



Stress tests and information disclosure [☆]

Itay Goldstein ^{a,*}, Yaron Leitner ^b

^a Finance Department, Wharton School, University of Pennsylvania, 2253 Steinberg-Dietrich Hall, 3620 Locust Walk, Philadelphia, PA 19104, United States

^b Research Department, Federal Reserve Bank of Philadelphia, Ten Independence Mall, Philadelphia, PA 19106, United States

Received 15 August 2017; final version received 24 May 2018; accepted 27 May 2018

Available online 31 May 2018

Abstract

We study an optimal disclosure policy of a regulator that has information about banks (e.g., from conducting stress tests). In our model, disclosure can destroy risk-sharing opportunities for banks (the Hirshleifer effect). Yet, in some cases, some level of disclosure is necessary for risk sharing to occur. We provide conditions under which optimal disclosure takes a simple form (e.g., full disclosure, no disclosure, or a cutoff rule). We also show that, in some cases, optimal disclosure takes a more complicated form (e.g., multiple cutoffs or nonmonotone rules), which we characterize. We relate our results to the Bayesian persuasion literature.

© 2018 Elsevier Inc. All rights reserved.

JEL classification: D82; G01; G28

Keywords: Bayesian persuasion; Optimal disclosure; Stress tests; Bank regulation; Adverse selection

1. Introduction

In the new era of financial regulation following the crisis of 2008, central banks around the world will conduct periodic stress tests for financial institutions to assess their ability to withstand

[☆] The views expressed in this paper are those of the authors and do not necessarily reflect those of the Federal Reserve Bank of Philadelphia or the Federal Reserve System.

* Corresponding author.

E-mail addresses: itayg@wharton.upenn.edu (I. Goldstein), yaronleitner@gmail.com (Y. Leitner).

future shocks. An important aspect of this regulation is that the results of these stress tests are meant to be disclosed publicly. A key question that occupies policymakers and bankers is whether such disclosure is indeed optimal and, if so, at what level of detail.¹

A classic concern about disclosure is based on the Hirshleifer effect (Hirshleifer, 1971). According to the Hirshleifer effect, disclosure might be harmful because it reduces future risk-sharing opportunities for economic agents. This is indeed a relevant concern in the context of banks and stress tests. A vast literature (e.g., Allen and Gale, 2000) studies risk-sharing arrangements among banks. If banks are exposed to random liquidity shocks, they will create arrangements among themselves or with outside markets to insure against such shocks. Public information about the state of each individual bank and its ability to withstand future shocks could limit these hedging opportunities, thereby generating a welfare loss.

Given this logic, one would think that no disclosure is desired. Yet, during the crisis, interbank markets were not performing well, and there was a sense that some disclosure was necessary to prevent a breakdown in financial activity. Such a breakdown can occur when market participants have asymmetric information (e.g., Akerlof, 1970) but can also occur when market participants share the same information. For example, in Leitner (2005), risk-sharing arrangements among banks can break down when the aggregate endowment in the banking system is expected to be low.

Hence, disclosure involves a tradeoff and may be desirable in some circumstances but not in others. Indeed, this is apparent in the choices of policymakers during the crisis. While the Federal Reserve revealed the results of its stress tests, it did not reveal the identities of banks that used its special lending facilities.²

We set up a simple stylized model that captures the two forces above. In our model, disclosure can destroy risk-sharing opportunities for banks. Yet, there are cases in which some level of disclosure is necessary for risk sharing to occur. We study how optimal disclosure looks like in this setting under different circumstances. We distinguish between two cases, as described below. In the first case, the entire tradeoff originates from risk-sharing concerns, which provide the cost and benefit of disclosure. In the second case, we add another force: the bank has private information. We show that, in general, this force pushes for more disclosure.

In the model, banks suffer losses if their capital falls below a certain level. Part of a bank's capital can be forecasted based on current analysis and will become clear to the regulator examining the bank. However, there are also future shocks that cannot be forecasted. Banks can engage in risk sharing to guarantee that their capital does not fall below the critical level. The regulator sets a disclosure policy to minimize expected losses in the banking system.

We first consider an environment in which the information discovered by the regulator is not already known to the bank. We show that if banks are perceived, on average, to have capital above the critical level ("normal" times), it is possible to achieve risk sharing without any disclosure; so, the regulator does not need to disclose anything. Consistent with the Hirshleifer effect, disclosure can even be harmful because it can prevent optimal risk-sharing arrangements from taking place. However, if banks are perceived, on average, to have capital below the critical level ("bad" times), then risk-sharing arrangements that insure them against falling below that level cannot arise without some disclosure. In this case, optimal disclosure is in the spirit of the Bayesian persuasion literature, as in Kamenica and Gentzkow (2011).

¹ The debate over this question is illustrated in "Lenders Stress over Test Results," *Wall Street Journal*, March 5, 2012.

² Gorton (2015) provides more examples of suspension of information during a crisis.

Specifically, it is optimal to separate banks into two groups. The first group includes all the banks with forecastable capital above the critical level as well as some banks with forecastable capital below the critical level, such that, on average, the group's forecastable capital equals the critical level. The second group includes all the other banks. Banks in the first group engage in risk sharing, while banks in the second group do not. The regulator can implement this type of disclosure by assigning a high score to banks in the first group and a low score to banks in the second group.

Interestingly, the optimal disclosure rule is not necessarily monotone; i.e., it is not always the case that banks below a certain threshold receive the low score and banks above the threshold receive the high score. There is a gain and a cost from giving a bank the high score. The gain is enabling the bank to participate in risk sharing, thereby preventing a welfare-decreasing drop in capital. The cost is that giving the high score to a weak bank reduces the average capital in the group, thereby preventing other weak banks from receiving that score. The allocation of banks into the high-score group depends on the gain-to-cost ratio, which does not always generate a monotone rule; it depends on the distribution of shocks that banks are exposed to. We provide conditions under which the disclosure rule is monotone.

The second environment we consider is one in which the information discovered by the regulator is already known to the bank itself but not to other market participants. This case introduces a new important feature. Banks with a forecastable level of capital that is significantly above the critical level will agree to participate in a risk-sharing agreement with other banks only if the average forecastable capital for the group is sufficiently high. This participation constraint brings new results. First, under some circumstances, some disclosure is necessary even if, on average, banks are perceived to have capital above the critical level. Second, it may no longer be possible to implement optimal disclosure by separating banks into just two groups.

We show that, in general, optimal disclosure separates banks into multiple groups. As before, one group includes only banks that are below the critical capital level, and these banks are shunned from risk-sharing arrangements. Each of the other groups pools together banks below the critical level with a bank above the critical level to enable risk sharing. Multiple groups are required to accommodate the different reservation utilities of different banks above the critical level of capital.

Interestingly, in this environment, nonmonotonicity becomes a general feature of optimal disclosure rules. When considering banks below the critical level of capital, it turns out that the stronger ones are pooled with a bank that has a level of capital only slightly above the critical level, while the weaker ones are pooled with a bank that has a level of capital significantly above the critical level. As we show in this paper, the increase in cost from pooling with a moderately strong bank to pooling with a very strong bank is not significant for the weakest banks but is significant for the moderately weak banks. This leads to the nonmonotonicity result.

Nonmonotone rules can lead to outcomes in which weaker banks end up with higher equilibrium payoffs compared to stronger banks. This is problematic if banks can affect the value of their assets by freely disposing assets (Innes, 1990). To address this issue, we enrich our model by adding a constraint that stronger banks must end up with higher equilibrium payoffs. We also add a constraint that the regulator cannot randomize; i.e., the regulator must follow a deterministic rule. This can be viewed as a practical constraint on regulators. We show that once we add these two constraints, optimal disclosure rules become monotone and generally involve two thresholds. Banks below the lower threshold are excluded from risk sharing, those in the middle are pooled together in the same risk-sharing arrangement, and those above the higher threshold are pooled together in a different risk-sharing agreement (or agreements) to reflect their higher reservation

utilities. Interestingly, under some circumstances, full disclosure emerges as the unique optimal disclosure rule.

We discuss a few extensions, such as the case in which the regulator takes into account externalities that banks impose on the rest of society, or the case in which the regulator can inject money to banks. We illustrate that, in some cases, it is optimal to inject money not only to weak banks but also to strong banks so that the market cannot distinguish between the two. The model insights can also be used to analyze optimal capital injection programs when banks cannot be forced to participate. In the context of credit rating agencies, our model suggests that coarse ratings, as well as low types receiving high ratings, may be a feature of a socially optimal outcome.

2. Related literature

Our paper is related to different strands of the literature. The first strand is the literature on the disclosure of regulatory information in financial markets. This literature is reviewed in Goldstein and Sapra (2013) and Leitner (2014),³ who show that, while disclosure can enhance market discipline, disclosure can also create problems, such as reducing the regulator's ability to collect information from banks (Prescott, 2008; Leitner, 2012), reducing the regulator's ability to learn from market prices (Bond and Goldstein, 2015), inducing bank managers to window dress, or leading economic agents to put too much weight on public signals (Morris and Shin, 2002; Angeletos and Pavan, 2007). Other papers study how the regulator's disclosure policy is affected by reputational concerns (Morrison and White, 2013; Shapiro and Skeie, 2015), commitment issues (Bouvard et al., 2015), and fiscal capacity (Faria-e-Castro et al., 2017).⁴

Our paper analyzes a different tradeoff, which originates from risk-sharing concerns and which is based on the idea that, while disclosure reduces risk-sharing opportunities as in Hirshleifer (1971), it is sometimes necessary to prevent risk-sharing markets from breaking down.

Andolfatto et al. (2014) study risk sharing in a monetary search model. In their model, information has no social value, and it is optimal to disclose information only when this is done to prevent individuals from wastefully acquiring (this same information) on their own. Diamond (1985) studies risk sharing in a competitive rational expectations model. In his setting, optimal disclosure not only reduces the social cost of acquiring information but also prevents investors from acquiring different pieces of information on their own; so, disclosure enhances trade by reducing information asymmetries among investors. Our setting does not involve information acquisition or information production by market participants.

Dang et al. (2017) use the insights from the Hirshleifer effect to explain bank opaqueness, while Alvarez and Barlevy (2015) show in a model of financial contagion and information spillovers that, under some conditions, forcing banks to disclose balance sheet information is undesirable. Both papers are silent about the regulator's optimal disclosure policy. Monnet and Quintin (2017) study an information design problem faced by long-term investors that can liquidate their investment by either scrapping the project or selling it to other investors at cash-in-the-market prices. The tradeoff in their setting is between the Hirshleifer effect, which pushes for secrecy, and the desire to obtain information to make efficient scrapping decisions. The in-

³ See also Goldstein and Yang (2017) for a recent survey on information disclosure in financial markets.

⁴ Gick and Pausch (2014) study optimal disclosure as a Bayesian persuasion problem in which the regulator's preferred outcome is that some depositors withdraw their funds (so that banks do not engage in excessive risk taking) and some depositors keep their funds (to avoid a run).

sights from the Hirshleifer effect are also present in Bouvard et al. (2015), who study how the regulator's disclosure policy affects the possibility of bank runs. In their model, banks are passive. In contrast, in our model, banks decide whether to engage in risk sharing and can refuse to participate in risk-sharing arrangements proposed by the regulator. In a different context, Marin and Rahi (2000) use the insights from the Hirshleifer effect to provide a theory of market incompleteness.

Another related literature studies optimal disclosure by certification intermediaries. Lizzeri (1999) shows that, to extract more rents, a monopolist intermediary will commit to reveal only the minimum information that is required for an efficient exchange. Kartasheva and Yilmaz (2013) extend Lizzeri (1999) by adding outside options to firms as well as information asymmetries among potential buyers. In both papers, full disclosure achieves the socially efficient outcome, as there are no risk-sharing concerns. Instead, in our setting, the socially efficient outcome typically involves pooling.

Our paper also relates to the literature on whether it is optimal to force firms to disclose information (e.g., Admati and Pfleiderer, 2000; Fishman and Hagerty, 2003).⁵ Our paper illustrates a case in which it is socially optimal to restrict information flow from firms. A strong firm ignores the fact that revealing information destroys risk-sharing opportunities for weak firms, but the regulator takes this negative externality into account.⁶

Another strand is the literature on Bayesian persuasion and information design, which is reviewed in Bergemann and Morris (2017). In our basic setting, the regulator discloses information to persuade a competitive market to make a sufficiently high price offer to the bank so that the bank's capital does not fall below the critical level. If the bank knows its type, the regulator also needs to take into account the bank's participation constraint, namely, the bank should be willing to accept the offer. In a different context, Calzolari and Pavan (2006a, 2006b) and Dworzak (2017) also study settings in which disclosure is followed by a game in which an uninformed receiver (the market in our setting) makes a price offer to a third party that has the information that the sender has, but in contrast to our paper, in those papers the sender must first elicit information from the third party.

On a technical level, our basic setting maps into a general Bayesian persuasion problem (Kamenica and Gentzkow, 2011) with one sender and one receiver. However, as we explain in Section 7, the standard concavification approach has limited applicability in characterizing optimal disclosure rules in our setting. The characterization and closed-form solutions we provide for the case in which the bank knows its type are completely new to this literature. The solution for the case in which the bank does not know its type and the condition for the optimality of a simple cutoff rule using the gain-to-cost ratio are also new.

A recent body of literature (e.g., Goldstein and Huang, 2016; Inostroza and Pavan, 2017; Mathevet et al., 2017) examines persuasion with multiple receivers, focusing on issues that arise when actions need to be coordinated or when receivers are heterogeneously informed. These issues do not arise in our setting, which, in its reduced form, has only one receiver.

Somewhat related, in a cheap talk setting, Chakraborty and Harbaugh (2007) show that, under some conditions, a sender can influence a receiver by ranking the different alternatives, and that the sender can benefit by providing partial ranking (i.e., pooling a few alternatives together).

⁵ For surveys, see Verrecchia (2001) and Beyer et al. (2010).

⁶ In a different setting, Fishman and Hagerty (1990) show that when an informed seller can verifiably disclose some, but not all, of its information, it may be optimal to restrict the *type* of information that can be disclosed, so that the seller cannot cherry-pick positive information.

Finally, the empirical literature shows that stress tests convey new information to market participants, and that this information affects market prices.⁷ However, it is hard to draw implications from the empirical literature regarding the cost and benefit of disclosure. One difficulty is that, in practice, disclosure takes various forms and is often coupled with other government interventions. For example, the stress tests conducted in the U.S. in 2009 were coupled with a government guarantee to make capital available to banks that were unable to raise private capital. The main purpose of disclosure at that time was probably to convince the market that these government guarantees were credible.

3. Model

Economic environment. There is a bank, a regulator, and a perfectly competitive market. The bank has an asset, which yields a random cash flow $\tilde{\theta} + \tilde{\varepsilon}$, where $\tilde{\theta}$ is the bank's type and $\tilde{\varepsilon}$ is the bank's idiosyncratic risk, which is independent of its type. The bank can sell its asset in the market for price x , which is derived endogenously. Everyone is risk neutral, and the risk-free rate is normalized to be zero percent. Therefore, x is the expected value of $\tilde{\theta} + \tilde{\varepsilon}$, conditional on the information available to the market. We use z to denote the bank's final cash holdings. Hence, $z = x$ if the bank sells the asset and $z = \tilde{\theta} + \tilde{\varepsilon}$ if the bank keeps the asset. The bank's final payoff is

$$R(z) = \begin{cases} z & \text{if } z < 1 \\ z + r & \text{if } z \geq 1, \end{cases} \quad (1)$$

for a parameter $r > 0$.

The payoff function (1) captures the general idea that a bank derives some gains when its cash holdings are (weakly) above some threshold. One can think of several motivations: (1) The bank has a project that yields a positive net present value r but requires a minimum level of investment. For various reasons (e.g., the project's cash flows are nonverifiable), the bank cannot finance the project if it does not have sufficient cash in hand. (2) The bank has a debt liability of 1. Not paying it leads to loss of future income r . (3) The bank faces a run if its cash holdings fall below some threshold.

