., 2012;

Kashyap and Kovrijnykh, 2016), the interaction of ratings with financial regulations (Opp et al., 2013), and the incentives for firms to block the publication of unfavorable ratings Sangiorgi and Spatt (2017). Daley et al. (2020), in perhaps the closest paper to ours in this literature, show how ratings affect firms’ incentives to signal their quality via asset retention, with further implication for screening and credit supply.

All of the studies cited above consider a simple, usually exogenous structure for the informative signals that are conveyed by credit ratings. In particular, the standard approach is to assume that the firm has a binary type and investors observe a binary distributed signal of this type.²² These assumptions are geared towards conveying intuition about various incentive problems in the clearest possible way.

Our model contributes a different perspective to this literature. Applying our general results to Example 2, we can analyze how credit ratings should be designed *in principle* when there are many possible types of firm, who can also make additional corporate disclosures to increase their borrowing capacity. Once again, the problem of designing an optimal rating boils down to ensuring public speech and private silence. This strategy, which crowds out all voluntary additional disclosures, ensures that the overall impact of financial frictions is minimized, and also economizes on any deadweight costs associated with corporate disclosure.

Unlike existing studies, we show how a wide spectrum of optimal ratings emerges endogenously from the objective to minimize financial frictions. First, there is a range of low “junk” ratings, whose recipients end up financially constrained with probability one. Second, there is significant bunching of firms around an “investment grade” rating, at which financing constraints become slack. Finally, there is a range of “premium” ratings, which deliver additional surplus to the firms who receive them. This array of ratings resembles the wide range of grades that the major credit ratings agencies deliver in practice. We would not expect credit ratings in the data to satisfy our optimality conditions exactly, since the ratings industry is subject to the various frictions and agency problems mentioned above. However, we view this result as an interesting and novel benchmark for ratings design.

5 Conclusion

We have characterized the optimal informational policy of a regulator whose goal is to alleviate financial frictions. Our results provide a characterization of optimal tests that is robust to an informational version of the “Lucas critique”: One should not assume that

²²A notable exception is Manso (2013), who models a stochastic process of more granular ratings but still takes the structure of ratings as given.

regulators have full control over the informational environment in which financial markets operate, and should expect that any change in informational policy will trigger further changes in the endogenous production of information.

We demonstrate that optimal tests take a markedly different shape once we account for this feature. We show that regulators should aim for public speech and private silence, refraining from private communication of test results to firms, and preserving firms' incentives not to communicate with investors. As a result, optimal tests become more granular, more informative and involve novel types of grades that pass firms "with flying colors" in order to preserve their incentives to stay quiet. Moreover, they should maintain some uncertainty about the likely grade firms will receive.

We further illustrated the applied insights that arise for the design of bank stress tests during financial crises, and for the design of an idealized credit ratings agency that maximizes social surplus. Beyond the applications covered in this paper, we hope that our approach and methodology will spark further work on constrained information design, which could consider additional applications in which the information designer does not have full control of the informational environment.

References

- Acharya, V., M. Pagano, and P. Volpin (2016). Seeking alpha: Excess risk taking and competition for managerial talent. *The Review of Financial Studies* 29(10), 2565–2599.
- Acharya, V. V., P. DeMarzo, and I. Kremer (2011). Endogenous information flows and the clustering of announcements. *American Economic Review* 101(7), 2955–2979.
- Ahnert, L., P. Vogt, V. Vonhoff, and F. Weigert (2018). The impact of regulatory stress testing on bank's equity and cds performance.
- Bannier, C. E., E. Feess, and N. Packham (2013). Competition, bonuses, and risk-taking in the banking industry. *Review of Finance* 17(2), 653–690.
- Barth, M. E. and W. R. Landsman (2010). How did financial reporting contribute to the financial crisis? *European accounting review* 19(3), 399–423.
- Bertomeu, J., E. Cheynel, and D. Cianciaruso (2021). Strategic withholding and imprecision in asset measurement. *Journal of Accounting Research* 59(5), 1523–1571.
- Best, J. and D. Quigley (2020). Persuasion for the long run. *Available at SSRN 2908115*.

- Billingsley, P. (1995). *Probability and measure* (3. ed ed.). John Wiley & Sons.
- Blackwell, D. (1953). Equivalent comparison of experiments. *Annals of Mathematical Statistics* 24(2), 265–272.
- Bolton, P., X. Freixas, and J. Shapiro (2012). The credit ratings game. *The Journal of Finance* 67(1), 85–111.
- Bond, P. and Y. Zeng (2021). Silence is safest: information disclosure when the audiences preferences are uncertain. *Journal of Financial Economics*.
- Bouvard, M., P. Chaigneau, and A. de Motta (2015). Transparency in the financial system: Rollover risk and crises. *Journal of Finance* 70(4), 1805–1837.
- C el erier, C. and B. Vall ee (2019). Returns to talent and the finance wage premium. *The Review of Financial Studies* 32(10), 4005–4040.
- Daley, B. and B. Green (2014). Market signaling with grades. *Journal of Economic Theory* 151, 114–145.
- Daley, B., B. Green, and V. Vanasco (2020). Securitization, ratings, and credit supply. *The Journal of Finance* 75(2), 1037–1082.
- DeMarzo, P. M., I. Kremer, and A. Skrzypacz (2019, June). Test design and minimum standards. *American Economic Review* 109(6), 2173–2207.
- Diamond, D. W. and P. H. Dybvig (1983). Bank runs, deposit insurance, and liquidity. *Journal of Political Economy* 91(3), 401–19.
- Dworzak, P. and G. Martini (2019). The simple economics of optimal persuasion. *Journal of Political Economy* 127(5), 1993–2048.
- Einhorn, E. (2018). Competing information sources. *The Accounting Review* 93(4), 151–176.
- Faria-e-Castro, M., J. Martinez, and T. Philippon (2016). Runs versus lemons: information disclosure and fiscal capacity. *The Review of Economic Studies*.
- Feltovich, N., R. Harbaugh, and T. To (2002). Too cool for school? Signalling and countersignalling. *RAND Journal of Economics* 33(4), 630–649.
- Fernandes, M., D. Igan, and M. Pinheiro (2020). March madness in wall street:(what) does the market learn from stress tests? *Journal of Banking & Finance* 112, 105250.

- Flannery, M., B. Hirtle, and A. Kovner (2017). Evaluating the information in the federal reserve stress tests. *Journal of Financial Intermediation* 29, 1–18.
- Frenkel, S., I. Guttman, and I. Kremer (2018). The effect of exogenous information on voluntary disclosure and market quality.
- Friedman, H., J. Hughes, and B. Michaeli (2015). Bayesian persuasion in the presence of discretionary disclosure. *Mimeo*.
- Gentzkow, M. and E. Kamenica (2017). Competition in persuasion. *The Review of Economic Studies* 84(1), 300–322.
- Gigler, F. and T. Hemmer (1998). On the frequency, quality, and informational role of mandatory financial reports. *Journal of Accounting Research* (36), 117–147.
- Goldstein, I. and Y. Leitner (2018). Stress tests and information disclosure. *Journal of Economic Theory* 177, 34–69.
- Goldstein, I. and A. Pauzner (2005). Demand–deposit contracts and the probability of bank runs. *Journal of Finance* 60(3), 1293–1327.
- Goldstein, I. and L. Yang (2017). Information disclosure in financial markets. *Annual Review of Financial Economics* 9, 101–125.
- Grossman, S. J. and O. D. Hart (1980). Disclosure laws and takeover bids. *Journal of Finance* 35(2), 323–334.
- Guo, Y. and E. Shmaya (2019). The interval structure of optimal disclosure. *Econometrica* 87(2), 653–675.
- Holmstrom, B. and J. Tirole (1997). Financial intermediation, loanable funds, and the real sector. *the Quarterly Journal of economics* 112(3), 663–691.
- Inostroza, N. (2019). Persuading multiple audiences: An information design approach to banking regulation. *Available at SSRN 3450981*.
- Inostroza, N. and A. Pavan (2022). Adversarial coordination and public information design. *Mimeo*.
- Iren, P., A. K. Reichert, and D. Gramlich (2014). Information disclosure, bank performance and bank stability. *International Journal of Banking, Accounting and Finance* 5(4), 374–417.

- Kamenica, E. and M. Gentzkow (2011). Bayesian persuasion. *American Economic Review* 101(6), 2590–2615.
- Kashyap, A. K. and N. Kovrijnykh (2016). Who should pay for credit ratings and how? *The Review of Financial Studies* 29(2), 420–456.
- Li, F. and P. Norman (2021). Sequential persuasion. *Theoretical Economics* 16(2), 639–675.
- Lyu, C. (2021). Information design for selling search goods and the effect of competition. *Working Paper*,.
- Manso, G. (2013). Feedback effects of credit ratings. *Journal of Financial Economics* 109(2), 535–548.
- Milgrom, P. R. (1981). Good news and bad news: Representation theorems and applications. *Bell Journal of Economics* 12(2), 380–391.
- Morgan, D. P., S. Peristiani, and V. Savino (2014). The information value of the stress test. *Journal of Money, Credit and Banking* 46(7), 1479–1500.
- Morris, S. and H. S. Shin (2000). Rethinking multiple equilibria in macroeconomic modeling. *NBER Macroeconomics Annual* 15, 139–182.
- Opp, C. C., M. M. Opp, and M. Harris (2013). Rating agencies in the face of regulation. *Journal of financial Economics* 108(1), 46–61.
- Orlov, D., P. Zryumov, and A. Skrzypacz (2021). Design of macro-prudential stress tests.
- Parlasca, M. (2021). Time inconsistency in stress test design. *Available at SSRN 3803876*.
- Petrella, G. and A. Resti (2013). Supervisors as information producers: Do stress tests reduce bank opaqueness? *Journal of Banking & Finance* 37(12), 5406–5420.
- Philippon, T. and V. Skreta (2012). Optimal interventions in markets with adverse selection. *American Economic Review* 102(1), 1–28.
- Rothschild, M. and J. E. Stiglitz (1970). Increasing risk: I. A definition. *Journal of Economic Theory* 2(3), 225–243.
- Sahin, C., J. de Haan, and E. Neretina (2020). Banking stress test effects on returns and risks. *Journal of Banking & Finance* 117, 105843.

- Sangiorgi, F. and C. Spatt (2017). Opacity, credit rating shopping, and bias. *Management Science* 63(12), 4016–4036.
- Shahhosseini, M. (2016). The unintended consequences of bank stress tests. *Mimeo*.
- Shin, H. S. (2003). Disclosures and asset returns. *Econometrica* 71(1), 105–133.
- Skreta, V. and L. Veldkamp (2009). Ratings shopping and asset complexity: A theory of ratings inflation. *Journal of Monetary Economics* 56(5), 678–695.
- Sowerbutts, R., P. Zimmerman, and I. Zer (2013). Banks disclosure and financial stability. *Bank of England Quarterly Bulletin*, Q4.
- Thanassoulis, J. (2012). The case for intervening in bankers pay. *The Journal of Finance* 67(3), 849–895.
- Verrecchia, R. E. (1983). Discretionary disclosure. *Journal of Accounting and Economics* 5, 179–194.

A Proofs of Main Results

A.1 Preliminaries

In this section we introduce some auxiliary concepts and solution methods used in our main arguments. The proofs of results presented in this section are provided in Appendix B.

With a slight abuse of notation, for any test/signal structure $\{G(s | \theta)\}_{s \in S, \theta \in \Theta}$, we write $G(z|\theta)$ for the associated conditional distribution of grades $z = E[\theta|s]$ when the firm does not make inside disclosures. We similarly write $G(z) = \int G(z|\theta) dF(\theta)$ for the *unconditional* distribution of grades induced by a test G .

A.1.1 Threshold-separable tests

We introduce a class of tests important to our main analysis, characterize them, and provide a preliminary result on which our main analysis draws.

Definition 1. A test is *threshold-separable* with threshold $\theta' \in [\underline{\theta}, \bar{\theta}]$ if for any pair $\theta_1, \theta_2 \in [\underline{\theta}, \bar{\theta}]$ such that $\theta_1 < \theta' \leq \theta_2$,

$$\text{supp } G(\cdot | \theta_1) \cap \text{supp } G(\cdot | \theta_2) = \emptyset.$$

We derive an equivalent definition of threshold-separable tests as follows:

Lemma 5. *A test is threshold-separable with threshold θ' if and only if*

$$\int_{\underline{\theta}}^r G(z) dz \leq \int_{\underline{\theta}}^r F(z) dz \tag{15}$$

for all $r \in [\underline{\theta}, \bar{\theta}]$, with equality at $r = \theta', \bar{\theta}$. Moreover, any test G satisfying (15) must also satisfy $G(\theta') = F(\theta')$.

We can also show that one can focus on threshold-separable tests in the regulator's Problem 8:

Lemma 6. *For any solution $\{G(z | \theta)\}$ to Problem (8), there exists an alternative threshold-separable solution $\{G''(z | \theta)\}$ with threshold $\theta' \leq t$. Moreover, G'' induces a distribution over grades $z \geq t$ if $\theta \geq \theta'$, and otherwise reveals θ .*

Lemma 6 is not sufficient to pin down optimal tests: it says little about the distribution of z on either side of the threshold. However, in our main analysis, we leverage it to draw yet stronger conclusions about the form of *every* optimal test.

A.1.2 Solution Methods for Relaxed Problem

We now develop solution methods for the regulator's relaxed problem (9), which involves fewer constraints than her main problem (8). As Lemma 3 establishes, the solution to (9) indeed yields a weakly higher value than (8). In this sense, the latter is a relaxed version of the former.

Yet because it involves both first and second order stochastic dominance constraints, Problem (9) itself is not amenable to analysis by the standard tools of Bayesian persuasion. Nonetheless, the next lemma shows that the methods of Dworzak and Martini (2019) can be readily extended to deal with the additional FOSD constraints.²³

Lemma 7. *If there exist a convex function $\gamma : [\underline{\theta}, \bar{\theta}] \rightarrow \mathbb{R}$, an decreasing function $\psi : [\underline{\theta}, \bar{\theta}] \rightarrow \mathbb{R}$ and a distribution $H^* \preceq_{co} F$ such that:*

1. $\gamma(x) + \psi(x) \geq v(x)$ for all $x \in [\underline{\theta}, \bar{\theta}]$, with equality for $x \in \text{supp } H^*$
2. $\int \gamma(x) dH^* = \int \gamma(x) dF$
3. $\int \psi(x) dH^* = \int \psi(x) dF \circ \delta$

and H^* first-order stochastically dominates $F \circ \delta$, then H^* solves (8). Moreover, any solution H' to (8) must satisfy 1.-3. for the same pair (γ, ψ) .

Making use of Lemma 7, we now describe the properties of any optimal distribution H^* over the firm's interim expected grades. We use these findings to prove our main results, Propositions 1 and 2, which translate these findings into implications for optimal tests. The next result establishes that an optimal solution to Problem (9) exists:

Lemma 8. *The distribution*

$$H^*(x) = \begin{cases} F(x), & x < \theta^* \\ F(\theta^*), & \theta^* \leq x < t \\ F \circ \delta(x), & t \leq x, \end{cases}$$

where θ^* is chosen to satisfy:

$$\int_{\underline{\theta}}^{\theta^*} x dF(x) + t(F(b) - F(\theta^*)) + \int_t^{\bar{\theta}} x dF \circ \delta(x) = \int \theta dF, \quad (16)$$

²³In a search context without any equivalent of our disclosure problem, Lyu (2021) uses a similar extension. In our setting, the presence of disclosure constraints means that reframing the problem in terms of higher-order beliefs, x , is necessary for identifying this kind of relaxed problem. Indeed, we find that optimal tests are meaningfully stochastic, where most persuasion problems with a continuous state variable find optimal signals which are deterministic.

solves problem (9).

H^* describes a distribution over the firm's conditional expected payoffs which solves Problem (9). We emphasize that θ^* refers to a particular optimal threshold (i.e. that consistent with H^*), where by contrast θ' introduced in Definition 1 denotes an arbitrary (not necessarily optimal) threshold. Of course, H^* is not a solution to (8). In particular, note that Lemma 8 does not identify whether a test G that induces this distribution of private beliefs exists, nor does it describe what such a G looks like. We establish the existence of such a test in the proof of Proposition 1 (see Appendix A.2).

Nevertheless, Lemma 7 allows us to identify some necessary properties of any solution to Problem (9). These are important to the proof of Proposition 2:

Proposition 5. [*Properties of any solution H' to (9)*]. *Up to a set of measure zero:*

1. $\text{supp } H' \subset [\underline{\theta}, \theta^*] \cup [t, \bar{\theta}]$,
2. $H'(x) = F \circ \delta(x)$ for $x \geq t$,
3. $x \geq t$ if and only if $\theta \geq \theta^*$. If v is strictly convex on $[\underline{\theta}, \theta^*]$ then $H'(\theta) = F(\theta)$ for all $x \leq \theta^*$,

where θ^* is unique and identified by Lemma 8.

In particular, we find that all optimal tests must involve (i) the same set of types enjoying an expected grade greater than t (those above the threshold θ^* identified in Lemma 8), and (ii) the FOSD constraint binds everywhere above t .