Our results do not hinge on the particular specification for $R(z)$ above. For example, our results extend to the case in which r depends on the bank's type (we discuss this more later). The results also extend to other payoff functions that exhibit discontinuity, such as assuming that the bank obtains az for some $a \in [0, 1)$ if $z < 1$, and $z + r$ if $z \geq 1$ (where r can be set to zero). The case $a = r = 0$ may best capture the idea that when the asset value falls below some threshold, there is a bank run and the bank is left with nothing. The key to all of these specifications is that disclosure of high types allows them to create surplus they otherwise wouldn't. In our specification, this benefit is due to the discontinuity we assume,⁸ although there may be other ways of modelling it. We discuss the role this discontinuity plays in our model in Section 5.5 and in Section 7.

The bank chooses whether to keep its asset or sell it in the market. The bank does so in a way that maximizes its expected final payoff $R(z)$, conditional on the information available to it. Selling the asset at a price of at least 1 guarantees that the bank's cash holdings will not fall below the threshold and, hence, can be thought of as a simple form of insurance or risk sharing,

⁷ See, for example, Petrella and Resti (2013), Morgan et al. (2014), and Flannery et al. (2017).

⁸ A similar discontinuity in payoffs appears in Leitner (2005) and in Elliott et al. (2014).

in which the bank replaces a stochastic cash flow with a deterministic cash flow.⁹ The nature of our model continues to hold for other forms of risk sharing, such as the case in which banks enter into more complicated derivative contracts, or the case in which banks share risk among themselves (see Section 7).

The bank's type $\tilde{\theta}$ is drawn from a finite set $\Theta \subset \mathbb{R}$ according to a probability distribution function $p(\theta) = \Pr(\tilde{\theta} = \theta)$. The idiosyncratic risk $\tilde{\varepsilon}$ is drawn from a continuous cumulative distribution function F that satisfies $E(\tilde{\varepsilon}) = 0$. The probability structure (i.e., the functions p and F) is common knowledge.

We denote the types in Θ by $\theta_{\max} = \theta_1 > \theta_2 > \dots > \theta_m = \theta_{\min}$. We assume that there are $k \geq 1$ types at or above 1, with at least one type strictly above 1 ($\theta_{\max} > 1$).

Assumption 1. $F(1 - \theta_{\min}) < 1$ and $F(1 - \theta_{\max}) > 0$.

Assumption 1 implies that even for the lowest type, there is a positive probability that the asset cash flow will be above 1, and even for the highest type, there is a positive probability that the asset cash flow will be below 1. It is easy to relax this assumption (e.g., the discussion of the case $\tilde{\varepsilon} = 0$ in Section 4 and Section 5).

The focus of this paper is on the optimal disclosure policy of a regulator who has information about the bank. Specifically, the regulator observes the realization of $\tilde{\theta}$ (denoted by θ) and discloses information according to a disclosure rule, which is chosen before θ is observed by the regulator. The regulator can commit to the chosen disclosure rule and his objective is to maximize expected total surplus. Since the market breaks even on average, maximizing expected total surplus is the same as maximizing the bank's expected payoff across the different types.

The market does not observe θ . As for the bank, we focus on two cases: (1) The bank does not observe θ . (2) The bank observes θ . In both cases, we assume that no one observes the realization of $\tilde{\varepsilon}$ (denoted by ε), which is residual noise. The first case captures the idea that the regulator may have some information advantage relative to banks.¹⁰ The second case captures the idea that the regulator and the bank share the same information, which is unobservable to other market participants. Throughout most of the analysis, we assume that the bank cannot affect what the regulator observes. We relax this assumption in Section 6.

Disclosure rules. A *disclosure rule* is defined by a finite set of "scores" S and a function that maps each type to a distribution over scores.¹¹ We use $g(s|\theta)$ to denote the probability, according to the disclosure rule, that the regulator assigns a score $s \in S$ when he observes type θ . That is, $g(s|\theta) = \Pr(\tilde{s} = s|\tilde{\theta} = \theta)$. Of course, for every $\theta \in \Theta$, the following has to hold: $\sum_{s \in S} g(s|\theta) = 1$.

To gain intuition on how disclosure rules work, note that full disclosure is obtained when, for every type θ , the regulator assigns some score $s_\theta \in S$ with probability 1, such that $s_\theta \neq s_{\theta'}$ if $\theta \neq \theta'$. No disclosure is obtained when the regulator assigns the same distribution over scores to all types (e.g., each type obtains the same score).

⁹ Unlike the bank, the market is not affected by the discontinuity in payoffs and only gets $\theta + \varepsilon$ if it buys the asset. Hence, this transfer of risk can increase surplus.

¹⁰ This might be the case if the stress test involves an assessment of bank exposure to the state of other banks, which is known to the regulator but not to the individual banks.

¹¹ In our setting, the optimal disclosure rule can be implemented with a finite number of scores. Hence, there is no loss of generality in assuming that S is finite.

For use below, denote $\mu(s) = E[\tilde{\theta} + \tilde{\varepsilon} | \tilde{s} = s]$. This is the expected value of the bank's asset, conditional on the bank obtaining score s . Since $E(\tilde{\varepsilon}) = 0$ and $\tilde{\varepsilon}$ is independent of $\tilde{\theta}$ (and hence, \tilde{s}), we obtain that

$$\mu(s) = E[\tilde{\theta} | \tilde{s} = s] = \sum_{\theta \in \Theta} \theta \Pr(\tilde{\theta} = \theta | \tilde{s} = s) = \frac{\sum_{\theta \in \Theta} \theta p(\theta) g(s | \theta)}{\sum_{\theta \in \Theta} p(\theta) g(s | \theta)}, \quad (2)$$

where the last equality follows from Bayes' rule.

Sequence of events. The sequence of events is as follows:

1. The regulator chooses a disclosure rule (S, g) and publicly announces it.
2. The bank's type θ is drawn and observed by the regulator. In case 2, θ is also observed by the bank.
3. The regulator assigns the bank a score s according to the disclosure rule and publicly announces s .
4. The market offers to purchase the asset at a price $x(s)$.
5. The bank chooses whether to keep its asset or sell it for a price $x(s)$.
6. The residual noise ε is realized. As a result, the bank's cash holdings z and the bank's final payoff $R(z)$ are determined.

The regulator's disclosure rule and assigned score specify a game between the bank and the market. We focus on perfect Bayesian equilibria of this game. Specifically, the bank chooses whether to sell or keep its asset to maximize its expected payoff conditional on its information. The market chooses a price $x(s)$, such that it breaks even, in expectation, conditional on the publicly announced score, taking as given the bank's equilibrium strategy.¹² We assume that, if the bank is indifferent between selling and not selling, it sells, and if the market is indifferent between two prices, it offers the highest price. This is the outcome preferred by the regulator. The regulator chooses a disclosure rule that maximizes the bank's expected payoff across all types, taking as given the equilibrium strategies of the market and the bank. Note that assigning scores is equivalent to recommending prices to the market. We discuss the assumption that the regulator can commit to a disclosure rule, as well as other regulator's objective functions, in Section 7.

4. Bank does not observe its type

We start with the case in which the bank does not observe θ . In this case, the bank's decision to sell does not convey any additional information to the market beyond the information contained in the score s . Consequently, the market sets a price $x(s) = \mu(s)$. It then follows from the payoff structure in (1) that:

Lemma 1. *In equilibrium, the bank sells the asset if, and only if, it obtains a score s such that $\mu(s) \geq 1$.*

The proof of Lemma 1 and all other proofs are in the Appendix. The idea behind Lemma 1 is simple. If $\mu(s) \geq 1$, selling guarantees that the bank's cash holding will be at least 1 and the bank will obtain r , but keeping the asset could lead to cash holdings below 1 (by the second part of Assumption 1). Hence, the bank acts like a risk-averse agent and is happy to replace the asset's

¹² Formally, we assume Bertrand competition among at least two market participants, where the payoff for each participant is $\theta + \varepsilon - x(s)$ if it purchases the asset, and zero otherwise.

random cash flow with its expected value. If, instead, $\mu(s) < 1$, the bank prefers to keep the asset because if the bank sells the asset at a price below 1, the bank’s cash holdings will surely be below 1; but, if the bank keeps the asset, there is a positive probability that the asset’s cash flow will turn out to be more than 1 (by the first part of Assumption 1). In this case, the bank acts like a risk-loving agent.

The expected payoff for a bank of type θ , given disclosure rule (S, g) , is then

$$u(\theta) \equiv \sum_{s:\mu(s)<1} [\theta + r \Pr(\tilde{\varepsilon} \geq 1 - \theta)]g(s|\theta) + \sum_{s:\mu(s)\geq 1} [\mu(s) + r]g(s|\theta). \tag{3}$$

The first term represents the cases in which the bank keeps the asset. The second term represents the cases in which the bank sells the asset.

The regulator’s problem is to choose a disclosure rule (S, g) to maximize the bank’s ex-ante expected payoff $\sum_{\theta \in \Theta} p(\theta)u(\theta)$.

Lemma 2. *The regulator’s problem reduces to choosing a disclosure rule (S, g) to maximize*

$$\sum_{\theta \in \Theta} p(\theta) \left[\Pr(\tilde{\varepsilon} < 1 - \theta) \sum_{s:\mu(s)\geq 1} g(s|\theta) \right]. \tag{4}$$

The term $\sum_{s:\mu(s)\geq 1} g(s|\theta)$ in the objective function (4) is the probability that a bank of type θ sells its asset. The term $\Pr(\tilde{\varepsilon} < 1 - \theta)$ represents the social gain from having type θ sell its asset: type θ can guarantee that its cash holdings are at least 1 even if the asset cash flow turns out to be less than 1 (when $\tilde{\varepsilon} < 1 - \theta$).

By Lemma 2, disclosure affects the regulator’s payoff only by affecting the probability that each type sells its asset. Hence, we can focus, without loss of generality, on disclosure rules that assign at most two scores, s_1 and s_0 , such that $\mu(s_1) \geq 1$ and $\mu(s_0) < 1$. The bank sells its asset if, and only if, it obtains score s_1 . Formally:

Lemma 3. *Consider a disclosure rule (S, g) and a disclosure rule (\hat{S}, \hat{g}) , defined by $\hat{S} = \{s_0, s_1\}$, $\hat{g}(s_1|\theta) = \sum_{s:\mu(s)\geq 1} g(s|\theta)$, and $\hat{g}(s_0|\theta) = 1 - \sum_{s:\mu(s)\geq 1} g(s|\theta)$. Then the probability that type θ sells its asset is the same under both rules. The value of the regulator’s objective function is also the same under both rules.*

We refer to scores s_1 and s_0 as “high” and “low,” respectively, and denote the probability that type θ obtains the high score s_1 by $h(\theta)$. That is, $h(\theta) \equiv \hat{g}(s_1|\theta)$.

Lemma 4. *The regulator’s problem reduces to finding a function $h : \Theta \rightarrow [0, 1]$ to maximize*

$$\sum_{\theta \in \Theta} p(\theta) \Pr(\tilde{\varepsilon} < 1 - \theta)h(\theta), \tag{5}$$

subject to

$$\sum_{\theta \in \Theta} p(\theta)(\theta - 1)h(\theta) \geq 0. \tag{6}$$

The objective function (5) follows from Lemma 2. Constraint (6) follows from $\mu(s_1) \geq 1$, using equation (2). The constraint says that the average θ conditional on obtaining the high score must be at least 1.

We can think of constraint (6) as a persuasion constraint. The regulator wants to persuade the market to purchase the bank’s asset at a price $x \geq 1$. For this, the average type that sells the asset must be at least 1. Essentially, by giving a high score, the regulator implements a cross subsidy from types with $\theta > 1$ to types with $\theta < 1$, so a high type sells its asset for less than the true value, and a low type sells its asset for more than the true value. This is beneficial because more types can ensure that their cash holdings are at least 1.

While there is no real transfer of resources in our basic setting, it might be useful to think of constraint (6) also as a resource constraint, where high types provide resources to low types. For example, this would be the case if multiple banks share risk among themselves, as discussed in Section 7. Then the banks that ended up with high realizations of cash flows would indeed transfer resources to those that ended up with low realizations.

The solution to the regulator’s problem is as follows. If $E(\tilde{\theta}) \geq 1$, assigning $h(\theta) = 1$ for every $\theta \in \Theta$ satisfies constraint (6) and, hence, is optimal. Otherwise, if $E(\tilde{\theta}) < 1$, it is impossible to assign $h(\theta) = 1$ for every $\theta \in \Theta$, and so, constraint (6) is binding. That is, the average type getting the high score is exactly 1. The optimal disclosure rule then has to determine the probability with which each type gets the high score. This depends on comparing the “gain-to-cost ratio” from increasing $h(\theta)$ for different types. The gain for type θ is the term $\Pr(\tilde{\varepsilon} < 1 - \theta)$ in the objective function (5). The cost is that type θ requires resources in the amount $1 - \theta$, as in equation (6). For a type $\theta \geq 1$, the cost is negative (or zero). Hence, it is optimal to assign $h(\theta) = 1$ for every $\theta \geq 1$. In contrast, for a type $\theta < 1$, the cost is positive. In this case, the gain-to-cost ratio from increasing $h(\theta)$ is

$$G(\theta) \equiv \frac{\Pr(\tilde{\varepsilon} < 1 - \theta)}{1 - \theta}. \tag{7}$$

It then follows from the linearity of the problem that, for types below 1, it is optimal to set a cutoff G^* such that types with a gain-to-cost ratio above the cutoff are assigned $h(\theta) = 1$, and types with a gain-to-cost ratio below the cutoff are assigned $h(\theta) = 0$. The optimal G^* is the lowest cutoff possible that satisfies constraint (6). For types with a gain-to-cost ratio that equals G^* , the probability of obtaining the high score could be between 0 and 1 and is set such that constraint (6) is satisfied with equality.

The following proposition summarizes the optimal disclosure rule.

Proposition 1. *When the bank does not observe its type, the optimal disclosure rule is such that*

1. *If $E(\tilde{\theta}) \geq 1$, then $h(\theta) = 1$ for every $\theta \in \Theta$.*
2. *If $E(\tilde{\theta}) < 1$, then*

$$h(\theta) = \begin{cases} 1 & \text{if } \theta \geq 1 \text{ or if } \theta < 1 \text{ and } G(\theta) > G^* \\ 0 & \text{if } \theta < 1 \text{ and } G(\theta) < G^*, \end{cases} \tag{8}$$

where G^* is the lowest $G \in \{G(\theta)\}_{\theta < 1}$ that satisfies $\sum_{\theta \geq 1} p(\theta)(\theta - 1) + \sum_{\theta < 1: G(\theta) > G} p(\theta)(\theta - 1) \geq 0$, if such G exists; otherwise, $G^* \equiv \max_{\theta < 1} G(\theta)$. If $G(\theta) = G^*$, then $h(\theta) \in [0, 1)$, such that (6) is satisfied with equality.

An interesting question is whether and when full disclosure is optimal, and whether and when no disclosure is optimal. If $E(\tilde{\theta}) \geq 1$, we know from Proposition 1 that every type must sell its asset with probability 1. The regulator can implement this outcome by giving all types the same score, i.e., with no disclosure. There are other ways to implement the optimal outcome, assigning more than one score such that the average θ of types receiving each score is at least 1.

In the special case $\theta_{\min} \geq 1$, the regulator can even assign a different score to each type, i.e., provide full disclosure.¹³ In contrast, if $E(\tilde{\theta}) < 1$, the regulator must assign at least two scores. Some disclosure is necessary because, without disclosure, the price would be less than 1 and no type would sell its asset. Yet, full disclosure is suboptimal because under full disclosure, only types above 1 sell their assets, while under the optimal disclosure rule, some types that are below 1 also sell their assets. Hence, partial disclosure that pools together types above 1 with some types below 1 is the only way to achieve the optimal outcome. The following corollary summarizes the results above.

Corollary 1. *When the bank does not observe its type, full disclosure achieves the optimal outcome if, and only if, $\theta_{\min} \geq 1$. No disclosure achieves the optimal outcome if, and only if, $E(\tilde{\theta}) \geq 1$. If $E(\tilde{\theta}) < 1$, then partial disclosure is the only way to achieve the optimal outcome.*

Another interesting question is under what conditions a simple cutoff rule with respect to θ is optimal. Interestingly, in the case of $E(\tilde{\theta}) < 1$ (summarized in the second part of Proposition 1), the types that obtain the low score are not necessarily the lowest. So, a simple cutoff rule that assigns the high score to high types and the low score to low types is not necessarily optimal. Intuitively, the gain from giving the high score is higher for lower types because low types are more likely to end up with low realizations of cash flow. That is, the numerator of (7) is decreasing in θ . But the cost of giving the high score to low types is also higher because low types require more resources. That is, the denominator of (7) is also decreasing in θ . Hence, it is unclear whether $G(\theta)$ is increasing or decreasing, or whether it is even monotone. The function $G(\theta)$, and hence the optimal disclosure rule, depends on the distribution of the idiosyncratic risk $\tilde{\varepsilon}$.

The optimal rule will involve a simple cutoff with respect to θ when $G(\theta)$ is increasing when $\theta < 1$.¹⁴ In this case, types above the cutoff obtain the high score with probability 1, and types below the cutoff obtain the low score with probability 1. A simple example in which this happens is when there is no idiosyncratic risk (i.e., $\varepsilon = 0$). Then, the gain-to-cost ratio is simply $G(\theta) = \frac{1}{1-\theta}$. More generally, a sufficient condition for obtaining a cutoff rule is that the cumulative distribution function of $\tilde{\varepsilon}$ satisfies Condition 1 below. This condition is satisfied by any cumulative distribution function that is concave on the positive region. Examples include a normal distribution and a uniform distribution (both with mean zero).

Condition 1. $F(\varepsilon)/\varepsilon$ is decreasing when $\varepsilon > 0$.

Corollary 2. *If $E(\tilde{\theta}) < 1$ and Condition 1 holds, the optimal disclosure rule involves a cutoff, such that types below the cutoff obtain the low score and types above the cutoff obtain the high score.*

Another case in which the optimal rule involves a simple cutoff with respect to θ is when r in the payoff function (1) depends on θ , according to some function $r(\theta)$ that is increasing in θ

¹³ In fact, any disclosure rule is optimal in this special case; but see also Corollary 6.

¹⁴ This condition implies that the regulator's virtual utility from assigning score s to type θ is a quasi-supermodular function, a condition that Mensch (2018) identifies as sufficient for a monotone disclosure rule to be optimal in a general Bayesian persuasion environment. In our setting, the virtual utility from assigning the low score is 0, and the virtual utility from assigning the high score is $\Pr(\tilde{\varepsilon} < 1 - \theta) - \lambda(1 - \theta)$, where $\lambda > 0$ is the Lagrange multiplier for constraint (6). Under the optimal disclosure rule, $\lambda = G^*$.

sufficiently strongly. This has a simple and intuitive economic interpretation: good banks have better investment opportunities in addition to having better assets in place. In this case, the gain from giving the high score is $r(\theta) \Pr(\tilde{\varepsilon} < 1 - \theta)$ and the gain-to-cost ratio is $r(\theta)G(\theta)$. So, no matter what shape $G(\theta)$ has, if $r(\theta)$ is increasing sufficiently strongly, the gain-to-cost ratio will be monotonically increasing, and the disclosure rule will look like a cutoff rule. For example, if $r(\theta) = \frac{1}{\Pr(\tilde{\varepsilon} < 1 - \theta)}$, then $r(\theta)G(\theta) = \frac{1}{1 - \theta}$, which is increasing in θ .

A cutoff rule will also be optimal in a variation of our model in which the bank can obtain the extra return r only if it sells its asset.¹⁵ For example, the bank’s investment opportunity may expire before the asset cash flows are obtained. In this case, the social gain from giving a high score is r , independent of type, and the resulting gain-to-cost ratio ($\frac{r}{1 - \theta}$) is increasing in type.

Finally, an example in which the optimal disclosure rule does not involve a simple cutoff, as in Corollary 2, is when $G(\theta)$ is decreasing when $\theta \leq \theta_{k+1}$. In this case, the optimal disclosure rule is nonmonotone in type. It includes a cutoff such that types below the cutoff and types above 1 obtain the high score, while types in the middle obtain the low score. A sufficient condition for this to happen is that $F(\varepsilon)/\varepsilon$ is increasing when $\varepsilon \geq 1 - \theta_{k+1}$.¹⁶

5. Bank observes its type

So far, we assumed that the bank does not observe its type. We showed that it is possible to implement the optimal disclosure rule with two scores, such that the regulator pools all the types that sell under the same score. In this section, we show that this conclusion may no longer be true when the bank observes its type. The difference is that now each type has a “reservation price,” i.e., a minimum price at which it is willing to sell. When different types have different reservation prices, the regulator may need to assign more than two scores to distinguish among them. We also discuss how the regulator should assign these multiple scores to low types that are pooled with high types.

5.1. Derivation of the regulator’s problem

We first derive the reservation price for each type. Define

$$\rho(\theta) = \begin{cases} \max\{1, \theta - r \Pr(\tilde{\varepsilon} < 1 - \theta)\} & \text{if } \theta \geq 1 \\ \min\{1, \theta + r \Pr(\tilde{\varepsilon} \geq 1 - \theta)\} & \text{if } \theta < 1. \end{cases} \tag{9}$$

Lemma 5. *In equilibrium, a bank of type θ agrees to sell at price x if and only if $x \geq \rho(\theta)$.*

Lemma 5 is derived as follows. If type θ keeps its asset, its expected payoff is $E[R(\theta + \tilde{\varepsilon})] = \theta + r \Pr(\tilde{\varepsilon} \geq 1 - \theta)$. If type θ sells at price x , its payoff is $R(x)$; i.e., it is x when $x < 1$ and $x + r$ when $x \geq 1$. Hence, type θ agrees to sell if, and only if, $R(x) \geq \theta + r \Pr(\tilde{\varepsilon} \geq 1 - \theta)$. In the proof, we show that this reduces to $x \geq \rho(\theta)$.

¹⁵ That is, the bank’s payoff is $E(\tilde{\theta} + \tilde{\varepsilon}) = E(\tilde{\theta})$, if it keeps the asset, and $R(x)$, if it sells the asset.

¹⁶ An example of a probability distribution function that satisfies the condition above is a truncated Cauchy distribution (Nadarajah and Kotz, 2006) on the interval $[-A, 0]$ minus its mean, where the lower bound A depends on the model parameters. Intuitively, for the sufficient condition above to hold, we must have $F(1 - \theta_{k+1}) < \frac{1 - \theta_{k+1}}{1 - \theta_{\min}} F(1 - \theta_{\min}) < \frac{1 - \theta_{k+1}}{1 - \theta_{\min}}$. So, if θ_{k+1} is close to 1, then $F(1 - \theta_{k+1}) \simeq 0$, and the probability distribution of $\tilde{\varepsilon}$ puts a low mass on negative values (and a high mass on positive values). To also satisfy $E(\tilde{\varepsilon}) = 0$, the distribution must have a fat left tail.




Fig. 1. Reservation prices (bold line).

We refer to $\rho(\theta)$ as type θ 's reservation price and denote $\rho_i = \rho(\theta_i)$. In the special case in which there is no idiosyncratic risk ($\tilde{\varepsilon} = 0$), the reservation price is $\rho(\theta) = \theta$ for every $\theta \in \Theta$, and the regulator cannot implement cross-subsidies from high types to low types because no type would be willing to sell at a price that reflects the average of a group that includes that type and lower types. In this case, every type $\theta \in \Theta$ ends up with a payoff $R(\theta)$ independently of the disclosure rule, and so, every disclosure rule is optimal.

The rest of our paper focuses on the more interesting case in which $\tilde{\varepsilon}$ has a nondegenerate distribution. In this case, the reservation price satisfies three properties, which we use later (see Fig. 1). First, $\rho(\theta)$ is increasing in θ . Second, $\rho(\theta) < \theta$ when $\theta > 1$. Third, $\rho(\theta) > \theta$ when $\theta < 1$. Intuitively, types above 1 are willing to sell below their true valuations to guarantee that their cash holdings do not fall below 1. This is the insurance motive. In contrast, types below 1 will agree to sell only above their true valuation, because if they sell for the true valuation, they lose the option value of ending up with cash holdings above 1.

The next lemma simplifies the analysis.

Lemma 6. *Under an optimal disclosure rule:*

1. Every type $\theta_i \geq 1$ sells its asset with probability 1.
2. If the highest type that obtains score s is $\theta_i \geq 1$, then every type sells its asset upon obtaining that score and the price satisfies $x(s) = \mu(s) \geq \rho_i$.
3. If the highest type that obtains score s is below 1, then every type keeps its asset upon obtaining that score.

The first part in Lemma 6 follows because if a type $\theta \geq 1$ did not sell its asset, the regulator could strictly increase type θ 's payoff, without affecting the payoffs of other types, by fully revealing θ 's type. Then, the market would offer to buy type θ 's asset at price θ , and type θ would accept the offer. Moreover, the price that the market would offer for the group of types that included θ initially would remain unchanged, because it is based on the average type with a given score that sells and not the average type with a given score.

The second part follows from the first part and the observation that the reservation price is increasing in θ . In particular, if the highest type that obtains score s is $\theta_i \geq 1$, we must have $x(s) \geq \rho_i$, so that type θ_i agrees to sell; and because the reservation price is increasing in θ , lower types also sell upon obtaining score s . Hence, the market prices the asset at the expected value given the score: $x(s) = \mu(s)$.

The third part reflects the fact that, if no type above 1 obtains score s , the price $x(s)$ must be less than 1. But then the bank will sell only if the price is strictly above the true value, which cannot be an equilibrium outcome because the market will overpay in expectation.

Lemma 6 implies that for each score that induces selling, the highest type that obtains the score must be at least 1. Moreover, the sale price must be at least as high as the reservation price of that type.

Since there are k types above 1, we can focus, without loss of generality, on disclosure rules that assign, at most, $k + 1$ scores. One score, which we denote by s_0 , is reserved for the types that do not sell their asset. The other k scores, which we denote by s_1, \dots, s_k , are reserved for the types that sell their asset, such that the highest type that obtains score s_i is θ_i .¹⁷

Formally, denote $S_i = \{s \in S : g(s|\theta_i) > 0 \text{ and } g(s|\theta) = 0 \text{ for every } \theta > \theta_i\}$. That is, S_i is the set of scores, such that the highest type that obtains each score in S_i is θ_i . Then:

Lemma 7. *If (S, g) is an optimal disclosure rule, then (\hat{S}, \hat{g}) , defined by $\hat{S} = \{s_0, s_1, \dots, s_k\}$, $\hat{g}(s_i|\theta) = \sum_{s \in S_i} g(s|\theta)$ for $i \in \{1, \dots, k\}$, and $\hat{g}(s_0|\theta) = 1 - \sum_{i=1}^k \hat{g}(s_i|\theta)$, is also optimal.*

For $i \in \{1, \dots, k\}$, denote the probability that type θ obtains score s_i by $h_i(\theta)$. We can write down the regulator’s problem as follows:

Proposition 2. *When the bank observes its type, the regulator’s problem reduces to choosing a set of functions $\{h_i : \Theta \rightarrow [0, 1]\}_{i=1, \dots, k}$ to maximize*

$$\sum_{\theta \in \Theta} p(\theta) \Pr(\tilde{\epsilon} < 1 - \theta) \sum_{i=1}^k h_i(\theta), \tag{10}$$

such that (11)–(13) hold:

$$\sum_{\theta \in \Theta} p(\theta)(\theta - \rho_i)h_i(\theta) \geq 0 \text{ for every } i \in \{1, \dots, k\}, \tag{11}$$

$$\sum_{i=1}^k h_i(\theta) \leq 1 \text{ for every } \theta \in \Theta, \tag{12}$$

$$h_i(\theta) = 0 \text{ for every } i \in \{1, \dots, k\} \text{ and } \theta > \theta_i. \tag{13}$$

The regulator chooses the probability $h_i(\theta)$ for every $\theta \in \Theta$ and $i \in \{1, \dots, k\}$. The objective function (10) is as in Lemma 4, noting that the probability that type θ sells its asset is $\sum_{i=1}^k h_i(\theta)$, rather than $h(\theta)$. Equation (11) is a generalization of constraint (6). Now there is a constraint for every score that induces selling. The constraint for score s_i (namely, constraint i) says that the average θ conditional on obtaining score s_i must be at least ρ_i . This constraint follows since $\mu(s_i) \geq \rho_i$ (and (2)). Equation (12) simply says that the probability that type θ sells its asset is at most 1. Equation (13) says that the highest type that obtains score s_i is type θ_i .

The problem in Proposition 2 is a linear programming problem. Because the feasible region is bounded and closed ($h_i(\theta) \in [0, 1]$) and is nonempty,¹⁸ a solution exists.

¹⁷ Note that in Section 4, when a bank does not know its type, everyone is willing to sell for the same price, and so, there are only 2 scores.

¹⁸ Setting $h_i(\theta) = 1$ if $\theta = \theta_i$, and $h_i(\theta) = 0$ if $\theta \neq \theta_i$, satisfies all the constraints.

5.2. Full disclosure vs. no disclosure

As in Corollary 1, full disclosure achieves an optimal outcome if, and only if, there are no types below 1. No disclosure achieves an optimal outcome if, and only if, $E(\tilde{\theta})$ is sufficiently high. However, the condition for no disclosure to be optimal is stricter than in Corollary 1, reflecting the reservation price of the highest type.

Corollary 3. *When the bank observes its type, no disclosure achieves the optimal outcome if, and only if, $E(\tilde{\theta}) \geq \rho_1$. Full disclosure achieves the optimal outcome if, and only if, $\theta_{\min} \geq 1$.*

Note that when r is higher, the reservation price ρ_1 is lower, and it is easier to satisfy the condition for no disclosure in Corollary 3. When r is sufficiently high, we obtain $\rho_1 = 1$, and the condition for no disclosure is the same as in Corollary 1.

5.3. Properties of optimal disclosure rules

The rest of this section focuses on the case $E(\tilde{\theta}) < \rho_1$, in which some disclosure is necessary to achieve an optimal outcome. As we illustrate below, the solution to the regulator's problem continues to depend on a gain-to-cost ratio, but now there is a different gain-to-cost ratio for every score.

Specifically, for $i \in \{1, \dots, k\}$, the gain from assigning score s_i to type θ is the term $\Pr(\tilde{\varepsilon} < 1 - \theta)$ in the objective function (10). The cost is $\rho_i - \theta$, which follows from constraint i in (11). Hence, the gain-to-cost ratio from assigning score s_i to type $\theta \neq \rho_i$ is

$$G_i(\theta) \equiv \frac{\Pr(\tilde{\varepsilon} < 1 - \theta)}{\rho_i - \theta}. \quad (14)$$

While the gain does not depend on the specific score and is exactly the same as in Section 4, the cost is different. It is more costly to assign scores that correspond to higher reservation prices. For types below ρ_i , the cost is positive, as these types require resources when they obtain score s_i (they reduce the group average). For types above ρ_i , the cost is negative, as these types provide resources (they increase the group average). For use below, we refer to a score that is associated with a higher reservation price as being higher. So score s_1 is the highest, score s_2 is the second highest, and so on.

A simple case is when there is only one type above 1. Then, there is only one score that induces selling, and the solution is similar to that in the second part of Proposition 1, except that the gain-to-cost ratio is $G_1(\theta)$ instead of $G(\theta)$. However, if there are two or more scores that induce selling, the solution is more complicated. Now the regulator needs to decide not only the probability that each type sells its asset but also the allocation of the k scores that induce selling to the types that sell.

The next proposition simplifies the analysis. It shows that there is a solution that satisfies the following two properties. First, each type above 1 obtains the score that is associated with its reservation price. That is, type $\theta_i > 1$ obtains score s_i with probability 1. Second, and perhaps counterintuitive, among the types below 1 that sell their assets, lower types obtain higher scores.

Proposition 3. *The regulator's problem in Proposition 2 has a solution that satisfies the following properties:*

1. $h_i(\theta_i) = 1$, for every $i \in \{1, \dots, k\}$.
2. $h_i(\theta_v)h_j(\theta_w) = 0$, for every v, ω, i, j , such that $\theta_v < \theta_w < 1 < \theta_i < \theta_j$. That is, if the lower type θ_v obtains the lower score s_i , then the higher type θ_w does not obtain the higher score s_j .

To prove Proposition 3, we show that, if a solution does not satisfy a property, we can construct an alternate solution that satisfies the property. The alternate solution “saves on resources” in the sense that it weakly relaxes the constraints in (11). Moreover, if the types above 1 have different reservation prices ($\rho_1 > \rho_2 > \dots > \rho_k$), the alternate solution strictly relaxes at least one of the constraints.¹⁹

The first part in Proposition 3 reflects the fact that it is more costly to give higher scores. Accordingly, the regulator gives each type above 1 the lowest possible score under which the type will agree to sell. This increases the amount of resources that types above 1 provide to cross-subsidize lower types, and it also reduces the amount of resources that the types below 1 that are matched with types above 1 end up with.