For purposes of reference, we end this section with a brief recount of the (well-known) solution to the benchmark design problem in which the firm cannot disclose its type:

$$\begin{aligned} V^{FB} &= \max_G \int v(z) dG(z) \\ \text{subject to} \quad & G \preceq_{co} F \end{aligned} \tag{17}$$

Corollary 2. *Suppose $\mathbb{E}[\theta] < t$. Any solution to (17) is characterized by a threshold θ^{FB} in which, with probability 1, types $\theta \geq \theta^{FB}$ (and only these types) are pooled into a single grade $z = t$, where θ^{FB} satisfies*

$$\mathbb{E}[\theta \mid \theta \geq \theta^{FB}] = t.$$

For every type $\theta < \theta^{FB}$, $\Pr[\tilde{z} \geq t \mid \theta] = 1 - \lim_{z \uparrow t} G(z \mid \theta) = 0$. Moreover, full revelation of such types is weakly optimal.

A.2 Proofs

Proof of Lemmas 1 and 2

Fix a test $\{G(s | \theta)\}_{s \in S, \theta \in \Theta}$ and some corresponding equilibrium disclosure strategies $\mu : \Theta \rightarrow \Delta M$, where $\mu(M(\theta)^c | \theta) = 0$.²⁴ The pair (G, μ) induces a distribution \hat{G} over posterior means $z = \mathbb{E}[\theta | s, m]$, defined by

$$\begin{aligned}\hat{G}(z | \theta) &= \int_{\underline{\theta}}^{\bar{\theta}} \mu(M_s(z) | \theta) dG(s | \theta) \\ \hat{G}(z) &= \int_{\underline{\theta}}^{\bar{\theta}} \hat{G}(z | \theta) dF(\theta),\end{aligned}\tag{18}$$

where $M_s(z) := \{m \in M : \mathbb{E}[\theta | s, m] \leq z\}$. We note that, as μ is an equilibrium strategy, it satisfies

$$m \in \text{supp } \mu(\cdot | \theta) \implies m \in \arg \max_{m \in M(\theta)} \int u(\mathbb{E}[\theta | s, m], \theta) dG(s | \theta) - c(m)\tag{19}$$

for all $\theta \in [\underline{\theta}, \bar{\theta}]$. Let the maximized value in (19) be $U^e(\theta)$.

We show that the test $\{\hat{G}(z | \theta)\}_{(z, \theta) \in [\underline{\theta}, \bar{\theta}]^2}$ and disclosure strategy $\hat{\mu}(\{m = \emptyset\} | \theta) = 1$, for all θ , is an equilibrium, supported by the skeptical off-path beliefs identified in section 2 following messages $m \neq \emptyset$. To that end, denote $\underline{\theta}(m) = \inf\{\theta : m \in M(\theta)\}$. The payoff to $m = \emptyset$ for type θ is

$$\begin{aligned}\int u(z, \theta) d\hat{G}(z | \theta) &= \int \int u(\mathbb{E}[\theta | s, m], \theta) d\mu(m | \theta) dG(s | \theta) \\ &\geq U^e(\theta)\end{aligned}$$

where the first line uses $z = \mathbb{E}[\theta | s, m]$, (18) and Fubini's theorem, and the inequality follows from (19) and $c(m) \geq 0$. By contrast, for any $m \in M$ we have

$$U^e(\theta) \geq \int u(\mathbb{E}[\theta | s, m], \theta) dG(s | \theta) - c(m) \geq \int u(\underline{\theta}(m), \theta) dG(s | \theta) - c(m)$$

which implies $U^e(\theta) \geq \bar{u}(\theta)$, for all $\theta \in [\underline{\theta}, \bar{\theta}]$, and hence $\hat{\mu}$ is an equilibrium strategy. Moreover, it is obvious by definition that \hat{G} induces the same distribution over z as does the pair (G, μ) and hence the regulator's payoff is unchanged.

²⁴ ΔX denotes the set of probability measures on a set X , and X^c the complement of X .

Since we have just shown that the use of skeptical off-path beliefs is without loss of generality, Lemma 2 follows immediately.

Proof of Lemma 3

Let G be feasible in Problem (8). Using Lemma 6, assume without loss that G is threshold-separable with a failure threshold at some level θ' ; moreover, we may assume that (i) $z = \theta$, and so $x = \theta$, for all $\theta < \theta'$, and (ii) $\Pr_G[z \geq t \mid \theta] = 1$ for $\theta \geq \theta'$. Letting H be the distribution of $x(\theta) = \mathbb{E}[z \mid \theta]$ induced by G , the value of the objective in the relaxed problem evaluated at G can be written:

$$\begin{aligned}
\int v(x)dH(x) &= \Pr[\theta < \theta'] \mathbb{E}[v(\theta) \mid \theta < \theta'] + \Pr[\theta \geq \theta'] \mathbb{E}[v(\mathbb{E}[z \mid \theta]) \mid \theta \geq \theta'] \\
&= \Pr[\theta < \theta'] \underbrace{\mathbb{E}[v(z) \mid \theta < \theta']}_{z=\theta, \text{ for } \theta < \theta'} + \Pr[\theta \geq \theta'] \underbrace{\mathbb{E}[\mathbb{E}[z \mid \theta] \mid \theta \geq \theta']}_{\text{by linearity for } z \geq t} \\
&= \Pr[\theta < \theta'] \mathbb{E}[v(z) \mid \theta < \theta'] + \Pr[\theta \geq \theta'] \underbrace{\mathbb{E}[v(z) \mid \theta \geq \theta']}_{\text{by LIE + linearity}} \\
&= \mathbb{E}[v(z)] \equiv V
\end{aligned}$$

Now we show that G satisfies the constraints in Problem (9) if it satisfies those in Problem (8). By assumption G satisfies the incentive constraints in (8), and so

$$\begin{aligned}
E[u(z, \theta) \mid \theta] &= E[(\phi z + (1 - \phi)\theta) \mid \theta] \\
&= \phi E[z \mid \theta] + (1 - \phi)\theta \\
&= \phi x(\theta) + (1 - \phi)\theta \\
&\geq \bar{u}(\theta), \forall \theta \geq \theta'.
\end{aligned}$$

Fix some $x' \geq t$. For any type $\theta \in \delta(x')$, $\bar{u}(\theta) \geq u(x', \theta)$ holds by definition. By contrast, $\bar{x}(\theta)$ satisfies $\bar{u}(\theta) = u(\bar{x}(\theta), \theta)$. Given that utility is increasing in its first argument, this implies $\bar{x}(\theta) \geq x'$ if and only if $\theta \in \delta(x')$. Similarly, the above chain of inequalities implies $x(\theta) \geq \bar{x}(\theta)$. Taken together, this means that for any value x' , G must satisfy

$$\begin{aligned}
\Pr_G[x(\theta) \geq x'] &\geq \Pr[\bar{x}(\theta) \geq x'] \\
&= \Pr[\theta \in \delta(x')]
\end{aligned}$$

or, equivalently,

$$1 - H(x) \geq 1 - F \circ \delta(x)$$

Since x' was arbitrary, this also holds for all $x' \geq t$.

Finally, since x is a conditional expectation of z , it is trivially a mean preserving contraction of z . Hence, $H' \preceq_{co} G'$. As G' is feasible in (8), we have $G' \preceq_{co} F$. The convex order is transitive, so that $H' \preceq_{co} F$.

We briefly remark that $\mathbb{E}[w(z, \theta) | z] = z\lambda(z) = \mathbb{E}[u(z, \theta) | z]$. Hence, maximizing $\mathbb{E}[v]$ is equivalent to maximizing $\mathbb{E}[u]$. Hence, Problem (9) also constitutes a relaxed problem if the objective in Problem (8) is replaced with the firm's ex ante expected payoff.

Proof of Proposition 1

By Lemma 8 (Appendix A.1.2), we need only verify that the candidate test induces the distribution H^* over $x(\theta) = \mathbb{E}[z | \theta]$. First, for any value $x' \leq \theta^*$, the test clearly induces $\Pr[x(\theta) \leq x'] = F(x')$. Moreover, the test ensures $E[v(z) | \theta] = E[z | \theta] = x(\theta) = \bar{x}(\theta)$ for any type $\theta \in \delta(t)$ (where $\bar{x}(\theta) > t$), and $x \leq t$ for $\theta \notin \delta(t)$. This implies $H^*(x) = F \circ \delta(x)$ for $x > t$. Since probabilities must sum to 1, this pins down the required mass at t .

Proof of Proposition 2

Precisely, we prove the following two statements, which imply Proposition 2.

Lemma 9. [*Properties of any optimal test.*] *Up to a set of measure zero:*

1. $z \geq t$ if and only if $\theta \geq \theta^*$. If w is strictly convex on $[\bar{\theta}, t]$ then each $\theta < \theta^*$ is fully revealed,
2. Any type $\theta \in \delta^c(t) \cap [\theta^*, \bar{\theta}]$ receives grade t with probability 1,
3. Any type $\theta \in \delta(t)$ receives an expected grade $x(\theta) = \mathbb{E}[z | \theta] = \bar{x}(\theta)$,

where θ^* is unique and identified by Lemma 8.