The second part in Proposition 3 reflects the fact that for the types below 1, the relative cost of giving a higher score rather than a lower score is increasing in type. That is, if $\rho_j > \rho_i$, the ratio $\frac{\rho_j - \theta}{\rho_i - \theta}$ is increasing in θ . The idea behind the proof is as follows. Suppose there is a solution in which a lower type ($\theta_v < 1$) obtains a lower score (s_i), and a higher type ($\theta_w < 1$) obtains a higher score (s_j). We can save on resources, without affecting the value of the objective function, by making the following two changes. First, reduce the probability that the higher type obtains the higher score and increase the probability that the lower type obtains that score. We can choose these probabilities, such that the resource constraint for the higher score (and hence the price for that score) is unaffected. Second, to ensure that the probabilities that the two types sell their assets remain unchanged, increase the probability that the higher type obtains the lower score and reduce the probability that the lower type obtains that score. Since the relative cost of giving the higher score rather than the lower score is increasing in type, this second change will relax the constraint for the lower score.

Another way to see the intuition for the result above is by separating the cost $\rho_i - \theta$ into two components: $\rho_i - 1$ and $1 - \theta$. The latter is the cost of bringing the bank up to the threshold of 1, and the former is the cost of bringing it up further from 1 to the reservation price ρ_i , which is associated with the score. For the types that are slightly below 1, the second component is relatively negligible, while the first component is relatively first order. In contrast, for the very low types, the second component is relatively first order. Hence, to save on resources, it is more beneficial to reduce the first component for the types that are closer to 1. This can be done by giving these types scores that are associated with lower reservation prices.

In the remainder of this subsection, we focus on two subcases: $E(\tilde{\theta}) < 1$ and $E(\tilde{\theta}) \in [1, \rho_1)$.

If $E(\tilde{\theta}) < 1$, it is impossible to implement an outcome in which every type sells with probability 1, and so, all the constraints in (11) are satisfied with equalities. The intuition is similar to that in Section 4, except that now it is even harder to implement the desired outcome in which every type sells with probability 1 because types with reservation prices above 1 will not agree to sell for a price of 1. Formally:

Lemma 8. *If $E(\tilde{\theta}) < 1$, every solution to the regulator’s problem satisfies the following:*

1. *There exists a type $\theta < 1$ for which $\sum_{i=1}^k h_i(\theta) < 1$.*
2. *All the constraints in (11) are satisfied with equalities.*

¹⁹ The only case in which two types above 1 have the same reservation price is when they both have a reservation price of 1 (i.e., r is sufficiently high). In this case, we can assume, without loss of generality, that only one score corresponds to that reservation price and that both types obtain that score with probability 1.

An immediate implication of Lemma 8 is that, for every score that induces selling, the sale price equals to the reservation price of the highest type that obtains the score. That is, for every $i \in \{1, \dots, k\}$, $x(s_i) = \rho_i$. Another implication is that, if the types above 1 have different reservation prices, then *any* solution to the regulator's problem must satisfy the two properties in Proposition 3. If this were not the case, we could construct an alternate solution that strictly relaxes one of the constraints in (11), which would violate the second part in Lemma 8. More generally, we can show that, if there are two types above 1 with different reservation prices, then these two types cannot obtain the same score. Combining the implications above, we obtain the following:

Corollary 4. *If $E(\tilde{\theta}) < 1$, then under an optimal disclosure rule:*

1. *Every type above 1 sells for its reservation price.*
2. *Among the types above 1, higher types sell for (weakly) higher prices.*
3. *Among the types below 1 that sell with a positive probability, lower types sell for (weakly) higher prices.*

Corollary 4 implies that, unless $\rho_1 = \rho_2 = \dots = \rho_k$, the sale price is non-monotone in type. Among the types below 1 that sell their assets, lower types sell for higher prices. However, among the types above 1, the opposite is true, as these types end up selling exactly for their reservation price, which is increasing in type. An immediate implication of Corollary 4 is that when types above 1 have different reservation prices, the regulator must assign more than two scores that induce selling, so that each type above 1 sells exactly for its reservation price.

Corollary 5. *If $E(\tilde{\theta}) < 1$ and $\rho_1 > \rho_2 > \dots > \rho_k$, the regulator must assign at least $k + 1$ scores.*

Finally, if $E(\tilde{\theta}) \in [1, \rho_1)$, then for some parameter values, it is possible to implement an outcome in which every type sells with probability 1, but for some parameter values, it is impossible to implement such an outcome. Proposition 3 suggests a simple algorithm to check whether it is possible to implement an outcome in which every type sells with probability 1. Assign score s_1 to the highest type and to the lowest types below 1. Assign this score to as many types below 1 as possible subject to the constraint that the average type receiving the score is ρ_1 (i.e., the constraint for score s_1 in (11)). Next, assign score s_2 to type θ_2 and to as many of the remaining lowest types below 1, subject to constraint for score s_2 ; and so on.²⁰ If the process ends when all types below 1 are assigned scores with probability 1, then it is possible to achieve such an outcome. Otherwise, it is impossible to implement such an outcome, and the implications are the same as in the case $E(\tilde{\theta}) < 1$.

5.4. Closed-form solutions and examples

Propositions 2 and 3 provide general properties of optimal disclosure rules, along with a general algorithm (linear programming) that can be used to determine an optimal disclosure rule for every set of parameters and distribution functions covered by our model. To get a better idea of how optimal disclosure rules work, we illustrate optimal disclosure rules in some special cases.

²⁰ Equivalently, we can start pooling the lowest type above 1 (type θ_k) with the highest types below 1 until the average θ for the group equals ρ_k ; then pool the second lowest type above 1 (type θ_{k-1}) with the remaining highest types below 1 until the average θ for the group equals ρ_{k-1} ; and so on.

Case 1. $\rho_1 = 1$: Here, the highest reservation price is 1. As we know from (9), this can be consistent with having multiple types above 1, but either r is sufficiently high or θ_{\max} is sufficiently low, so they are willing to sell the asset at a price of 1. In this case, the optimal disclosure rule is identical to the one when the bank does not observe its type, as in Proposition 1. In particular, if $h(\theta)$ solves the regulator’s problem in Lemma 4, then $h_1(\theta) = h(\theta)$ and $h_i(\theta) = 0$ for every $i > 1$ solves the regulator’s problem in Proposition 2; and vice versa: if $\{h_i(\theta)\}$ solves the problem in Proposition 2, then $h(\theta) = \sum_{i=1}^k h_i(\theta)$ solves the problem in Lemma 4.

Case 2. $k = 1$: Here, there is only one type above 1. In this case, the optimal disclosure rule is similar to that in Proposition 1, except the gain-to-cost ratio is $G_1(\theta)$ instead of $G(\theta)$. Corollary 2, which provides a sufficient condition for a simple cutoff rule to be optimal, then holds only if ρ_1 is sufficiently small. Otherwise, $G_1(\theta)$ is decreasing when $\theta < 1$, even if Condition 1 is satisfied, and so, a simple cutoff rule is not optimal.²¹ Intuitively, when ρ_1 is very high, the cost $\rho_1 - \theta$ of giving the high score to type $\theta < 1$ is very high, no matter how high θ is, and so, the dominant factor in deciding which types should be included in risk sharing is that the gain $\Pr(\tilde{\varepsilon} < 1 - \theta)$ is decreasing in θ . So, when ρ_1 is sufficiently high, the low types and the only type above 1 obtain the high score, while the types in the middle obtain the low score.

Case 3. $k \geq 2$ and $G_i(\theta)$ is increasing in θ for every $\theta < 1$ and every $i \in \{1, \dots, k\}$ ²²: Using similar logic as in Section 4, one can show that the lowest types are excluded from risk sharing. The optimal disclosure rule can be found by applying the two properties in Proposition 3. First, pool the lowest type above 1 (type θ_k) with the highest types below 1 until all the resources from type θ_k are exhausted (that is, until the average θ for the group equals ρ_k). Next, pool the second lowest type above 1 (type θ_{k-1}) with the remaining highest types below 1 until the resources from type θ_{k-1} are exhausted. And so on, until we exhaust the resources from the highest type θ_1 . The following example illustrates the solution:

Example 1. There are five types $\theta_1 > \theta_2 > 1 > \theta_3 > \theta_4 > \theta_5$, and $\rho_1 > \rho_2$. So, without loss of generality, there are three scores: s_0, s_1 , and s_2 . Suppose $G_i(\theta)$ is increasing in θ for every $\theta < 1$ and $i \in \{1, 2\}$. Suppose in addition that:

$$p(\theta_2)(\theta_2 - \rho_2) = p(\theta_3)(\rho_2 - \theta_3) \tag{15}$$

$$p(\theta_1)(\theta_1 - \rho_1) = p(\theta_4)(\rho_1 - \theta_4). \tag{16}$$

Then, the optimal disclosure rule is as follows (an element in the table is the probability that type θ_j obtains score s_i):

	θ_5	θ_4	θ_3	θ_2	θ_1
score s_1 (sell at price ρ_1)		1			1
score s_2 (sell at price ρ_2)			1	1	
score s_0 (keep asset)	1				

In particular, equation (15) implies that θ_3 gets score s_2 with probability 1, and equation (16) implies that θ_4 gets score s_1 with probability 1, so that the two constraints in (11) are satisfied

²¹ Formally, $G_1(\theta)$ is increasing when $\theta < 1$ if, and only if, $-F'(1 - \theta)(\rho_1 - \theta) + F(1 - \theta) \geq 0$. This reduces to $\rho_1 \leq \max_{\{\theta \in \Theta: \theta < 1\}} \{\theta + \frac{F(1-\theta)}{F'(1-\theta)}\}$.

²² A sufficient condition for this to happen is that $\rho_1 < \max_{\{\theta \in \Theta: \theta < 1\}} \{\theta + \frac{F(1-\theta)}{F'(1-\theta)}\}$.

with equality. As we can see, θ_1 and θ_4 are pooled together at the highest price, θ_2 and θ_3 are pooled together at the lower price, and θ_5 does not sell and does not participate in risk sharing.²³

Case 4. $k \geq 2$ and $G_i(\theta)$ is decreasing in θ for every $\theta < 1$ and every $i \in \{1, \dots, k\}$ ²⁴: In this case, the types in the “middle” are excluded from risk sharing, and the optimal disclosure rule can be found as follows: Pool the highest type θ_1 with the lowest types until all the resources from type θ_1 are exhausted. Next, pool the second highest type θ_2 with the remaining lowest types, and so on, until all the resources of the types above 1 are exhausted.

Case 5. $k \geq 2$ and there exists $\hat{k} \in \{1, \dots, k\}$, such that, for every $\theta < 1$, $G_i(\theta)$ is decreasing in θ if $i \in \{1, \dots, \hat{k}\}$ and increasing in θ if $i \in \{\hat{k} + 1, \dots, k\}$ ²⁵: In this case, the optimal disclosure rule can be found by combining the procedures in cases 3 and 4. The following example illustrates this.

Example 2. Suppose $\theta_1 > \theta_2 > 1 > \theta_3 > \theta_4 > \theta_5$ and $\rho_1 > \rho_2$, as in Example 1. Suppose $G_2(\theta)$ is increasing when $\theta < 1$, but $G_1(\theta)$ is decreasing. Suppose that equation (15) holds, but, instead of equation (16), we have $p(\theta_1)(\theta_1 - \rho_1) = p(\theta_5)(\rho_1 - \theta_5)$. Then, the optimal disclosure rule is as follows:

	θ_5	θ_4	θ_3	θ_2	θ_1
score s_1 (sell at price ρ_1)	1				1
score s_2 (sell at price ρ_2)			1	1	
score s_0 (keep asset)		1			

5.5. Discussion of nonmonotonicity

Optimal disclosure rules may exhibit two forms of nonmonotonicity. First, the probability of selling the asset may be nonmonotone in type (Example 2 and the discussion at the end of Section 4). Second, for the types that sell, the sale price may be nonmonotone in type (Corollary 4 and Examples 1 and 2).

The first form of nonmonotonicity arises when the gain-to-cost ratio is decreasing in θ . A necessary condition for this is that the gain, which follows from the coefficient of $\sum_{i=1}^n h_i(\theta)$ in the regulator’s objective function, is decreasing in θ when $\theta < 1$. In other words, the regulator has a preference for helping lower types. In our model, this happens because, without engaging in risk sharing, lower types are more likely to end up with cash holdings below 1.

The second form of nonmonotonicity follows from the tension between the regulator’s desire to pool different types together and the types’ participation constraints. For this tension, it is important that there is some social benefit from insurance across different types and that the bank has some private information related to the information that the regulator has. However, the specific form of insurance is not crucial. For example, nonmonotonicity will continue to hold if we replace the coefficient $\Pr(\bar{\varepsilon} < 1 - \theta)$ in the regulator’s objective function with some

²³ In general, a type below 1 could be pooled with more than one type above 1. For example, if we changed equations (15) and (16) so that $p(\theta_2)(\theta_2 - \rho_2) = \frac{1}{3}p(\theta_3)(\rho_2 - \theta_3)$ and $p(\theta_1)(\theta_1 - \rho_1) = \frac{2}{3}p(\theta_3)(\rho_1 - \theta_3) + \frac{4}{3}p(\theta_4)(\rho_1 - \theta_4)$, then θ_3 will obtain score s_1 with probability $\frac{2}{3}$ and score s_2 with probability $\frac{1}{3}$, and θ_4 will obtain s_1 with probability $\frac{4}{3}$ (and s_0 with probability $\frac{1}{3}$).

²⁴ A sufficient condition for this to happen is that $\rho_k > \max_{\{\theta \in \Theta: \theta < 1\}} \{\theta + \frac{F(1-\theta)}{F'(1-\theta)}\}$.

²⁵ A sufficient condition for this to happen is that $\rho_{\hat{k}} > \max_{\{\theta \in \Theta: \theta < 1\}} \{\theta + \frac{F(1-\theta)}{F'(1-\theta)}\} > \rho_{\hat{k}+1}$.

other positive coefficient that depends on θ . Nonmonotonicity will also prevail if we alter our model so that the bank must sell its asset at price $x \geq 1$ to obtain r (footnote 15). In this case, the idiosyncratic risk ε does not affect the bank's payoff function, yet pooling types together continues to be socially desirable to ensure that the price is at least 1.²⁶ If, however, the bank could obtain r by selling at any price, full disclosure would clearly be optimal, as there would be no social surplus from insurance; this case boils down to a simple setting of gains from trade with adverse selection.

Another case in which nonmonotonicity will prevail is if the bank is risk averse, namely the bank's payoff is given by some concave function rather than the discontinuous function in (1).²⁷ In this case, the optimal disclosure rule can be found using a similar algorithm as in case 4 in Section 5.4, but with the following adjustment. As long as there remain types that have not been assigned a score, we will continue to match the highest type of the remaining types with the lowest remaining types. So in contrast to Lemma 6, every type will end up selling its asset.²⁸

6. Bank can affect the value of its asset

A natural question is whether we should expect to see nonmonotone rules in practice. Alternatively, one could ask whether nonmonotone rules will continue to prevail if we enrich our model. In this section, we show that once we add two additional assumptions to our model, optimal disclosure rules become monotone.

The first assumption is that the bank can affect the value of its assets. Specifically, we assume that before the regulator learns the value of θ , the bank can reduce its θ by freely disposing its assets. This assumption leads to an additional constraint that the bank's expected equilibrium payoff is weakly increasing in type.²⁹

The second assumption is that the regulator cannot randomize. That is, $g(s|\theta)$ can take only the values 0 and 1. This could be viewed as a practical constraint on what regulators can do.

Proposition 4 below characterizes optimal disclosure under the two assumptions above. The proposition shows that an optimal disclosure rule involves two cutoffs $z_L \leq z_H$. The role of z_L is to determine which types sell their assets. Types below z_L sell with probability 0, while all the other types sell with probability 1. The role of z_H is to determine which types are pooled together. Types on the interval $[z_L, z_H]$ sell for the same price, and they can all obtain the same score. Types above z_H sell at higher prices to reflect their higher reservation prices. If the average θ of the types above z_H is at least ρ_1 , it is possible to implement optimal disclosure by giving all the types above z_H the same score. Otherwise, it is impossible to do so, and for some parameter values, the only way to implement optimal disclosure is by giving each type above z_H its own score.

The optimal cutoffs z_L, z_H can be found as follows. For a given z_H , it is optimal to set z_L as low as possible, subject to the constraint that the average cash flow for the types on $[z_L, z_H]$ is

²⁶ In this case, type θ 's reservation price is $\rho(\theta) = \theta$, if $\theta < 1$, and $\rho(\theta) = \max\{1, \theta - r\}$, if $\theta \geq 1$, which does not satisfy the third property of the reservation price in Section 5.1. Hence, the third part in Lemma 6 will not necessarily hold. In particular, for any disclosure rule in which a type below 1 keeps its asset, there will be a payoff equivalent rule in which that type sells its asset for the true value. The wording in Proposition 4 and the third part of Corollary 4 will need to be adjusted to reflect that.

²⁷ For more details, see the last paragraph in the proof of Lemma A.1 in the appendix.

²⁸ For a complete analysis of our model with a concave payoff function, see a follow-up paper by Garcia et al. (2017).

²⁹ Note that this constraint is violated in Examples 1 and 2. For instance, in Example 1, type θ_4 's expected equilibrium payoff is higher than type θ_3 's expected equilibrium payoff.

at least $\rho(z_H)$. This constraint can be satisfied only if $z_H \geq 1$, and it reduces to $z_L = \hat{z}_L(z_H)$, where

$$\hat{z}_L(z_H) = \min\{z \in \Theta : \sum_{\theta \in [z, z_H]} p(\theta)[\theta - \rho(z_H)] \geq 0\}. \tag{17}$$

The optimal z_H minimizes $\hat{z}_L(z_H)$.

Formally, denote $z_H^* = \arg \min_{z_H \in \{\theta: \theta \geq 1\}} \hat{z}_L(z_H)$.³⁰ The optimal outcome is summarized in the next proposition.

Proposition 4. *Suppose the bank can freely dispose assets and the regulator must follow a deterministic disclosure rule. Then under an optimal disclosure rule:*

1. Types below $\hat{z}_L(z_H^*)$ do not sell their assets, types that belong to the interval $[\hat{z}_L(z_H^*), z_H^*]$ sell at the same price, and types above z_H^* sell at a higher price (or prices).
2. If $\hat{z}_L(\theta) = \theta$ for every $\theta > z_H^*$, then every type above z_H^* must obtain a different score and sell at a different price.

Intuitively, the choice of z_H involves a tradeoff. A higher z_H increases the resources that are available to cross subsidize types below 1, but it also requires an increase in the resources that each type on $[\hat{z}_L(z_H), z_H]$ should end up with, because these required resources are determined by the reservation price of z_H . The optimal z_H balances these two forces.³¹

In general, an optimal disclosure rule must include at least three scores to distinguish among the types that do not sell, the types that sell at the same price, and the types that sell at a higher price (or prices). In some cases, more than three scores are needed to distinguish among the high types that sell.

Interestingly, for some parameter values, an optimal disclosure rule must assign each type a different score, i.e., provide full disclosure.³² Formally, recall that the lowest type is $\theta_{\min} = \theta_m$ and the second lowest type is θ_{m-1} . Then:

Corollary 6. *Suppose the bank can freely dispose assets and the regulator must follow a deterministic disclosure rule. Then, if $\theta_{m-1} \geq 1$ and $\hat{z}_L(\theta) = \theta$ for every $\theta \geq \theta_{m-1}$, the only way to implement an optimal outcome is by full disclosure.*

Corollary 6 illustrates two cases in which full disclosure is necessary. In the first case, all types are above 1 ($\theta_m \geq \theta_{m-1} \geq 1$) but have very different reservation prices. Full disclosure is necessary because no type is willing to sell at a price that reflects the average of a group that includes that type and lower types. The second case is similar to the first case, but there is an additional type below 1 ($\theta_m < 1$), which ends up keeping its asset.³³

Finally, one could ask what happens if banks can freely dispose assets and the regulator can randomize. In an earlier draft, we analyze this case. We show that, perhaps surprisingly, for some parameter values, the two forms of nonmonotonicity in Section 5.5 continue to hold. In particular, for some parameter values, lower types continue to sell at higher prices. However,

³⁰ If the set $\arg \min_{z_2 \in \Theta} \hat{z}_1(z_2)$ contains more than one number, then any of them is optimal and any of them satisfies Proposition 4.

³¹ Example 3 in the appendix illustrates this tradeoff.

³² This last result holds even if the bank cannot freely dispose asset, but under a smaller set of parameters. (The proof contains more details.)

³³ The second case is consistent with $E(\tilde{\theta}) < 1$ as well as with $E(\tilde{\theta}) \geq 1$.

to ensure that the equilibrium payoff is weakly increasing in type, these lower types sell with a probability that is less than 1. We also show that, for some parameter values, some types above 1 sell above their reservation prices. This is beneficial because it allows the lower types that obtain high scores to sell with a higher probability, while still ensuring that the equilibrium payoff is increasing in type.³⁴

7. Discussion and extensions

In this section, we discuss some of the assumptions, interpretations, and possible extensions of the model.

1. We assumed that the bank's payoff function (1) exhibits some discontinuity. This assumption is important for the basic tradeoff in our model when the bank does not observe its type, because it implies that disclosure can not only reduce risk-sharing opportunities as in Hirshleifer (1971) but is also sometimes necessary to prevent risk-sharing markets from breaking down. The latter effect is a direct consequence of our discontinuity assumption and would not be obtained if we had a traditional specification of insurance motive, namely a risk-averse bank with a standard concave utility function. In contrast, the discontinuity in the bank's payoff function is not necessary for our main results for the case in which the bank knows its type. As we explained in Section 5.5, a concave payoff function will lead to similar results.

2. In our model, the bank shares risk with the market. The nature of our results remains the same if banks enter risk-sharing arrangements among themselves rather than with the market (e.g., as in Allen and Gale, 2000; Leitner, 2005), provided that idiosyncratic risk is fully diversified within a group. This is the case if there is a continuum of banks of each type, or if there are two banks of each type, one with cash flow $\theta + \varepsilon$ and one with $\theta - \varepsilon$. This is also the case if the regulator provides insurance against idiosyncratic risk within a group.

3. In our model, the regulator discloses information about the value of existing assets. In practice, stress tests are designed to assess how the bank would perform under some hypothetical stress scenario. Yet, information from stress tests can still have a significant impact on current values. First, if the probability of the stress scenario is sufficiently high, stress test results will clearly affect current asset values. Second, stress tests results will also have a significant impact on current values if the regulator requires the bank to take some action (e.g., increase capital), which is costly to the bank, say, by preventing the bank from implementing its desired investment strategy. Our model suggests that, rather than provide detailed information about how each bank would perform under stress scenarios, it might be optimal to just say what action each bank should take. Moreover, it might be optimal to require banks of different levels of strength to take the same action (see more below).

4. An interesting extension of our model would allow the regulator to provide funds to banks. Such an extension would suggest that, in some cases, it is optimal to inject money not only to weak banks but also to strong banks so that the market cannot distinguish among them.³⁵ For example, suppose there are two banks: strong ($\theta_1 = 1.2$, $\rho_1 = 1$) and weak ($\theta_2 = 0.4$), and the

³⁴ Note that some types above 1 could sell above their reservation prices also in the case discussed in Proposition 4.

³⁵ Indeed, one of the first uses of the Troubled Asset Relief Program (TARP) funds was providing capital to nine major financial institutions as part of the Capital Purchase Program, a program designed to infuse capital to "healthy" banks. During the audit, former Federal Reserve Chair Ben Bernanke told the Special Inspector General for the TARP that "there were differences in the nine banks in terms of strength and weakness, but that the selection was generalized in order to avoid stigmatizing any one bank as being a weak bank and creating panic." (See SIGTARP report 10-001, "Emergency

regulator has a bailout fund in the amount of 0.4. Suppose that the regulator would like each bank to end up with at least one dollar and that both banks can raise cash by issuing equity. Giving all the money to the weak bank identifies this bank as weak; and because the value of the weak bank after the cash injection is 0.8, that bank will not be able to raise one dollar by issuing equity. Splitting the money equally between the two banks leads to a better outcome. Now, after the cash injection, the value of the strong bank is 1.4, and the value of the weak bank is 0.6. But since the market cannot distinguish between the two banks, each bank can sell its equity for a price of 1, which is the average value of both banks. So, both banks can guarantee cash holdings of 1.³⁶

5. In our model, all the economic surplus is captured by the banking sector, and so, the regulator sets a disclosure rule aiming to maximize the surplus in the banking sector. Our model can easily capture externalities that banks impose on the rest of society. For example, suppose that when a bank of type θ fails ($\theta + \varepsilon < 1$), society suffers some exogenous loss $l(\theta)$. Then, the social gain from having a bank sell its asset is higher by $l(\theta)$. We can include this gain in the regulator's objective function and take it into account in the design of the disclosure rule.³⁷ Our main results continue to hold in this case. Clearly, now the regulator will have a stronger motive to help banks with a high $l(\theta)$. This may capture the familiar “too big to fail” argument, whereby regulators try to save institutions whose failure could pose a substantial damage to the economy.

6. As in any standard mechanism design problem, we assumed that the regulator can commit to assigning scores according to the announced disclosure rule; in particular, the regulator does not lie. In our basic setting, this assumption is crucial because ex post the regulator could gain by giving a score that induces selling instead of a score that does not induce selling. In some cases, the commitment not to lie arises endogenously. For example, if banks share risk among themselves, rather than with a third party, and the optimal disclosure rule has two scores, as in Proposition 1, then the regulator cannot gain ex post by lying. If he lies, banks will have insufficient funds to honor the agreements, and they could all fail by contagion.³⁸ However, if the ex-ante optimal disclosure rule has a group whose average θ is more than 1 (i.e., $\rho_1 > 1$), the regulator *could* gain ex post by adding more types to that group, as long as the group average is at least 1. Anticipating that, types with reservation prices above 1 will not agree to participate in risk sharing. To overcome this problem, the regulator could supplement the disclosure policy with other instruments, such as guaranteeing payments under these risk-sharing arrangements, or providing convincing evidence that the average θ for a group receiving a given score is as high as announced.

7. Another aspect of commitment is that the regulator cannot alter the disclosure rule ex post, once he learns the bank's type. This type of commitment arises naturally if there is a continuum of banks and the probability $p(\theta)$ of being a type θ represents the proportion of banks of this type in the continuum. Then, maximizing each bank's ex-ante expected payoff is the same as

capital injections provided to support the viability of Bank of America, other major banks, and the U.S. financial system,” October 2009.)

³⁶ We can extend the example to show that, in some cases, it is optimal to create two groups of banks, each containing strong banks and weak banks, and inject cash only to banks in one group. For example, suppose there are four banks: $\theta_1 = 2.2$, $\theta_2 = 1.2$, $\theta_3 = 0.4$, $\theta_4 = 0.2$, with $\rho_1 = 1.2$ and $\rho_2 = 1$. Then it is optimal to have θ_2 and θ_3 in one group (as before) and θ_1 and θ_4 in another group, and it is optimal to split the bailout money (0.4) equally only among banks in the first group. Then every bank can raise at least \$1 by selling equity. Banks θ_2 and θ_3 will sell for \$1, and banks θ_1 and θ_4 will sell for \$1.2.

³⁷ Specifically, we can replace the regulator's objective function (10) with $\sum_{\theta \in \Theta} p(\theta) \Pr(\tilde{\varepsilon} < 1 - \theta) [r + l(\theta)] \sum_{i=1}^k h_i(\theta)$.

³⁸ This would be true, even if the regulator privately observed the aggregate state.

maximizing the sum of banks' ex-post payoffs. Since the regulator is interested in that, he has no incentive to deviate to another disclosure rule ex post.³⁹ In this sense, the regulator is different from a single bank. A strong bank will have incentive to reveal itself as strong, even when it is socially optimal that the bank is pooled together with weaker banks.

8. The discussion above suggests that studying disclosure by the regulator and not by individual banks is very natural in the context of our model. First, ex ante, the regulator and banks may have different objective functions because the regulator cares about externalities that banks impose on the rest of society.⁴⁰ Second, ex post, the regulator's commitment to a disclosure rule arises more naturally than that of an individual bank.

9. Assigning a score is equivalent to recommending a price to the market. A high score does not necessarily mean that the bank is strong. It only means that the average cash flow conditional on obtaining the score equals the recommended price. One can also think of scores more broadly. Scores separate banks into groups, and assigning scores is isomorphic to recommending to the banks which groups to form. For example, one can think of scores as suggesting mergers among banks; or if we applied our ideas to Leitner's (2005) setting, scores could suggest joint liability arrangements. Our model provides a characterization of optimal groups that could emerge following disclosure when each bank can choose whether to join the recommended group or stay in autarky and under the assumption that idiosyncratic risk is fully diversified within a group. In one interpretation, our model illustrates how the regulator's disclosure policy affects the financial networks that banks form.⁴¹

10. We believe that our model can be used as a benchmark to think of credit rating agencies. Our model suggests that coarse ratings as well as low types receiving high ratings may be a feature of a socially optimal outcome. This is different from other papers (e.g., Kartasheva and Yilmaz, 2013; Goel and Thakor, 2015), in which coarse ratings reduce welfare. An interesting question is what the optimal disclosure rule looks like when the regulator faces competition from credit rating agencies or whether it is possible to implement risk sharing when the regulator and credit rating agencies have different objectives.

11. We could interpret our paper as one about the design of stress tests rather than disclosure. In particular, choosing a disclosure rule is the same as designing an experiment (e.g., stress tests) that provides a public signal $s \in S$ according to some distribution $g(s|\theta)$. With this interpretation, the regulator's commitment boils down to committing not to manipulate the public signal.

12. Our basic setup maps into a general Bayesian persuasion game (Kamenica and Gentzkow, 2011) with one sender (the regulator) and one receiver (the market) whose equilibrium action (price offered) equals the expected state (bank's type) conditional on posterior beliefs.⁴² The standard concavification approach has limited applicability in characterizing optimal disclosure rules in our setting because if the bank observes its type, the equilibrium price generally depends on the entire distribution of types given the market's posterior beliefs. In particular, the equilibrium price x is a fixed point of $\hat{E}(\theta|\rho(\theta) \leq x)$, where \hat{E} denotes the expected value given the market's posterior beliefs, and the condition $\rho(\theta) \leq x$ reflects the fact that the bank agrees to

³⁹ This aspect of commitment holds, even if risk sharing is between the bank and the market, as is explicitly assumed in our basic setting.

⁴⁰ The idea that the regulator and banks may have different incentives to disclose information is also explored in Alvarez and Barlevy (2015).

⁴¹ We do abstract, however, from other issues of group formation, such as whether a bank receiving one score will attempt to form a link with a bank receiving a different score.

⁴² So, the market can be viewed as having a quadratic utility function: $-(x - \theta)^2$.

sell only if the price is above its reservation price.⁴³ If the bank does not observe its type, the equilibrium price is simply $x = \hat{E}(\theta)$, but the standard concavification approach still has limited applicability in characterizing optimal disclosure rules, because the regulator's payoff depends on the bank's type: it is r if $x \geq 1$ and $r \Pr(\tilde{\varepsilon} \geq 1 - \theta)$ if $x < 1$.⁴⁴

13. Our results on optimal disclosure rules could be applied to other settings of Bayesian persuasion in which disclosure is followed by a game in which an uninformed receiver makes a price offer to a third party that has the information that the sender has and that can either accept or reject the offer. For example, consider schools that grade students with different abilities and potential employers who care about the average ability of the students they hire. Suppose that students know their own abilities and can open their own businesses instead of getting hired. In this case, a student's reservation price is the benefit from opening his own business. Our analysis can shed light on the way schools will communicate information about students.⁴⁵ In different settings, reservation prices could also reflect payoffs that informed third parties could obtain by signaling their types.

8. Conclusion

We provide a model of an optimal disclosure policy of a regulator who has information about banks. The disclosure policy affects whether banks can take corrective actions, particularly whether banks can engage in risk sharing to ensure that their capital does not fall below some critical level.