Proof. Property 1 follows from Lemma 6 and Proposition 5.3. If $\theta \in \delta(t)$ then $\theta \geq t$ and $\bar{x}(\theta) > t$. Combining this observation with Proposition 5.2, property 3 follows. Moreover, this also implies $\Pr[x > t | \theta \in \delta^c(t) \cap [\theta^*, \bar{\theta}]] = 0$. Hence, $x(\theta) = \mathbb{E}[z | \theta] = t$ for all $\theta \in \delta^c(t) \cap [\theta^*, \bar{\theta}]$. Since such types cannot receive a grade less than t , this implies they must also have a zero probability of earning a grade strictly in excess of t . \square

Lemma 10. *There does not exist a deterministic test which solves (8).*

Proof. We show that no deterministic test can simultaneously satisfy the necessary optimality conditions of Lemma 9 and induce a distribution G over z which constitutes a mean

preserving contraction of F . For the sake of a contradiction, suppose such a deterministic mapping $\pi_d : \Theta \rightarrow S$ did exist. Then, Lemma 9.3 implies that $\pi_d(\theta) = z(\theta) = \bar{x}(\theta) < \theta$, for all $\theta \in \delta(t)$.²⁵ By Lemma 9.1-2 we can further conclude $z > t \iff \theta \in \delta(t)$.

Consider the conditional distributions of z and θ , respectively, given $z > t$. Since $z(\theta) = \bar{x}(\theta) < \theta$ for all $\theta \in \delta(t)$ and $z(\theta) = t$ for all $\theta \in \delta^c(t) \cap [\theta^*, \bar{\theta}]$, we necessarily have $\Pr[z < z' \mid z > t] > \Pr[\theta < z' \mid z > t]$ for all $z' \in (t, \bar{\theta}]$. Hence, $\mathbb{E}[z \mid z > t] < \mathbb{E}[\theta \mid z > t] = \mathbb{E}[\theta \mid \theta \in \delta(t)]$, and:

$$\begin{aligned} \mathbb{E}[z] &= \Pr[z \leq t] \mathbb{E}[z \mid z \leq t] + \Pr[z > t] \mathbb{E}[z \mid z > t] \\ &= \Pr[\theta \in \delta^c(t)] \mathbb{E}[\mathbb{E}[\theta \mid z] \mid z \leq t] + \Pr[\theta \in \delta(t)] \mathbb{E}[z \mid z > t] \\ &< \Pr[\theta \in \delta^c(t)] \mathbb{E}[\theta \mid \theta \in \delta^c(t)] + \Pr[\theta \in \delta(t)] \mathbb{E}[\theta \mid \theta \in \delta(t)] \\ &= \mathbb{E}[\theta], \end{aligned}$$

where the second line uses $\Pr[z \leq t] = \Pr[\theta \in \delta^c(t)]$ and $z = \mathbb{E}[\theta \mid z]$, and the inequality follows after applying LIE for $z \leq t$ and $\mathbb{E}[z \mid z > t] < \mathbb{E}[\theta \mid \theta \in \delta(t)]$. We conclude $\mathbb{E}[\theta] = \mathbb{E}[\mathbb{E}[\theta \mid z]] = \mathbb{E}[z] < \mathbb{E}[\theta]$ —a contradiction. \square

Proof of Proposition 3

Let $G_{\bar{x}}^*$ denote an optimal test when the outside option is described by the function \bar{x} , and let the corresponding equilibrium threshold be $\theta_{\bar{x}}^*$. Fix two functions \bar{x}, \bar{x}' satisfying $\bar{x}(\theta) \leq \bar{x}'(\theta)$ for all $\theta \in [\theta, \bar{\theta}]$. To show property 1, recall condition (11) defining the threshold can be written

$$\int_{\theta_{\bar{x}}^*}^{\bar{\theta}} (\theta - \max\{t, \bar{x}(\theta)\}) dF(\theta) = 0.$$

Direct comparison shows that $\int_k^{\bar{\theta}} (\theta - \max\{t, \bar{x}'(\theta)\}) dF(\theta) \leq 0$ for $k = \theta_{\bar{x}}^*$. Since $\theta_{\bar{x}}^* \leq t$, the integral is locally increasing in k and hence $\theta_{\bar{x}'}^* \geq \theta_{\bar{x}}^*$. Note that the inequality is strict if and only if $\Pr[\{\theta : \bar{x}'(\theta) > \bar{x}(\theta)\}] > 0$.

To see property 2, note that, by the linearity of v on $[t, \bar{\theta}]$, the regulator's payoff from test $G_{\bar{x}}^*$ can be written

$$V(\bar{x}) = \int_{\theta}^k v(\theta) dF(\theta) + \int_k^{\bar{\theta}} \theta dF(\theta)$$

for $k = \theta_{\bar{x}}^*$. As $v(\theta) < \theta$ for all $\theta < t$, it is clear that the right side of this expression strictly

²⁵Since the test is deterministic, we can identify the grade $z(\theta)$ as a function of θ .

decreases in k for $k \leq t$. Hence $V(\bar{x}') \leq V(\bar{x})$, with strict inequality if $\theta_{\bar{x}'}^* < \theta_{\bar{x}}^*$.

Finally, for comparative statics of informativeness in \bar{x} we focus on the implementation of the optimal test described in Proposition 1. In this case, the result is almost immediate. By property 1, $\theta_{\bar{x}'}^* > \theta_{\bar{x}}^*$. Moreover, note that $\delta_{\bar{x}}(t) \subset \delta_{\bar{x}'}(t)$ for all x . The distribution $G_{\bar{x}}^*$ satisfies

$$G_{\bar{x}}^*(z) = \begin{cases} F(z), & z < \theta_{\bar{x}}^* \\ F(\theta_{\bar{x}}^*) + \int_{\theta \in \delta_{\bar{x}}(t) \cap [\underline{\theta}, z]} \alpha(\theta) dF(\theta), & \theta_{\bar{x}}^* \leq z. \end{cases}$$

As they are clearly ordered by single-crossing and have the same mean, $G_{\bar{x}'}^*$ second-order stochastically dominates $G_{\bar{x}}^*$ (see, for instance, Rothschild and Stiglitz (1970)).

Proof of Lemma 4

As described in section 2, message space M , costs c , and ϕ induce the function $\bar{u}(\theta)$, which defines the disclosure-equivalent grades $\bar{x}(\theta) = \phi^{-1}(\bar{u}(\theta) - (1 - \phi)\theta)$.

Consider comparative statics with respect to $M(\theta)$: fix two message spaces $(M(\theta))_{\theta \in \Theta}$, $(M'(\theta))_{\theta \in \Theta}$ such that $\bigcup_{\theta \in \Theta} M(\theta) = M = \bigcup_{\theta \in \Theta} M'(\theta)$ and $M^{-1}(m) \subset M'^{-1}(m)$ for all $m \in M$. Since $M^{-1}(m) \subset M'^{-1}(m)$, we have $\underline{\theta}'(m) \leq \underline{\theta}(m)$ for all $m \in M$, where $\underline{\theta}(m) = \min\{\theta : m \in M(\theta)\}$ and $\underline{\theta}'(m)$ is similarly defined. Since this holds for all m , it follows that $\bar{u}'(\theta) \leq \bar{u}(\theta)$, where $\bar{u}'(\theta)$ has the obvious definition. From here, $\bar{x}'(\theta) \leq \bar{x}(\theta)$ follows immediately. By similar consideration of (5), it is easy to see that $\bar{x}'(\theta) \geq \bar{x}(\theta)$ if $c'(m) \leq c(m)$ for all $m \in M$.

Finally, we establish comparative statics in ϕ . Note that

$$\begin{aligned} \bar{x}(\theta) &= \phi^{-1}\left(\max_{m \in M(\theta)} \{u(\underline{\theta}(m), \theta) - c(m)\} - (1 - \phi)\theta\right) \\ &= \max_{m \in M(\theta)} \{\phi^{-1}(u(\underline{\theta}(m), \theta) - c(m) - (1 - \phi)\theta)\} \\ &= \max_{m \in M(\theta)} \left\{ \underline{\theta}(m)(1 - \lambda(\underline{\theta}(m))) - \frac{1 - \phi}{\phi} \theta \lambda(\underline{\theta}(m)) - \frac{1}{\phi} c(m) \right\}, \end{aligned}$$

where the third line uses the definition of $u(z, \theta)$. Clearly, the maximand in the last line increases in ϕ for each m , and therefore $\bar{x}(\theta)$ increases in ϕ too.