We show that, if the average forecasted capital is sufficiently high, $E(\tilde{\theta}) > \rho_1$, then no disclosure is necessary. Otherwise, some disclosure is necessary, which generally takes the form of partial disclosure, pooling together banks of different levels of strength. If banks do not know their types, two scores are sufficient, and, in many cases, a simple cutoff rule is optimal. If a bank knows its type, more scores are needed, and, in general, an optimal disclosure is nonmonotone in that the price at which a bank can sell its asset is nonmonotone in type. However, this nonmonotonicity disappears if banks can freely dispose assets and the regulator must follow a deterministic disclosure rule. In this case, an optimal disclosure rule generally involves two thresholds, such that, banks below the lower threshold are excluded from risk sharing, those in the middle are pooled together at a price that reflects the average capital for the group, and those above the higher threshold are valued at higher prices that reflect their higher reservation values. In some cases, full disclosure is necessary.

We discuss a few extensions, such as the case in which the regulator can inject money to banks. There are many other interesting extensions, some of which have been explored recently. For example, Williams (2017) studies how optimal disclosure affects a bank's ex-ante portfolio choice; Inostroza and Pavan (2017) study a case in which market participants have heterogeneous beliefs; and Orlov et al. (2017) study the joint problem of optimal disclosure and capital requirements when the banks' portfolios are correlated. We believe the model results could also be used to characterize optimal signals in other Bayesian persuasion problems, outside the banking arena.

Finally, there are important issues related to disclosure in stress tests beyond the disclosure of the test results. Leitner and Williams (2017) study whether a regulator should reveal its stress-

⁴³ More recent approaches (e.g., Gentzkow and Kamenica, 2016; Kolotilin, 2018; Dworzak and Martini, 2017) also have limited applicability in characterizing optimal disclosure rules in our setting.

⁴⁴ If the bank observes its type, the regulator's payoff is r if $x \geq \rho(\theta)$ and $r \Pr(\tilde{\varepsilon} \geq 1 - \theta)$ if $x < \rho(\theta)$.

⁴⁵ Ostrovsky and Schwarz (2010) study a similar problem, but without such reservation prices.

testing model to banks. Leitner and Yilmaz (2018) study the extent to which a regulator should rely on bank internal risk models.

Acknowledgments

We thank Briana Chang, Pierre Chaigneau, Mehran Ebrahimian, David Feldman, Emir Kamenica, Alfred Lehar, Stephen Morris, George Pennacchi, Paul Pfleiderer, Pegaret Pichler, Ned Prescott, Francesco Sangiorgi, Joel Shapiro, Nikola Tarashev, Jianxing Wei, an associate editor, and the editor (Alessandro Pavan) for constructive comments. We also thank seminar participants at the Bank of Canada; Board of Governors; Boston University; Dartmouth College; Drexel; Federal Reserve Banks of Cleveland, New York, and Philadelphia; LBS; LSE; UCLA; University of Arizona; University of Calgary; Wharton; AEA; Barcelona GSE Summer Forum workshop “Information, Competition, and Market Frictions”; Fed System conference on financial structure and regulation; Federal Reserve Bank of Philadelphia “Enhancing Prudential Standards in Financial Regulations”; FIRS; MFA; NBER credit rating agency meeting; NFA; NY Fed/NYU Stern Conference on Financial Intermediation; NYU International Network on Expectations and Coordination Conference; OxFIT; Tel Aviv Finance Conference; UBC Summer Finance Conference; and WFA.

Appendix

Proof of Lemma 1. From the text, the equilibrium price is $x(s) = \mu(s)$. Conditional on the bank’s information, the bank’s expected payoff is $E[R(\tilde{\theta} + \tilde{\varepsilon}|\tilde{s} = s)] = \mu(s) + r \Pr(\tilde{\theta} + \tilde{\varepsilon} \geq 1|\tilde{s} = s)$ if it keeps the asset, and $R(\mu(s))$ if it sells. From Assumption 1, $\mu(s) + r > E[R(\tilde{\theta} + \tilde{\varepsilon}|\tilde{s} = s)] > \mu(s)$. Hence, if $\mu(s) \geq 1$, it is optimal for the bank to sell because $R(\mu(s)) = \mu(s) + r > E[R(\tilde{\theta} + \tilde{\varepsilon}|\tilde{s} = s)]$. If $\mu(s) < 1$, it is optimal to keep the asset because $R(\mu(s)) = \mu(s) < E[R(\tilde{\theta} + \tilde{\varepsilon}|\tilde{s} = s)]$.

Proof of Lemma 2. The regulator chooses a disclosure rule (S, g) to maximize $\sum_{\theta \in \Theta} p(\theta)u(\theta)$. Observe that,

$$\begin{aligned} \sum_{\theta \in \Theta} p(\theta) \sum_{s:\mu(s) \geq 1} \mu(s)g(s|\theta) &= \sum_{s:\mu(s) \geq 1} \mu(s) \sum_{\theta \in \Theta} p(\theta)g(s|\theta) \\ &= \sum_{s:\mu(s) \geq 1} \sum_{\theta \in \Theta} \theta p(\theta)g(s|\theta) = \sum_{\theta \in \Theta} p(\theta) \sum_{s:\mu(s) \geq 1} \theta g(s|\theta), \end{aligned} \tag{A.1}$$

where the second equality follows from equation (2). Hence,

$$\begin{aligned} \sum_{\theta \in \Theta} p(\theta)u(\theta) &= \sum_{\theta \in \Theta} p(\theta) \sum_{s:\mu(s) < 1} [\theta + r \Pr(\tilde{\varepsilon} \geq 1 - \theta)]g(s|\theta) \\ &\quad + \sum_{\theta \in \Theta} p(\theta) \sum_{s:\mu(s) \geq 1} [\theta + r]g(s|\theta) \\ &= \sum_{\theta \in \Theta} p(\theta)[\theta + r \Pr(\tilde{\varepsilon} \geq 1 - \theta)] \sum_{s:\mu(s) < 1} g(s|\theta) \\ &\quad + \sum_{\theta \in \Theta} p(\theta)[\theta + r] \sum_{s:\mu(s) \geq 1} g(s|\theta). \end{aligned} \tag{A.2}$$

Since $\sum_{s:\mu(s)<1} g(s|\theta) = 1 - \sum_{s:\mu(s)\geq 1} g(s|\theta)$, we obtain that

$$\sum_{\theta \in \Theta} p(\theta)u(\theta) = \sum_{\theta \in \Theta} p(\theta)[\theta + r \Pr(\tilde{\varepsilon} \geq 1 - \theta)] + r \sum_{\theta \in \Theta} p(\theta) \Pr(\tilde{\varepsilon} < 1 - \theta) \sum_{s:\mu(s)\geq 1} g(s|\theta). \tag{A.3}$$

In the right-hand side of (A.3), only the second term is affected by the disclosure rule. Hence, maximizing $\sum_{\theta \in \Theta} p(\theta)u(\theta)$ is the same as maximizing $\sum_{\theta \in \Theta} p(\theta) \Pr(\tilde{\varepsilon} < 1 - \theta) \times \sum_{s:\mu(s)\geq 1} g(s|\theta)$.

Proof of Lemma 3. For $s \in \hat{S}$, denote $\hat{\mu}(s) = E[\tilde{\theta}|\tilde{s} = s]$, where the expectation is calculated under \hat{g} .

Suppose first that $\sum_{\theta \in \Theta} p(\theta)\hat{g}(s_1|\theta) > 0$ and $\sum_{\theta \in \Theta} p(\theta)\hat{g}(s_0|\theta) > 0$. From (2), the definition of \hat{g} , and the law of iterated expectations,

$$\hat{\mu}(s_1) = \frac{\sum_{\theta \in \Theta} \theta p(\theta)\hat{g}(s_1|\theta)}{\sum_{\theta \in \Theta} p(\theta)\hat{g}(s_1|\theta)} = \frac{\sum_{\theta \in \Theta} p(\theta) \sum_{s:\mu(s)\geq 1} \mu(s)g(s|\theta)}{\sum_{\theta \in \Theta} p(\theta) \sum_{s:\mu(s)\geq 1} g(s|\theta)}. \tag{A.4}$$

Hence, $\hat{\mu}(s_1) \geq 1$. Similarly, $\hat{\mu}(s_0) < 1$. From Lemma 1, under (S, g) , the bank sells its asset with probability $\sum_{s:\mu(s)\geq 1} g(s|\theta)$; and under (\hat{S}, \hat{g}) , the bank sells with probability $\hat{g}(s_1|\theta)$. From the definition of $\hat{g}(s_1|\theta)$, these two probabilities are the same. Lemma 2 then implies that the value of the objective function is also the same under both rules.

In the special case $\sum_{\theta \in \Theta} p(\theta)\hat{g}(s_1|\theta) = 0$, we have that $\hat{g}(s_1|\theta) = 0$ for every $\theta \in \Theta$. Hence, under (\hat{S}, \hat{g}) , $\Pr(\tilde{s} = s_1) = 0$ and $\Pr(\tilde{s} = s_0) = 1$; and under (S, g) , every type obtains a score with $\mu(\cdot) < 1$ with probability 1. Hence, from the law of iterated expectations, under (\hat{S}, \hat{g}) , $E(\tilde{\theta}) = \hat{\mu}(s_1) \Pr(\tilde{s} = s_1) + \hat{\mu}(s_0) \Pr(\tilde{s} = s_0) = \hat{\mu}(s_0)$; and under (S, g) , $E(\tilde{\theta}) = \sum_{s \in S} \mu(s) \Pr(\tilde{s} = s) < 1$. Hence, $\hat{\mu}(s_0) < 1$. Hence, under both disclosure rules, the bank sells with probability 0.

Similarly, if $\sum_{\theta \in \Theta} p(\theta)\hat{g}(s_0|\theta) = 0$, it follows that $\hat{\mu}(s_1) = E(\tilde{\theta}) \geq 1$, and under both disclosure rules, the banks sells with probability 1.

Proof of Lemma 4. As a preliminary, observe that from (2), $\mu(s) \geq 1$ if and only if $\sum_{\theta \in \Theta} p(\theta)(\theta - 1)g(s|\theta) \geq 0$.

From Lemma 3, we can focus, without loss of generality, on disclosure rules (\hat{S}, \hat{g}) that assign at most two scores \hat{s}_0, \hat{s}_1 , such that $\hat{\mu}(\hat{s}_1) \geq 1$ and $\hat{\mu}(\hat{s}_0) < 1$. (As in Lemma 3, $\hat{\mu}(\cdot)$ denotes the conditional expectation under \hat{g} .) We first show that, if such (\hat{S}, \hat{g}) solves the problem in Lemma 2, then $h(\theta) = \hat{g}(\hat{s}_1|\theta)$ solves the problem in Lemma 4. Specifically, since $\hat{\mu}(\hat{s}_1) \geq 1$, h satisfies constraint (6). Now suppose, to the contrary, that there exists $\tilde{h} : \Theta \rightarrow [0, 1]$ that satisfies (6) and gives a higher value for the objective function (5); that is, $\sum_{\theta \in \Theta} p(\theta) \Pr(\tilde{\varepsilon} < 1 - \theta)\tilde{h}(\theta) > \sum_{\theta \in \Theta} p(\theta) \Pr(\tilde{\varepsilon} < 1 - \theta)h(\theta)$. Construct an alternate disclosure rule (S', g') , defined by: $S' = \{s'_0, s'_1\}$, $g'(s'_1|\theta) = \tilde{h}(\theta)$, and $g'(s'_0|\theta) = 1 - \tilde{h}(\theta)$. Since \tilde{h} satisfies (6), then under (S', g') (and when $\mu'(\cdot)$ is calculated under g'), $\mu'(\cdot) \geq 1$, at least for score s'_1 . Hence, the value of the objective function (4) is at least $\sum_{\theta \in \Theta} p(\theta) \Pr(\tilde{\varepsilon} < 1 - \theta)g'(s'_1|\theta) = \sum_{\theta \in \Theta} p(\theta) \Pr(\tilde{\varepsilon} < 1 - \theta)\tilde{h}(\theta)$. In contrast, under (\hat{S}, \hat{g}) , the value of the objective function (4) is only $\sum_{\theta \in \Theta} p(\theta) \Pr(\tilde{\varepsilon} < 1 - \theta)\hat{g}(\hat{s}_1|\theta) = \sum_{\theta \in \Theta} p(\theta) \Pr(\tilde{\varepsilon} < 1 - \theta)h(\theta)$. But this contradicts the optimality of (\hat{S}, \hat{g}) .

Next, we show that if h solves the problem in Lemma 4, then (S, g) , defined by $S = \{s_0, s_1\}$, $g(s_1|\theta) = h(\theta)$, and $g(s_0|\theta) = 1 - h(\theta)$ solves the problem in Lemma 2. Suppose, to the contrary, that there is a disclosure rule (\check{S}, \check{g}) that gives a higher value to (4). Without loss of

generality, $\check{S} = \{\check{s}_0, \check{s}_1\}$, and when $\check{\mu}(\cdot)$ is calculated under \check{g} , $\check{\mu}(\check{s}_1) \geq 1$ and $\check{\mu}(\check{s}_0) < 1$. Hence, $\sum_{\theta \in \Theta} p(\theta) \check{g}(\check{s}_1|\theta) > \sum_{\theta \in \Theta} p(\theta) \Pr(\tilde{\varepsilon} < 1 - \theta) \check{g}(\check{s}_1|\theta) > \sum_{\theta \in \Theta} p(\theta) \Pr(\tilde{\varepsilon} < 1 - \theta) g(s_1|\theta) \geq 0$. Now let $\check{h}(\theta) = \check{g}(\check{s}_1|\theta)$. Then, \check{h} satisfies (6) and gives a higher value for (5) than h . But this contradicts the optimality of h .

Proof of Proposition 1. By Assumption 1, the coefficient of $h(\theta)$ in (5) is positive.

Part 1: Setting $h(\theta) = 1$ for every $\theta \in \Theta$ achieves the maximal attainable value for (5) and satisfies constraint (6). Any other $h : \Theta \rightarrow [0, 1]$ reduces the value of (5).

Part 2: Clearly, $h(\theta) = 1$ for every $\theta \geq 1$, because a higher $h(\theta)$ increases the value of (5) and weakly relaxes the constraint (6). However, setting $h(\theta) = 1$ for every $\theta \in \Theta$ violates (6). Hence, (6) is binding.

Next, we show that if $h(\hat{\theta}) > 0$ for some $\hat{\theta} < 1$, then $h(\theta) = 1$ for every $\theta \in \Theta$, such that $G(\theta) > G(\hat{\theta})$. That is, if type $\hat{\theta} < 1$ sells its asset with some positive probability, then every type with a higher gain-to-cost ratio sells with probability 1. Suppose, to the contrary, that a type $\theta' < 1$ exists, such that $G(\theta') > G(\hat{\theta})$ but $h(\theta') < 1$. We obtain a contradiction to the optimality of h by constructing an alternate solution:

$$\tilde{h}(\theta) = \begin{cases} h(\theta) & \text{if } \theta \notin \{\theta', \hat{\theta}\} \\ h(\theta) + \Delta & \text{if } \theta = \theta' \\ h(\theta) - \frac{p(\theta')(1-\theta')}{p(\hat{\theta})(1-\hat{\theta})} \Delta & \text{if } \theta = \hat{\theta} \end{cases} \tag{A.5}$$

In particular, if $\Delta > 0$ is sufficiently small, \tilde{h} is a function from Θ to $[0, 1]$. Moreover, \tilde{h} satisfies (6) and increases the value of (5) by $p(\theta') \Pr(\tilde{\varepsilon} < 1 - \theta') \Delta - p(\hat{\theta}) \Pr(\tilde{\varepsilon} < 1 - \hat{\theta}) \frac{p(\theta')(1-\theta')}{p(\hat{\theta})(1-\hat{\theta})} \Delta$, which equals $\Delta p(\theta')(1 - \theta')[G(\theta') - G(\hat{\theta})] > 0$.

Hence, there exists a type $\theta^* < 1$, such that, for every $\theta < 1$, $h(\theta) = 1$ if $G(\theta) > G(\theta^*)$; and $h(\theta) = 0$ if $G(\theta) < G(\theta^*)$. Let $G^* = G(\theta^*)$. Then to satisfy (6), we must have

$$\sum_{\theta \geq 1} p(\theta)(\theta - 1) + \sum_{\theta < 1: G(\theta) > G^*} p(\theta)(\theta - 1) + \sum_{\theta < 1: G(\theta) = G^*} p(\theta)(\theta - 1)h(\theta) = 0. \tag{A.6}$$

The result then follows immediately. (Note that there could be more than one type $\theta < 1$, such that $G(\theta) = G^*$. In that case, there is more than one combination of $\{h(\theta)\}_{\theta: G(\theta)=G^*}$ that satisfies (6).)

Proof of Corollary 1. Since $\theta_{\max} > 1$, it follows from Proposition 1 that, under an optimal disclosure rule, at least one type below 1 sells with a positive probability.

Under full disclosure, type θ is offered price θ (i.e., type θ obtains a score s such that $\mu(s) = \theta$). So, type $\theta \geq 1$ sells with probability 1, and type $\theta < 1$ sells with probability 0. If $\theta_{\min} \geq 1$, full disclosure is optimal because every type sells with probability 1. If $\theta_{\min} < 1$, full disclosure is suboptimal because no type below 1 sells.

Under no disclosure, $\mu(s) = E(\tilde{\theta})$ for every $s \in S$. If $E(\tilde{\theta}) \geq 1$, every type sells with probability 1, which is optimal. If $E(\tilde{\theta}) < 1$, every type sells with probability 0, which is suboptimal.

Finally, if $E(\tilde{\theta}) < 1$, then $\theta_{\min} < 1$. Hence, neither full disclosure nor no disclosure is optimal. In other words, partial disclosure is the only way to achieve the optimal outcome.

Proof of Corollary 2. Let $\theta_i < \theta_j < 1$. Then $1 - \theta_i > 1 - \theta_j > 0$. From condition 1, $G(\theta_i) < G(\theta_j)$. The result then follows from Proposition 1.

Proof of Lemma 5. The main argument is in the text. It remains to show that

$$R(x) \geq \theta + r \Pr(\tilde{\varepsilon} \geq 1 - \theta) \tag{A.7}$$

reduces to $x \geq \rho(\theta)$.

Case 1: $\theta \geq 1$. If $x \geq 1$, (A.7) is satisfied if, and only if, $x + r \geq \theta + r \Pr(\tilde{\varepsilon} \geq 1 - \theta)$. If $x < 1$, (A.7) is violated because $R(x) < 1$. Hence, (A.7) is satisfied if, and only if, $x \geq \max\{1, \theta - r \Pr(\tilde{\varepsilon} < 1 - \theta)\} = \rho(\theta)$.

Case 2: $\theta < 1$. If $x \geq 1$, (A.7) is satisfied because $R(x) \geq 1 + r$. If $x < 1$, (A.7) is satisfied if, and only if, $x \geq \theta + r \Pr(\tilde{\varepsilon} \geq 1 - \theta)$. Hence, (A.7) is satisfied if, and only if, $x \geq \min\{1, \theta + r \Pr(\tilde{\varepsilon} \geq 1 - \theta)\} = \rho(\theta)$.

Proof of Lemma 6. Consider an optimal disclosure rule (S, g) .

Part 1: Suppose, to the contrary, that there exist a type $\theta' \geq 1$ and a score $s' \in S$, such that $g(s'|\theta') > 0$, and such that θ' does not sell upon obtaining score s' . Construct an alternate rule, defined by: $\tilde{S} = S \cup \{\tilde{s}\}$,

$$\tilde{g}(s|\theta') = \begin{cases} g(s'|\theta') & \text{if } s = \tilde{s} \\ 0 & \text{if } s = s' \\ g(s|\theta') & \text{if } s \notin \{s', \tilde{s}\}, \end{cases}$$

and, for $\theta \neq \theta'$,

$$\tilde{g}(s|\theta) = \begin{cases} g(s|\theta) & \text{if } s \neq \tilde{s} \\ 0 & \text{if } s = \tilde{s}. \end{cases}$$

Under (\tilde{S}, \tilde{g}) , the only type that obtains score \tilde{s} is θ' . So, the equilibrium price for score \tilde{s} is θ' . By Lemma 5, θ' sells upon obtaining score \tilde{s} . Since the price of each score reflects the average θ of all the types that sell under the score, equilibrium prices for all other scores remain unchanged. Hence, the alternate rule increases the expected payoff for type θ' by $rg(s'|\theta') \Pr(\tilde{\varepsilon} < 1 - \theta')$, while keeping the payoffs for all other types unchanged. But this contradicts the optimality of (S, g) .

Part 2: Suppose the highest type that obtains score s is $\theta_i \geq 1$. That is, $\theta_i = \max\{\theta \in \Theta : g(s|\theta) > 0\}$. From part 1, type θ_i sells its asset upon obtaining score s . So, by Lemma 5, $x(s) \geq \rho_i$. Because $\rho(\theta)$ is increasing in θ , any other type that obtains score s also sells upon obtaining score s . Hence, selling does not convey additional information to the market. Hence, the price is $x(s) = \mu(s)$.

Part 3: Suppose the highest type that obtains score s is $\theta_i < 1$. Suppose, to the contrary, that some types sell upon obtaining score s , and that the highest type that sells is $\tilde{\theta} < 1$. Since the market does not expect to lose money, the price must satisfy $x(s) \leq \tilde{\theta}$. And since $\tilde{\theta} < 1$, we know that $\tilde{\theta} < \rho(\tilde{\theta})$. But then $x(s) < \rho(\tilde{\theta})$, which contradicts Lemma 5.

Proof of Lemma 7. Suppose (S, g) is an optimal disclosure rule. From Lemma 6 and the definition of S_i , type θ sells its asset upon obtaining score $s \in S$ if, and only if, $s \in \cup_{i=1}^k S_i$. Moreover, the sale price is $x(s) = \mu(s)$. Hence, type θ 's expected payoff is

$$u(\theta) = \sum_{s \notin \cup_{i=1}^k S_i} [\theta + r \Pr(\tilde{\varepsilon} \geq 1 - \theta)]g(s|\theta) + \sum_{s \in \cup_{i=1}^k S_i} [\mu(s) + r]g(s|\theta). \tag{A.8}$$

Following the steps in Lemma 2, but replacing the conditions $\mu(s) < 1$ and $\mu(s) \geq 1$ with $s \notin \cup_{i=1}^k S_i$ and $s \in \cup_{i=1}^k S_i$, respectively, we can show that the regulator's payoff reduces to

$$r \sum_{\theta \in \Theta} p(\theta) \Pr(\tilde{\varepsilon} < 1 - \theta) \sum_{s \in \cup_{i=1}^k S_i} g(s|\theta), \tag{A.9}$$

plus a constant term that does not depend on the disclosure rule.

Now consider (\hat{S}, \hat{g}) . It is easy to verify that (\hat{S}, \hat{g}) is a disclosure rule. For every $s_i \in \hat{S}$, such that $\sum_{\theta \in \Theta} p(\theta) \hat{g}(s_i|\theta) > 0$, denote $\hat{\mu}(s_i) = E[\tilde{\theta}|\tilde{s} = s_i]$, where the expectation is calculated under \hat{g} . (If $\sum_{\theta \in \Theta} p(\theta) \hat{g}(s_i|\theta) = 0$, then $\hat{g}(s_i|\theta) = 0$ for every $\theta \in \Theta$, and the set S_i is empty.)

From (2), the definition of \hat{g} , the law of iterated expectations, and part 2 in Lemma 6, it follows that, for each $i \in \{1, \dots, k\}$ such that $\sum_{\theta \in \Theta} p(\theta) \hat{g}(s_i|\theta) > 0$,

$$\hat{\mu}(s_i) = \frac{\sum_{\theta \in \Theta} \theta p(\theta) \hat{g}(s_i|\theta)}{\sum_{\theta \in \Theta} p(\theta) \hat{g}(s_i|\theta)} = \frac{\sum_{\theta \in \Theta} p(\theta) \sum_{s \in S_i} \mu(s) g(s|\theta)}{\sum_{\theta \in \Theta} p(\theta) \sum_{s \in S_i} g(s|\theta)} \geq \rho_i. \tag{A.10}$$

Since the highest type that obtains score s_i is θ_i , it then follows from Lemma 5 and the observation that $\rho(\theta)$ is increasing in θ , that every type that obtains score s_i will agree to sell at price $\hat{\mu}(s_i)$. Hence, if the market offers a price $x(s_i) \geq \hat{\mu}(s_i)$, every type that obtains score s_i will sell at that price. In equilibrium, the market will offer price $x(s_i) = \hat{\mu}(s_i)$, which is the highest price under which the market breaks even. As for score s_0 , we know that the highest type that obtains that score is less than 1. So, from Lemma 6, no type sells upon obtaining s_0 . Hence, under (\hat{S}, \hat{g}) , type θ 's expected payoff is

$$u(\theta) = \sum_{s=s_0} [\theta + r \Pr(\tilde{\varepsilon} \geq 1 - \theta)] \hat{g}(s|\theta) + \sum_{s \in \{s_1, \dots, s_k\}} [\hat{\mu}(s) + r] \hat{g}(s|\theta). \tag{A.11}$$

Again, using similar steps as in Lemma 2, the regulator's payoff reduces to

$$r \sum_{\theta \in \Theta} p(\theta) \Pr(\tilde{\varepsilon} < 1 - \theta) \sum_{s \in \{s_1, \dots, s_k\}} \hat{g}(s|\theta), \tag{A.12}$$

plus the same constant term as before. From the definition of (\hat{S}, \hat{g}) , the two payoffs (A.9) and (A.12) are the same. Hence, (\hat{S}, \hat{g}) is also optimal.

Proof of Proposition 2. From the proof of Lemma 7, the regulator's problem reduces to finding a disclosure rule (S, g) to maximize (A.9), such that $\mu(s) \geq \rho_i$ for every $i \in \{1, \dots, k\}$ and $s \in S_i$. Moreover, if (S, g) is optimal, a disclosure rule with $k + 1$ scores as defined in Lemma 7 is also optimal. Hence, the regulator's problem reduces to finding a disclosure rule $(\hat{S} = \{s_0, s_1, \dots, s_k\}, \hat{g})$ that maximizes (A.12), such that, for every $i \in \{1, \dots, k\}$, the following holds: (i) $\hat{g}(s_i|\theta) = 0$ if $\theta > \theta_i$; and (ii) $\hat{\mu}(s_i) \geq \rho_i$ if $\sum_{\theta \in \Theta} p(\theta) \hat{g}(s_i|\theta) > 0$.

Since $h_i(\theta) = \hat{g}(s_i|\theta)$ (by definition), $\hat{\mu}(s_i) \geq \rho_i$ reduces to (11), using (2). It is then easy to show that if (\hat{S}, \hat{g}) solves the problem above, then $h_i(\theta) = \hat{g}(s_i|\theta)$ solves the problem in Proposition 2; and vice versa, if h solves the problem in Proposition 2, then $\hat{g}(s_i|\theta) = h_i(\theta)$ solves the problem above.

Proof of Corollary 3. No disclosure reduces to setting $h_1(\theta) = 1$ for every $\theta \in \Theta$. If $E(\tilde{\theta}) \geq \rho_1$, this satisfies constraint (11) and is optimal because it achieves the maximal attainable value for (10) given constraint (12). If $E(\tilde{\theta}) < \rho_1$, constraint (11) is violated, and so, a policy of no disclosure cannot be optimal.

Full disclosure reduces to setting $h_i(\theta) = 1$ if $\theta = \theta_i$ and setting $h_i(\theta) = 0$ if $\theta \neq \theta_i$ ($i \in \{1, \dots, k\}$). If $\theta_{\min} \geq 1$, this achieves the highest attainable utility. If $\theta_{\min} < 1$, then since

$\theta_1 = \theta_{\max} > 1$, full disclosure cannot be optimal because increasing $h_1(\theta_{\min})$ by a small amount satisfies constraint (11) and increases the value of the objective function.

Proof of Part 1 in Proposition 3. The proof is by induction. From Lemma 6, type θ_1 sells its asset with probability 1. So, from (13), any solution to the regulator’s problem satisfies $h_1(\theta_1) = 1$.

Next, for every $z \in \{2, \dots, k\}$, we show that, if there is a solution that satisfies $h_i(\theta_i) = 1$ for every $i \in \{1, \dots, z - 1\}$, there is also a solution that satisfies $h_i(\theta_i) = 1$ for every $i \in \{1, \dots, z\}$. Consider a solution h that satisfies $h_i(\theta_i) = 1$ for every $i \in \{1, \dots, z - 1\}$, but not for $i = z$. We construct an alternate solution \tilde{h} , such that $\tilde{h}_i(\theta_i) = 1$ for every $i \in \{1, \dots, z\}$.

Specifically, for every $i < z$, let Δ_i be the lowest $h_i(\theta_z)$ that satisfies constraint i in (11). That is, if $\theta_z \leq \rho_i$ (so θ_z takes resources), then $\Delta_i = 0$; and if $\theta_z > \rho_i$ (so θ_z provides resources), then

$$\Delta_i = \max\left\{0, \frac{\sum_{\theta \in \Theta: \theta \neq \theta_z} p(\theta)(\rho_i - \theta)h_i(\theta)}{p(\theta_z)(\theta_z - \rho_i)}\right\}. \tag{A.13}$$

Observe that, for every $i < z$ such that $\theta_z > \rho_i$, there exists a set of types $\Theta'_i \subset \{\theta \in \Theta : \theta < \rho_i\}$ and numbers $\Delta_{i\theta} \in [0, h_i(\theta)]$, such that

$$p(\theta_z)(\theta_z - \rho_i)\Delta_i = \sum_{\theta \in \Theta'_i} p(\theta)(\rho_i - \theta)\Delta_{i\theta}. \tag{A.14}$$

That is, type θ_z provides resources, which are used by the types in the set Θ'_i . (Note that $\Delta_i = 0$ implies that $\Delta_{i\theta} = 0$ for every $\theta \in \Theta'_i$.)

To construct the alternate solution, start with $\tilde{h} = h$ and make the following changes. First, for every $i < z$, reduce $\tilde{h}_i(\theta_z)$ to Δ_i and transfer the mass to $\tilde{h}_z(\theta_z)$. Following this change, all the constraints continue to be satisfied, and the value of the objective function is unchanged. Next, make the following two changes for every $i < z$. First, reduce $\tilde{h}_i(\theta_z)$ to 0 and transfer the mass to $\tilde{h}_z(\theta_z)$. Second, for every $\theta \in \Theta'_i$, reduce $\tilde{h}_i(\theta)$ by $\Delta_{i\theta}$ and transfer the mass to $\tilde{h}_z(\theta)$. Following this change, constraint i in (11) continues to be satisfied and constraint z is relaxed (strictly relaxed if $\rho_i > \rho_z$). The last observation follows because $i < z$ implies that $\rho_i \geq \rho_z$, which implies that

$$p(\theta_z)(\theta_z - \rho_z)\Delta_i \geq p(\theta_z)(\theta_z - \rho_i)\Delta_i \tag{A.15}$$

and

$$\sum_{\theta \in \Theta'_i} p(\theta)(\rho_z - \theta)\Delta_{i\theta} \leq \sum_{\theta \in \Theta'_i} p(\theta)(\rho_i - \theta)\Delta_{i\theta}. \tag{A.16}$$

So, using (A.14),

$$p(\theta_z)(\theta_z - \rho_z)\Delta_i \geq \sum_{\theta \in \Theta'_i} p(\theta)(\rho_z - \theta)\Delta_{i\theta}. \tag{A.17}$$

It is easy to verify that all other constraints continue to be satisfied and that each \tilde{h}_i continues to be a function from Θ to $[0, 1]$. After repeating the above process for every $i < z$, we obtain that $\tilde{h}_z(z) = 1$. Hence, we constructed a solution \tilde{h} that satisfies $\tilde{h}_i(\theta_i) = 1$ for every $i \in \{1, \dots, z\}$.

Proof of Part 2 in Proposition 3. For a given solution h and $j \in \{1, \dots, k\}$, denote $\alpha_j(h) = \theta_{\min}$ and $\beta_j(h) = \theta_{\max}$, if $h_j(\theta) = 0$ for every $\theta < 1$; and $\alpha_j(h) = \max\{\theta \in \Theta : \theta < 1 \text{ and } h_j(\theta) > 0\}$ and $\beta_j(h) = \min\{\theta \in \Theta : \theta < 1 \text{ and } h_j(\theta) > 0\}$, otherwise. That is, $\alpha_j(h)$ is the highest type

below 1 that obtains score s_j , and, if no such type exists, then $\alpha_j(h) = \theta_{\min}$. Similarly, $\beta_j(h)$ is the lowest type below 1 that obtains score s_j , and, if no such type exists, then $\beta_j(h) = \theta_{\max}$.