Proof of Proposition 4

Consider a family $\{F_y\}_{y \in Y}$ of distributions on $[\underline{\theta}, \bar{\theta}]$ indexed by y . In order that we can compare distributions using MLRP, assume that each F_y has a corresponding density f_y ; for simplicity let $f_y(\theta) > 0$ for all $\theta \in [\underline{\theta}, \bar{\theta}]$.²⁶ For the purposes of describing some comparative

²⁶Our results do not rely on this simplification. It serves to simplify our description of the MLRP order.

statics, assume Y is a totally ordered set. We order distributions by the monotone likelihood ratio (MLR) order, so that if $y < y'$ and $\theta < \theta'$ then

$$\frac{f_y(\theta')}{f_y(\theta)} < \frac{f_{y'}(\theta')}{f_{y'}(\theta)}.$$

We can now prove the result. Consider condition (11)

$$\int_{\theta_y^*}^{\bar{\theta}} (\theta - \max\{t, \bar{x}(\theta)\}) f_y(\theta) d\theta = 0. \quad (20)$$

Fixing $\theta^* = \theta_y^*$, the integrand on the left-hand side again obeys the single-crossing property. Hence, for $y' > y$ the left-hand side is positive. Since $\theta_y^* < t$, this expression is increasing in θ^* . Hence, $\theta_{y'}^* < \theta_y^*$. To see that the regulator's utility is increasing in y , note that her optimal payoff can be written as a function of y as follows:

$$V(y) = \int_{\underline{\theta}}^{\theta_y^*} v(\theta) f_y(\theta) d\theta + \int_{\theta_y^*}^{\bar{\theta}} \theta f_y(\theta) d\theta.$$

As $v(\theta)$ is increasing in θ , the right-hand side is increasing in y for fixed θ_y^* . Moreover, as $\theta > v(\theta)$ for all $\theta < t$, the right-hand side is also decreasing in θ_y^* . Putting these together shows that $V(y) < V(y')$ whenever $y < y'$.

B Proofs of Preliminary Results

Proof of Lemma 5

We first prove that, if (15) holds for all r , with equality at $r = \theta', \bar{\theta}$ then $G(\theta') = F(\theta')$. For the sake of a contradiction, suppose first that $G(\theta') > F(\theta')$. By the right-continuity of distribution functions, there exists a $r > \theta'$ such that $F(z) < G(z)$ for all $z \in [\theta', r]$. But then,

$$\int_{\underline{\theta}}^{\theta'} F(z) dz + \int_{\theta'}^r F(z) dz < \int_{\underline{\theta}}^{\theta'} G(z) dz + \int_{\theta'}^r G(z) dz,$$

which violates (15). On the other hand, if $G(\theta') < F(\theta')$, then—since F is continuous by assumption—there is similarly a $r < \theta'$ such that $1 - G(z) > 1 - F(z)$ on $[r, \theta']$. From here,

an almost identical argument, applied to the integrals $\int_r^{\bar{\theta}}(1 - G(z))dz$, $\int_r^{\bar{\theta}}(1 - F(z))dz$, yields another contradiction.

We now prove the equivalence between threshold-separable signals and the condition in the Lemma. If a test is threshold-separable, then it follows immediately from the mean preserving contraction criterion for signals, applied to $[\underline{\theta}, \theta']$, $[\theta', \bar{\theta}]$, respectively, that (15) holds. Similarly, it follows from Rothschild and Stiglitz (1970) applied to $[\underline{\theta}, \theta']$, $[\theta', \bar{\theta}]$ respectively that for any distribution G satisfying (15)—with equality at $r = \theta', \bar{\theta}$ —there is a threshold separable signal which generates it.

The only remaining issue to prove is that no other test could generate G . For a contradiction, suppose there is a test $\{G(\cdot | \theta)\}_{\theta \in \Theta}$ with $\text{supp } G(\cdot | \theta) \cap \text{supp } G(\cdot | \theta) \neq \emptyset$ which induces a G satisfying (15). Hence, G must also satisfy $G(\theta') = F(\theta')$. Integrating by parts, $G(\theta') = F(\theta')$ and $\int_{\underline{\theta}}^{\theta'} G(z)dz = \int_{\underline{\theta}}^{\theta'} F(z)dz$ together imply that $\mathbb{E}_G[z | z \leq \theta'] = \mathbb{E}_F[z | z \leq \theta']$. But if G is not threshold-separable, then $\Pr_G[\theta \leq k, s \leq \theta'] \leq F(k)$ for all $k \leq \theta'$, with *strict inequality* for some k . Hence, $G(\theta') = F(\theta')$ also implies the conditional distribution $G(z | z \leq \theta')$ first order stochastically dominates $F(z | z \leq \theta')$, so that $\mathbb{E}_G[z | z \leq \theta'] > \mathbb{E}_F[z | z \leq \theta']$ —a contradiction.

Proof of Lemma 6

Fix some candidate optimal test inducing a distribution G . If $G(t) = 0$ then G satisfies the hypothesis of the Lemma with $\theta' = \underline{\theta}$. So suppose from now on that $G(t) > 0$. We argue that, if there is no such threshold, then we can find another feasible, incentive compatible structure G'' which strictly increases the objective in Problem (8). We construct G'' in two steps.

Step 1: We construct an intermediate test G' which weakly improves on G , and which ensures that no type $\theta \leq t$ induces a grade $z > t$ with positive probability. Define $T := \{\theta : \theta > t\}$. By the Radon-Nikodym theorem, the random variables $\Pr[T | z] = 1 - G(t | z)$, $h_z := \mathbb{E}[\theta | T, z]$ and $l_z := \mathbb{E}[\theta | T^c, z]$ are all measurable functions of z .²⁷ Notice that $h_z > l_z$ for all z , and $l_z \leq t < h_z$ for all $z > t$. We construct G' by splitting each grade $z > t$ into a pair $\{t, h_z\}$ where h_z is induced with a conditional probability $(1 - G(t | z))\alpha_z$, with α_z defined by

$$\alpha_z = \begin{cases} 0, & z \leq t \\ 1 - \frac{G(t|z)(t-l_z)}{(1-G(t|z))(h_z-t)}, & z > t \end{cases}$$

²⁷If $G(t | z) \in \{0, 1\}$, define $h_z = \bar{\theta}$ or $l_z = \underline{\theta}$ as appropriate.

α_z is continuous in $G(t | z)$, h_z and l_z for $z > t$, and so inherits measurability. Moreover, it lives in $[0, 1]$ and satisfies, for all $z > t$:

$$\Pr[T | z]\alpha_z h_z + (1 - \Pr[T | z]\alpha_z)t = z$$

α_z is a conditional probability of inducing grade h_z , given the outcome of test G and $\theta \in T$. Under this new test, the implied distribution of grades is

$$G'(z) = \begin{cases} G(z), & z < t \\ G(t) + \int (1 - \alpha_{z'} \Pr[T | z'])dG(z'), & z = t \\ G'(t) + \int_{\{z': h_{z'} < z\}} \alpha_{z'} \Pr[T | z']dG(z'), & z > t. \end{cases}$$

G' corresponds to augmenting G with a commitment to reveal $\{\theta \in T\}$ with conditional probability α_z . Since this is clearly less informative than augmenting the random variable z' induced by G with full revelation of θ , the conditional distributions trivially satisfy $G'(z | z') \preceq_{co} F(z | z')$ for each z' , where \preceq_{co} denotes the convex order.²⁸ But the convex order is preserved under integration, and so $G' \preceq_{co} F$ too.

Next, notice that

$$\begin{aligned} \int v(z)dG'(z) &= \int (\alpha_{z'} \Pr[T | z']v(h_{z'}) + (1 - \alpha_{z'} \Pr[T | z'])v(t))dG(z') \\ &= \int v(\alpha_{z'} \Pr[T | z']h_{z'} + (1 - \alpha_{z'} \Pr[T | z'])t)dG(z') \\ &= \int v(z')dG(z') \end{aligned}$$

since v is linear for $z \geq t$. Hence, G' is a feasible test which weakly improves the objective. We postpone verification of incentive compatibility to the next step.

Step 2: We construct a distribution G'' with the required properties which strictly improves on G' . Our aim is to alter the distribution of those grades $z \leq t$ that would have been realized under G' so that G'' becomes threshold-separable (Appendix A.1.1) with threshold θ' . To that end, consider the conditional distribution $G'(\theta | \tilde{z} \leq t)$ of types given a grade no higher than t . Applying Corollary 2 to $G'(\cdot | \tilde{z} \leq t)$, there is a threshold-separable distribution $G^t \preceq_{co} G'(\cdot | \tilde{z} \leq t)$ with some threshold θ' , such that $\int v(z)dG^t(z) \geq \int v(z)dG'(z | \tilde{z} \leq t)$. Moreover, by Corollary 2 we may further assume $G^t(z) = G'(z | \tilde{z} \leq t)$ for $z \leq \theta'$, so that G^t reveals types below θ' . Now consider the

²⁸With a slight abuse, here we briefly adopt the notation $G'(z | z')$, $F(z | z')$ to emphasize that we are referring to conditional distributions (i.e., given z') over *final grades* from the respective augmented tests.

distribution G'' defined by

$$G''(z) = \begin{cases} G'(t) \cdot G^t(z) & , \text{ if } z \leq t \\ G'(z) & , \text{ otherwise.} \end{cases}$$

Noting that $G'(z) = G'(t) \cdot G'(z | \tilde{z} \leq t)$, G' can be written

$$G'(z) = \begin{cases} G'(t) \cdot G'(z | \tilde{z} \leq t) & , \text{ if } z \leq t \\ G'(z) & , \text{ otherwise.} \end{cases}$$

It is clear by inspection that G'' adjusts G' by replacing the distribution $G'(z | \tilde{z} \leq t)$ of grades below t with G^t . Using this characterization, it follows from $G^t \preceq_{co} G'(\cdot | \tilde{z} \leq t)$ that $\int_{\underline{\theta}}^z G''(s)ds \leq \int_{\underline{\theta}}^z G'(s)ds$ for all z , and hence $G'' \preceq_{co} G' \preceq_{co} F$, so that G'' is a feasible distribution of grades. Moreover, by construction of G' , every type $\theta \leq t$ send signals $\tilde{z} \leq t$ with probability 1 so that $\Pr_{G'}[\tilde{z} \leq t, \tilde{\theta} \leq \theta] = \Pr_{G'}[\tilde{\theta} \leq \theta] = F(\theta)$. Since $G^t(z) = G'(\cdot | \tilde{z} \leq t)$ for $z \leq \theta'$ this implies $G'(t) \cdot G^t(z) = G'(t) \cdot G'(z | \tilde{z} \leq t) = F(z)$ for all $z \leq \theta'$. Applying Lemma 5, this proves that G'' is threshold separable with threshold θ' . Moreover, since G^t assigns grade t for $\theta \geq \theta'$, this shows that if $\theta \geq \theta'$ then $G''(\cdot | \theta)$ induces $z \geq t$ with probability 1.

To prove incentive compatibility, we first identify a test $\{G''(\cdot | \theta)\}_{\theta \in \Theta}$ that induces the marginal distribution G'' . In particular, consider the structure defined by

$$G''(z | \theta) = \mathbb{1}_{\theta \leq \theta'} \cdot \mathbb{1}_{z \geq \theta} + \mathbb{1}_{\theta > \theta'} \cdot \mathbb{1}_{z \geq t} \cdot G'(z | \theta).$$

To see intuitively that this test induces G'' , note that for $\theta \leq \theta'$, the signal $z = \theta$ occurs with probability 1 under $G''(\cdot | \theta)$. By contrast, for $\theta \geq \theta'$, $G''(\cdot | \theta)$ reassigns all the mass under $G'(\cdot | \theta)$ from grades $z < t$ to t , and otherwise leaves the conditional distribution unchanged. Hence, this test indeed adjusts G' in the way described in the previous paragraph. To prove this intuition formally, we must show two things: (i) the structure does induce marginal $\Pr_{G''}[\tilde{z} \leq z] = G''(z)$ and (ii) the grades satisfy $\mathbb{E}_{G''}[\theta | z] = z$, for all z .²⁹ Integrating across θ shows that

$$\Pr_{G''}[\tilde{z} \leq z] = \int_{\underline{\theta}}^{\bar{\theta}} G''(z | \theta) dF(\theta) = \int_{\underline{\theta}}^{\theta'} \mathbb{1}_{z \geq \theta} dF(\theta) + \int_{\theta'}^{\bar{\theta}} \mathbb{1}_{z \geq t} G'(z | \theta) dF(\theta).$$

For $z < t$, the second integral is 0 and the expression reduces to $\Pr_{G''}[\tilde{z} \leq z] = F(\min\{z, \theta'\}) =$

²⁹For each z , G'' is the sum of products of measurable functions, and hence measurable.

$G''(z)$. For $z \geq t$, the second integral reduces to $\int_{\theta'}^{\bar{\theta}} G'(z | \theta) dF(\theta) \equiv \Pr_{G'}[\tilde{z} \leq z, \tilde{\theta} > \theta']$. Recalling that for all $z \geq t$, $F(\theta') = G'(t) \cdot G'(z | \tilde{z} \leq t) = \Pr_{G'}[\tilde{z} \leq t, \theta \leq \theta'] = \Pr_{G'}[\tilde{z} \leq z, \theta \leq \theta']$, the right-hand side reduces to

$$\Pr_{G'}[\tilde{z} \leq z, \theta \leq \theta'] + \Pr_{G'}[\tilde{z} \leq t, \theta > \theta'] = G'(z).$$

Recalling $G'(z) = G''(z)$ for such z establishes that the signal structure induces G'' . To show (ii), we note that it suffices to prove $\mathbb{E}_{G''}[\theta | t] = t$: indeed, for any $z \neq t$ $G''(\cdot | \theta)$ is either fully revealing or identical to $G'(\cdot | \theta)$, and so it is not difficult to see that $\mathbb{E}_{G''}[\theta | z] = \mathbb{E}_{G'}[\theta | z] = z$.³⁰ If $\Pr_{G''}[\tilde{z} = t] = 0$, then without loss of generality we can simply assign $\mathbb{E}_{G''}[\theta | t] = t$.³¹ So suppose $\Pr_{G''}[s = t] = 0$. Then, by construction the distribution of grades $z \leq t$ under distribution G'' satisfies

$$\int_{\{\theta < t\}} \theta dF + t \Pr_{G''}[\tilde{z} = t] = G'(t) \cdot \int_{\underline{\theta}}^{\bar{\theta}} \theta dG'(\theta | \tilde{z} \leq t). \quad (21)$$

However, by the law of iterated expectations, the distribution of $\mathbb{E}[\theta | z]$ induced by test $\{G''(\cdot | \theta)\}$ satisfies

$$\int_{\{z < t\}} \int_{\underline{\theta}}^{\bar{\theta}} \theta G''(d\theta, dz) = \int_{\underline{\theta}}^t \mathbb{E}[\theta | z] dG''(z) = \int_{\{\theta < t\}} \theta dF + \mathbb{E}_{G''}[\theta | t] \Pr_{G''}[\tilde{z} = t], \quad (22)$$

where $G''(\theta, z)$ denotes the joint distribution of $(\tilde{\theta}, \tilde{z})$ induced by $\{G''(\cdot | \theta)\}$, and the second equality follows from the threshold-separability of G'' . Rearranging the right-hand side shows that

$$\int_{\{z < t\}} \int_{\underline{\theta}}^{\bar{\theta}} \theta G''(d\theta, ds) = \int_{\underline{\theta}}^{\bar{\theta}} \theta G''(d\theta, t) = G'(t) \cdot \int_{\underline{\theta}}^{\bar{\theta}} \theta dG^t(\theta)$$

But since G^t is a MPC of $G'(\cdot | \tilde{z} \leq t)$, $\int_{\underline{\theta}}^{\bar{\theta}} \theta dG^t(\theta) = \int_{\underline{\theta}}^{\bar{\theta}} \theta dG'(\theta | \tilde{z} \leq t)$. Given this, it is clear on inspection of equations (21) and (22) that $t = \mathbb{E}_{G''}[\theta | t]$.

Finally, we show that G'' is incentive compatible. For $\theta < \theta'$, G'' reveals θ . As $\bar{u}(\theta) \leq u(\theta, \theta)$ for all θ , incentive compatibility is trivially satisfied for all $\theta \leq \theta'$. For $\theta \in (\theta', t)$,

³⁰One can directly verify that the conditional expectation $\mathbb{E}_{G''}[\theta | z]$ satisfies the required properties for all z by finding a valid conditional distribution (Radon-Nikodym derivative) $G''(z | \theta)$ associated with $G''(z | \theta)$. While not difficult, it is tedious and hence we omit it. Details available on request.

³¹See, for example, Billingsley (1995), pg. 445.

$z = t$, which is again trivially incentive compatible. Finally, for $\theta \geq t$ we have

$$\begin{aligned} \int u(z, \theta) dG''(z | \theta) &\geq \int u(z, \theta) dG'(z | \theta) \\ &= \int (\alpha_{z'} u(h_{z'}, \theta) + (1 - \alpha_{z'}) t) dG(z' | \theta) \\ &\geq \int u(z', \theta) dG(z' | \theta) \end{aligned}$$

where the first inequality follows as u is increasing in z and $G''(\cdot | \theta)$ first-order stochastically dominates $G(\cdot | \theta)$, the second uses the definition of G' and the third uses the linearity of $u(z, \theta)$ for $z \geq t$, $h_z > t$ and $\alpha_{z'} \Pr[T | z'] h_{z'} + (1 - \alpha_{z'} \Pr[T | z']) t = z'$ to establish the integrands are ordered pointwise in z' . Hence, G'' is incentive compatible if G is.

Proof of Lemma 7

Suppose such a γ , ψ and H^* exist. Then, for any feasible distribution H we have

$$\begin{aligned} \int v dH &= \int (v - \gamma - \psi) dH + \int \gamma dH + \int \psi dH \\ &\leq \int \gamma dF + \int \psi dF \circ \delta \\ &= \int v dH^*. \end{aligned}$$

where the inequality follows by applying condition (1.) to the first term on the right-hand side, noting $H \preceq_{co} F$ for the second, and noting that H first-order stochastically dominates $F \circ \delta$ for the third. The final line follows immediately from application of (1.)-(3.). It is straightforward to see that any optimal H' must meet conditions (1.)-(3.) (and the constraints in Problem (9)) in order to satisfy the above with equality.

Proof of Lemma 8

We identify functions γ^* , ψ^* for which the verification result of Lemma 7 applies. Consider the following functions:

$$\gamma^*(x) = \begin{cases} v(x), & x \leq \theta^* \\ v(\theta^*) + \gamma^+(x - \theta^*), & x > \theta^* \end{cases}$$

where $\gamma^+ = \frac{v(t)-v(\theta^*)}{t-\theta^*}$, and

$$\psi^*(x) = \begin{cases} 0, & x < t \\ v(x) - \gamma(x), & x \geq t \end{cases}$$

Clearly, γ^* is convex, ψ^* is decreasing and $\gamma^* + \psi^* \geq v$ (with equality on the support of H^*). As H^* is identical to F on $[0, \theta^*]$ and γ^* linear on $[\theta^*, \bar{\theta}]$, we have $\int \gamma^* dH^* = \int \gamma^* dF$. By a similar argument, $\int \psi^* dH^* = \int \psi^* dF \circ \delta$. H^* trivially satisfies the FOSD constraint. Finally, given the definition of θ^* , and since $F(x) \leq F \circ \delta(x)$ for $x \geq t$, it is easy to verify (by a single-crossing argument) that $\int_{\underline{\theta}}^{x'} H^*(x) dx \leq \int_{\underline{\theta}}^{x'} F(x) dx$ for all $x' \in [\underline{\theta}, \bar{\theta}]$.

Proof of Proposition 5

Properties (1) and (2) follow immediately by application of conditions (1)-(3) of Lemma 7 to γ^* , ψ^* . To prove property (3), we note that γ^* can be written

$$\gamma^*(x) = \tilde{\gamma}(x) + \lambda \max\{0, x - \theta^*\}$$

where $\tilde{\gamma}$ is the convex function defined as follows: $\tilde{\gamma}(x) = v(x)$ for $x \leq \theta^*$, and $\tilde{\gamma}(x) = \gamma^*(\theta^*) + \frac{d\gamma^*}{d\theta^*} \cdot (x - \theta^*)$ where $\frac{d\gamma^*}{d\theta^*} := \lim_{x \uparrow \theta^*} \frac{\gamma^*(\theta^*) - \gamma^*(x)}{\theta^* - x}$ for $x > \theta^*$, and $\lambda = \gamma^+ - \frac{d\gamma^*}{d\theta^*} > 0$. If test H is not threshold-separable, then by Lemma 5 (Appendix A.1.1), $\int_{\underline{\theta}}^{\theta^*} H(x) dx < \int_{\underline{\theta}}^{\theta^*} F(x) dx$.

This implies

$$\begin{aligned} \int \gamma^* dH &= \int \tilde{\gamma} dH + \lambda \int \max\{0, x - \theta^*\} dH \\ &= \int \tilde{\gamma} dH + \lambda \int_{\theta^*}^{\bar{\theta}} (1 - H(x)) dx \\ &< \int \tilde{\gamma} dF + \lambda \int_{\theta^*}^{\bar{\theta}} (1 - F(x)) dx \\ &= \int \gamma^* dF \end{aligned}$$

where the strict inequality follows by convexity of $\tilde{\gamma}$ and $\int_{\underline{\theta}}^{\theta^*} H(x) dx < \int_{\underline{\theta}}^{\theta^*} F(x) dx$. But this contradicts condition 2. of Lemma 7.

C Extensions

C.1 Regulatory Previews

Here we prove an extended version of Lemma 1, in which we also allow the regulator to provide a preview to the firm about the likely outcome of the test. Specifically, we enrich the regulator's design problem from section 2 by allowing her to provide an additional private signal $\rho : \Theta \times S \rightarrow \Delta\mathbb{R}$ to the firm.³² Note that, in principle, ρ is flexible enough to allow the preview to depend both on the firm's type and on the grade; equivalently, one can think of the regulator being able to tailor information about grades to different types of firm. In particular, notice that this extension allows us to capture variations in the timing of test results as special cases of preview strategy: for instance, pre-disclosure tests correspond to the special case $\rho(\theta, s) = s$, while post-disclosure tests (i.e. the timing considered in the main body) can be captured with the constant function $\rho \equiv 0$. With a slight abuse, we use ρ to denote both the preview strategy and its realizations, where convenient and clear from context. We prove:

Lemma 11. [*Public Speech, Private Silence*] *For any valid test $\{G, \rho\}$, there exists an alternative test $\{\hat{G}, \hat{\rho}\}$ such that i) the regulator obtains the same expected utility, ii) the firm makes no verifiable disclosures in equilibrium under $\{\hat{G}, \hat{\rho}\}$, and iii) the regulator does not send an informative private signal to the firm with $\hat{\rho} \equiv 0$.*

Proof. Fixing ρ , an almost identical argument to the proof of Lemma 1 shows that (G, ρ, μ) can be replaced with a pair $(\hat{G}, \rho, \hat{\mu})$ such that (i) and (ii) hold.³³ The only difference is that the incentive constraints become:

$$\int u(z, \theta) d\hat{G}(z | \theta, r) \geq \bar{u}(\theta) \quad (23)$$

for all θ and $r \in \text{supp}\rho(\cdot | \theta)$. Integrating across r , (23) implies $\int u(z, \theta) d\hat{G}(| \theta) \geq \bar{u}(\theta)$ for all θ . That is, $\hat{\mu}$ is also an equilibrium strategy under $(\hat{G}, \hat{\rho})$, where $\hat{\rho} \equiv 0$. \square

C.2 On the strict optimality of private silence

We argue there does not exist a test G and corresponding equilibrium disclosures μ in which (i) a positive measure of types $\theta \geq \theta^*$ disclose with positive probability and (ii) the

³²More formally, we require ρ to be a measurable function.

³³In particular, since \hat{G} is a mean-preserving spread of G (see Lemma 1), it is feasible to provide the firm with the same information structure ρ . Indeed, if ρ is a Blackwell-garbling of G , then it can be constructed as a Blackwell-garbling of \hat{G} too.

regulator's equilibrium payoff is V .

We first argue that types $\theta \in \delta_+(t) := \{\theta : \bar{u}(\theta) > u(t, \theta)\}$ —i.e., the set where $\bar{x}(\theta) > t$ —cannot disclose with positive probability in equilibrium. Otherwise, the payoff of such a type would be

$$\mathbb{E}[u(\mathbb{E}[\tilde{\theta} \mid s, \tilde{m}], \theta) \mid \theta] - \mathbb{E}[c(\tilde{m}) \mid \theta] \geq \bar{u}(\theta)$$

where the inequality follows since type θ has the option to disclose \underline{m} , which is worth at least $\bar{u}(\theta)$ in equilibrium. If θ discloses with positive probability, then $\mathbb{E}[c(\tilde{m}) \mid \theta] > 0$ and so

$$\mathbb{E}[u(\mathbb{E}[\tilde{\theta} \mid s, \tilde{m}], \theta) \mid \theta] - \mathbb{E}[c(\tilde{m}) \mid \theta] > \bar{u}(\theta).$$

But if this inequality holds for a positive measure of types, then the implied distribution G^* of $z = \mathbb{E}[\theta \mid s, m]$ violates Proposition 9.3 and hence cannot be optimal. Therefore if the regulator's payoff is V , then types $\theta \in \delta_+(t)$ must not disclose with positive probability.

We now show this extends to all types $\theta \geq \theta^*$. By Proposition 9.3 any remaining type $\theta \in [\theta^*, \bar{\theta}] \cap \delta_+^c(t)$ must receive grade t with probability 1. Since disclosures are costly and all types in $[\theta^*, \bar{\theta}] \cap \delta_+^c(t)$ must earn the same grade, all types in this set must be disclosing (otherwise, non-disclosers expect a grade of t without paying the disclosure cost). Let m' be this disclosure.³⁴ Moreover, Proposition 9.3 implies there must be some grade s' such that $t = \mathbb{E}[\theta \mid m', s']$ and (m', s') is induced in equilibrium iff $\theta \in [\theta^*, \bar{\theta}] \cap \delta_+^c(t)$. By the previous argument, any type $\theta \in \delta_+(t)$ must not send m' with positive probability. We argue this implies $\mathbb{E}[\theta \mid m', s'] < t$ —a contradiction. To see this, recall that θ^* by definition satisfies

$$tq + (1 - q)\mathbb{E}[\bar{x}(\theta) \mid \bar{x}(\theta) > t] = \mathbb{E}[\max\{t, \bar{x}(\theta)\} \mid \theta \geq \theta^*] = \mathbb{E}[\theta \mid \theta \geq \theta^*],$$

where $q = \Pr[\theta \in \delta_+^c(t) \mid \theta \geq \theta^*]$. Further recall that $\bar{x}(\theta) < \theta$ for all θ , and hence

$$tq + (1 - q)\mathbb{E}[\bar{x}(\theta) \mid \bar{x}(\theta) > t] > \mathbb{E}[\theta \mid \theta \geq \theta^*]. \quad (24)$$

But, by the Law of Iterated Expectations, we have

$$\mathbb{E}[\theta \mid \theta \in [\theta^*, \bar{\theta}] \cap \delta_+^c(t)]q + (1 - q)\mathbb{E}[\bar{x}(\theta) \mid \bar{x}(\theta) > t] = \mathbb{E}[\theta \mid \theta \geq \theta^*], \quad (25)$$

where we have used the fact that $\bar{x}(\theta) > t \iff \theta \in \delta_+(t)$. A direct comparison of (24) and (25) shows that $t > \mathbb{E}[\theta \mid \theta \in [\theta^*, \bar{\theta}] \cap \delta_+^c(t), s'] = \mathbb{E}[\theta \mid m', s']$.

³⁴It is without loss of generality to consider a single disclosure message: if there are multiple disclosure messages, the same argument applies after taking expectations across such messages.

D Formal Analysis of Bank Runs

We rigorously derive Example 1 in the paper from a model of bank runs, which follows Diamond and Dybvig (1983) and Morris and Shin (2000).

Model Environment: The model has three periods $t = 0, 1, 2$. A bank interacts with a continuum of investors. Each investor is endowed with \$1 at $t = 0$. At date 1, a measure λ of investors finds out that they are either patient, obtaining utility $u(c_1) = \log(c_1)$ from consumption at date 1, and a unit measure of other depositors find out that they are patient, obtaining utility $u(c_1 + c_2) = \log(c_1 + c_2)$ from consumption at either $t = 1$ or $t = 2$. As in Diamond and Dybvig (1983), it is not possible to verify or contract on the realization of investors' preferences.

The bank has access to a riskless storage technology that transfers dollars across time periods, and also to a long-term investment technology that yields $Re^{-\kappa\ell}$ at $t = 2$ for every dollar invested at $t = 0$, where $R > 1$ denotes the return on investment if it is held to maturity, and where ℓ is the (endogenously determined) fraction of resources that are withdrawn at $t = 1$. In this formulation, κ is a parameter measuring the illiquidity of the bank's technology.

We assume that the baseline return on investment satisfies

$$\log R = \theta + \eta,$$

where $\theta \in [\underline{\theta}, \bar{\theta}]$ is the unknown quality of the bank, defined as in the paper as a random variable with cumulative distribution $F(\theta)$, and where $\eta \sim N\left(0, \frac{1}{\alpha}\right)$ is an idiosyncratic component of banks' returns, which is statistically independent of θ .

Optimal Contracts: Diamond and Dybvig (1983) and Morris and Shin (2000) show that an efficient allocation can be implemented by the following contract between the bank and its investors: Each investor deposits \$1 in the bank at $t = 0$. At $t = 1$, each investor has the right to withdraw her claim and obtain a fixed repayment of \$1, or to leave her claim in the bank and obtain a share of the final returns to investment, which are shared equally between all investors that do not withdraw at $t = 2$. The bank invests $\$ \lambda$ in storage and the remaining \$1 in the long-term technology. Hence, any withdrawals from impatient investors can be met without (inefficiently) withdrawing resources from the long-term investment.

Inside and Outside Information: At $t = 1$, investors observe a disclosure from the bank m and the regulatory signal s (outside information). These variables contain information

about θ but are independent of η , and are otherwise defined as in the paper. We let $z = E[\theta|s, m]$ denote investors' posterior expectation of θ after these disclosures.

In addition, each investor obtains an outside signal $x_i = \eta + \epsilon_i$ of the idiosyncratic component of bank returns, where $\epsilon_i \sim N\left(0, \frac{1}{\beta}\right)$ is independent of η and θ . This signal is introduced only to generate small deviations from common knowledge among investors, as in the literature on global games in banking (e.g., Goldstein and Pauzner, 2005). We will focus below on the “noiseless” limit where the precision of agent’s signals β tends to infinity. Therefore, we do not treat the auxiliary outside signal x as part of the information design problem. The regulator’s problem is again to design the test s so as to maximize investors’ joint utility.

The Bank Run Game: All impatient investors have a dominant strategy to withdraw at date 1, and the bank repays them in cash. If a measure ℓ of patient consumers also withdraw, a patient investor who withdraws at $t = 1$ obtains \$1, for utility $u(1) = 0$. A patient investor who waits until $t = 2$ obtains $Re^{-\kappa\ell}$, for utility $u\left(Re^{-\kappa\ell}\right) = \theta + \eta - \kappa\ell$. Therefore, patient investor i prefers to wait if

$$E[\theta + \eta|s, m, x_i] \geq \kappa\ell,$$

Equivalently, using the independence of θ and η , we can write

$$z + \rho_i \geq \kappa\ell,$$

where

$$\rho_i = E[\eta|x_i] = \frac{\beta x_i}{\alpha + \beta}.$$

We look for cutoff equilibria such that investor i withdraws if and only if $z + \rho_i \leq \rho^*$. In such an equilibrium, suppose that i is the an investor who is just indifferent between withdrawing and waiting, i.e., that i receives a signal x_i that ensures the posterior belief $\rho_i = \rho^*$.

Investor i expects that the mass of investors withdrawing is

$$\begin{aligned} E[\ell|x_i] &= Pr[\rho_j \leq \rho_i|x_i] \\ &= Pr\left[\frac{\beta}{\alpha + \beta}x_j \leq \rho_i|x_i\right] \\ &= Pr\left[x_j \leq \rho_i\left(1 + \frac{\alpha}{\beta}\right)|x_i\right] \end{aligned}$$

Notice also that

$$x_j|x_i \sim N\left(\rho_i, \frac{\alpha + 2\beta}{\beta(\alpha + \beta)}\right)$$

implying that

$$\begin{aligned} E[\ell|x_i] &= \Phi\left(\sqrt{\frac{\beta(\alpha+\beta)}{\alpha+2\beta}} \frac{\alpha}{\beta} \rho_i\right) \\ &= \Phi(\sqrt{\gamma} \rho_i), \end{aligned}$$

where we have defined

$$\begin{aligned} \gamma &= \frac{\alpha^2 \beta (\alpha + \beta)}{\beta^2 (\alpha + 2\beta)} \\ &= \frac{\alpha^2 (\alpha + \beta)}{\beta (\alpha + 2\beta)} \end{aligned}$$

To complete the verification of equilibrium, we require that investor i is indeed indifferent between withdrawing and waiting. This requirement translates to $z + \rho_i = \kappa E[\ell|x_i]$ or, equivalently,

$$\kappa \Phi(\sqrt{\gamma} \rho^*) = z + \rho^*. \quad (26)$$

We now focus on the noiseless limit of this model, which is defined as follows: Consider a sequence of parameter values (α, β) , describing the precisions of the idiosyncratic component η of bank returns and the noise in investors' private signals. Assume that along this sequence we have $\alpha, \beta \rightarrow \infty$, so that in the limit, η collapses to a point mass at zero and investors have perfect information about η . Assume also that $\gamma \rightarrow 0$, which requires that β diverges to infinity at a faster rate than α .

In the limit as $\gamma \rightarrow 0$, (26) implies that $\rho^* \rightarrow \frac{\kappa}{2} - z$. Moreover, as $\alpha, \beta \rightarrow \infty$, investors have near-perfect information, so that the distribution of each ρ_i collapses to a point mass at zero. Hence, there is a run in the noiseless limit whenever

$$\begin{aligned} \rho_i &\leq \frac{\kappa}{2} - z \\ \Leftrightarrow z &\leq \frac{\kappa}{2} \equiv t \end{aligned}$$

Hence, we have derived the specification of bank runs in Example 1 as the unique equilibrium in the noiseless limit of the bank run game.

Regulator's Objective Function: The regulator wants to maximize the joint utility of depositors. Here, we can ignore the utility of impatient investors, which is equal to $u(c_1) = u(1) = 0$ regardless of whether there is a bank run. The utility of a representative patient investor if there is a run is also $u(1) = 0$. If there is no run, then it is given by

$u(R) = \theta + \eta \rightarrow \theta$ in the noiseless limit. Hence, the joint utility of all investors is proportional to

$$\theta \cdot Pr[\text{bank run}] = \theta \cdot (1 - \lambda(z)),$$

where $\lambda(z) = 1\{z < t\}$. Hence, we have derived the objective function in Equation (2) in the paper, with $\lambda(z)$ defined as in Example 1.