To prove the proposition, we prove an equivalent result that there is a solution \tilde{h} that satisfies $\alpha_j(\tilde{h}) \leq \beta_i(\tilde{h})$ for every $i, j \in \{1, \dots, k\}$, such that $i > j$. The last result follows from Lemma A.1 below. To see that, suppose there is a solution h , such that $\beta_i(h) < \alpha_j(h)$ for some $i > j$. Since there is a finite number of types, we can repeat Lemma A.1 to construct a solution h' , such that $\alpha_j(h') \leq \beta_i(h')$. Moreover, if $\alpha_{j'}(h) \leq \beta_{i'}(h)$ for some $i' > j'$, we continue to have $\alpha_{j'}(h') \leq \beta_{i'}(h')$. We can repeat the process above for any pair i, j , for which $\beta_i(h) < \alpha_j(h)$. Since there is a finite number of scores, there is a finite number of pairs (i, j) . So, we end up with a solution that satisfies the desired property.

Lemma A.1. Consider a solution h , and suppose there exists $i > j$, such that $\beta_i(h) < \alpha_j(h)$. Then there is another solution \tilde{h} that satisfies the following: (i) $\alpha_j(\tilde{h}) \leq \alpha_j(h)$ and $\beta_i(\tilde{h}) \geq \beta_i(h)$, with at least one strict inequality; and (ii) $\alpha_s(\tilde{h}) = \alpha_s(h)$ and $\beta_s(\tilde{h}) = \beta_s(h)$, for every $s \notin \{i, j\}$.

Proof. Since $\beta_i(h) < \alpha_j(h)$, there exist types $\theta_v, \theta_w < 1$, such that $\theta_v = \beta_i(h)$ and $\theta_w = \alpha_j(h)$. So, $\theta_v < \theta_w < 1 \leq \theta_i < \theta_j$. In addition, $h_i(\theta_v) > 0$, and $h_j(\theta_w) > 0$. Let

$$\Delta = \min\{h_j(\theta_w), h_i(\theta_v) \frac{p(\theta_v)(\rho_j - \theta_v)}{p(\theta_w)(\rho_j - \theta_w)}\}. \tag{A.18}$$

To construct \tilde{h} , start with $\tilde{h} = h$. For type θ_w , reduce $h_j(\theta_w)$ and increase $h_i(\theta_w)$, both by Δ . This relaxes the constraint for score j by $\Delta p(\theta_w)(\rho_j - \theta_w)$ and tightens the constraint for score i by $\Delta p(\theta_w)(\rho_i - \theta_w)$. For type θ_v , reduce $h_i(\theta_v)$ and increase $h_j(\theta_v)$, both by $\frac{p(\theta_w)(\rho_j - \theta_w)}{p(\theta_v)(\rho_j - \theta_v)} \Delta$. (Observe that $h_i(\theta_v) - \frac{p(\theta_w)(\rho_j - \theta_w)}{p(\theta_v)(\rho_j - \theta_v)} \Delta \geq 0$.) This relaxes the constraint for score i by $\frac{p(\theta_w)(\rho_j - \theta_w)}{p(\theta_v)(\rho_j - \theta_v)} \Delta p(\theta_v)(\rho_i - \theta_v)$ and tightens the constraint for score j by $\frac{p(\theta_w)(\rho_j - \theta_w)}{p(\theta_v)(\rho_j - \theta_v)} \Delta p(\theta_v)(\rho_j - \theta_v)$. Overall, the net effect on the constraint for score j is zero:

$$\Delta p(\theta_w)(\rho_j - \theta_w) - \frac{p(\theta_w)(\rho_j - \theta_w)}{p(\theta_v)(\rho_j - \theta_v)} \Delta p(\theta_v)(\rho_j - \theta_v) = 0. \tag{A.19}$$

As for the constraint for score i , the net effect is nonnegative (strictly positive if $\rho_i < \rho_j$):

$$\begin{aligned} & -\Delta p(\theta_w)(\rho_i - \theta_w) + \frac{p(\theta_w)(\rho_j - \theta_w)}{p(\theta_v)(\rho_j - \theta_v)} \Delta p(\theta_v)(\rho_i - \theta_v) \\ & = \Delta p(\theta_w)(\rho_i - \theta_w) \frac{\rho_i - \theta_w}{\rho_j - \theta_v} \left[\frac{\rho_j - \theta_w}{\rho_i - \theta_w} - \frac{\rho_j - \theta_v}{\rho_i - \theta_v} \right] \geq 0. \end{aligned} \tag{A.20}$$

The last inequality follows because if $\rho_i < \rho_j$, the ratio $\frac{\rho_j - \theta}{\rho_i - \theta}$ is strictly increasing in θ . (If $\rho_i = \rho_j$, the ratio is constant.) Note that \tilde{h} satisfies all other constraints and keeps the value of the objective function unchanged. Hence, \tilde{h} solves the regulator’s problem. Moreover, either $\tilde{h}_i(\theta_v) = 0$ or $\tilde{h}_j(\theta_w) = 0$. Hence, $\alpha_j(\tilde{h}) \leq \alpha_j(h)$ and $\beta_i(\tilde{h}) \geq \beta_i(h)$, with at least one strict inequality. Also, since $\tilde{h}_s(\theta) = h_s(\theta)$ for every $s \notin \{i, j\}$ and $\theta \in \Theta$, it follows that $\alpha_s(\tilde{h}) = \alpha_s(h)$ and $\beta_s(\tilde{h}) = \beta_s(h)$ for every $s \notin \{i, j\}$.

As an aside, note that \tilde{h} decreases the probability that the bank obtains the higher score s_j and increases the probability that the bank obtains the lower score s_i , both by $p(\theta_w)\Delta -$

$p(\theta_v) \frac{p(\theta_w)(\rho_j - \theta_w)}{p(\theta_v)(\rho_j - \theta_v)} \Delta$, which is positive (since $\theta_w > \theta_v$). Moreover, under \tilde{h} , the price for score s_i increases because the constraint for that score is relaxed, but the average sale price across all types remains unchanged. Hence, if the bank’s payoff function (1) was concave, \tilde{h} would increase the value of the objective function.

Proof of Lemma 8. Part 1: Summing up all k constraints in (11) and changing the order of summation, we obtain $\sum_{\theta \in \Theta} \sum_{i=1}^k p(\theta)(\theta - \rho_i)h_i(\theta) \geq 0$. Since $\rho_i \geq 1$ for every $i \in \{1, \dots, k\}$, it follows that $\sum_{\theta \in \Theta} p(\theta)(\theta - 1) \sum_{i=1}^k h_i(\theta) \geq 0$. But since $E(\tilde{\theta}) < 1$, it follows that $\sum_{\theta \in \Theta} p(\theta)(\theta - 1) < 0$. Hence, we must have a type $\theta < 1$ for which $\sum_{i=1}^k h_i(\theta) < 1$.

Part 2: Suppose h solves the regulator’s problem, and suppose, to the contrary, that one of the constraints in (11), say constraint j , is satisfied with strict inequality. From part 1, there exists a type $\theta \in \Theta$ for which $\sum_{i=1}^k h_i(\theta) < 1$. Construct \tilde{h} from h by increasing $h_j(\theta)$ slightly. Since $\theta < 1$, constraint (13) continues to be satisfied. Then \tilde{h} solves the regulator’s problem and strictly increases the value of the objective function. But this contradicts the optimality of h .

Proof of Corollary 4. As explained in the text, if $E(\tilde{\theta}) < 1$, then $x(s_i) = \rho_i$ for every $i \in \{1, \dots, k\}$, since constraint (11) is binding. Moreover, if $\rho_1 > \rho_2 > \dots > \rho_k$, then any solution to the regulator’s problem satisfies the two properties in Proposition 3. More generally, we can show that, if $E(\tilde{\theta}) < 1$, then any solution satisfies the following:

1. For every $i \in \{1, \dots, k\}$, $h_j(\theta_i) > 0$ only if $\rho_i = \rho_j$.
2. $h_i(\theta_v)h_j(\theta_w) = 0$, for every v, ω, i, j , such that $\theta_v < \theta_w < 1 < \theta_i < \theta_j$ and $\rho_i < \rho_j$.

Hence, if (S, g) is an optimal disclosure rule, the corresponding solution (\hat{S}, \hat{g}) from Lemma 7 must satisfy the two properties above. The results then follow easily.

Proof of Corollary 5. From Corollary 4, type θ_i sells for price ρ_i ($i \in \{1, \dots, k\}$). So, if $\rho_1 > \rho_2 > \dots > \rho_k$, every type above 1 must obtain a different score. Moreover, since $E(\tilde{\theta}) < 1$, we know from Lemma 8 that there exists $\theta < 1$, which keeps its asset with a positive probability. To implement this, the regulator must assign at least one additional null score.

Proof of Proposition 4. Consider a disclosure rule (S, g) that satisfies the two assumptions in Section 6. Suppose (S, g) is optimal.

Part 1: We first show that there exist two cutoffs $z_L \leq z_H$, such that types below z_L (if any) do not sell, types that belong to the interval $[z_L, z_H]$ sell for the same price, and types above z_H (if any) sell for a higher price (or prices).

Suppose the lowest type that sells is $\check{\theta}$. (Since $\theta_{\max} > 1$, such a type exists.) Suppose $\check{\theta}$ obtains score s' , and suppose that the highest type that obtains score s' and sells upon obtaining that score is θ' . Since the market breaks even, on average, we must have $x(s') \leq \theta'$. Moreover, from Lemma 5, $x(s') \geq \rho(\theta')$. Hence, $\theta' \geq \rho(\theta')$. Hence, $\theta' \geq 1$, because $\theta' < 1$ would imply that $\theta' < \rho(\theta')$. Hence, $\rho(\theta') \geq 1$. Hence, $x(s') \geq 1$.

Suppose the highest type that sells at price $x(s')$ is $\hat{\theta}$. ($\hat{\theta} \geq \theta'$.) Then, the expected equilibrium payoff for types $\check{\theta}$ and $\hat{\theta}$ is $x(s') + r$. Since the equilibrium payoff is weakly increasing in type, all types between $\check{\theta}$ and $\hat{\theta}$ also end up with an expected payoff $x(s') + r$. Hence, all types between $\check{\theta}$ and $\hat{\theta}$ must sell at price $x(s')$. To see that, note that, if a type $\theta \in (\check{\theta}, \hat{\theta})$ sold for price $x \neq x(s')$, its expected payoff would be $x + r \neq x(s') + r$; and if a type $\theta \in (\check{\theta}, \hat{\theta})$ did not sell, it would end up with $\theta + r \Pr(\tilde{\varepsilon} \geq 1 - \theta)$, which is strictly less than $\hat{\theta} + r \Pr(\tilde{\varepsilon} \geq 1 - \hat{\theta})$, which is less than or equal to $x(s') + r$, since $\hat{\theta}$ is selling its asset.

Since $\hat{\theta} \geq 1$, all types above $\hat{\theta}$ also sell because, if this were not the case, the regulator could increase the value of the objective function, without violating the constraints, by giving each type above $\hat{\theta}$ its own score. From the way we defined $\hat{\theta}$, types above $\hat{\theta}$ (if exist) sell at prices above $x(s')$. Letting $z_L = \hat{\theta}$ and $z_H = \hat{\theta}$, we have established the first part.

Next, observe that Lemma 6 continues to hold in the case under consideration. Hence, for every score s that induces selling at price $x(s')$, $\mu(s) = x(s')$. By aggregating across all these scores, and since $x(s') \geq \rho(z_H)$, we obtain that $\sum_{\theta \in [z_L, z_H]} p(\theta)[\theta - \rho(z_H)] \geq 0$ (using (2)). That is, the average cash flow for types selling at price $x(s')$ is at least $\rho(z_H)$. It then follows that the optimal cutoffs are $z_L = \hat{z}_L(z_H^*)$ and $z_H = z_H^*$. In particular, for a given z_H , setting $z_L < \hat{z}_L(z_H)$ will violate the constraint above, and setting $z_L > \hat{z}_L(z_H)$ will reduce the value of the objective function. The optimal z_H minimizes $\hat{z}_L(z_H^*)$. Any other z_H will reduce the value of the objective function because fewer types will sell their assets. (Note that given (10), the objective function is just a function of selling probabilities, not prices.)

Part 2: Suppose, to the contrary, that two types above z_H^* obtain the same score. Then, both sell at the same price, say x . Suppose θ_l and θ_h are the lowest and highest types, respectively, that sell at price x . ($\theta_l < \theta_h$) Repeating the arguments in part 1, it follows that $x \geq \rho(\theta_h)$ and every type that belongs to the interval $[\theta_l, \theta_h]$ sells at price x . Moreover, $\sum_{\theta \in [\theta_l, \theta_h]} p(\theta)[\theta - \rho(\theta_h)] \geq 0$. Hence, $\hat{z}_L(\theta_h) \leq \theta_l < \theta_h$. But this contradicts the fact that $\hat{z}_L(\theta_h) = \theta_h$.

As an aside, note that Part 2 holds, even if $\hat{z}_L(\theta) = \theta$ does not hold for the lowest type above z_H (but holds for all other types above z_H). Also note that, if $\hat{z}_L(\theta) < \theta$ for some type θ that is above the lowest type above z_H , then there is a solution in which type θ is pooled together (i.e., gets the same score and sells for the same price) with at least one type below it; but there is also a solution in which every type above z_2 gets its own score and sells for a different price.

Proof of Corollary 6. Under an optimal disclosure rule, all types above 1 sell. Since $\hat{z}_L(\theta) = \theta$ for every $\theta \geq \theta_{m-1}$, it follows that $\hat{z}_L(\theta) > \theta_m$ for every $\theta > \theta_m$. Hence, if $\theta_m < 1$, the optimal cutoffs are $z_L = z_H = \theta_{m-1}$. If $\theta_m \geq 1$, the optimal cutoffs are $z_L = z_H = \theta_m$. The result then follows from Proposition 4.

To see why footnote 32 is true, note that if $p(\theta) = \frac{1}{m}$ for every $\theta \in \Theta$ (i.e., all types have the same probability), the condition $\hat{z}_L(\theta) = \theta$ implies not only that type θ cannot be pooled together with the type just below it, but also that type θ cannot be pooled together with any subset of types below it. So, assuming that $p(\theta) = \frac{1}{m}$ for every $\theta \in \Theta$ is sufficient to ensure that when the regulator must follow a deterministic rule, Corollary 6 continues to hold, even if the bank cannot freely dispose its assets.

Example 3. Suppose there are five types $\theta_1 > \theta_2 > 1 > \theta_3 > \theta_4 > \theta_5$, equation (15) holds, and $\rho_1 > \rho_2$. Then, $\hat{z}_L(\theta_2) = \theta_3$. There are three cases:

Case 1: $\sum_{i=1}^3 p(\theta_i)(\theta_i - \rho_1) > 0$. In this case, $\hat{z}_L(\theta_1) < \theta_3$, and it is optimal to set $z_H = \theta_1$ and $z_L < \theta_3$. Types below z_L obtain score s_0 and keep their assets. All other types obtain score s_1 and sell at price ρ_1 .

Case 2: $\sum_{i=1}^3 p(\theta_i)(\theta_i - \rho_1) < 0$. In this case, $\hat{z}_L(\theta_1) > \theta_3$, and it is optimal to set $z_H = \theta_2$ and $z_L = \theta_3$. Types below θ_3 keep their assets; type θ_1 sells at price ρ_1 ; types θ_2 and θ_3 sell at price ρ_2 .

Case 3: $\sum_{i=1}^3 p(\theta_i)(\theta_i - \rho_1) = 0$. In this case, $\hat{z}_L(\theta_1) = \theta_3$, and the regulator is indifferent between setting $z_H = \theta_1$ and $z_H = \theta_2$. In both cases, $z_L = \theta_3$.

Note that the condition $\sum_{i=1}^3 p(\theta_i)(\theta_i - \rho_1) > 0$ is equivalent to $p(\theta_1)(\theta_1 - \rho_1) > (\rho_1 - \rho_2) \sum_{i=2}^3 p(\theta_i)$. The term $p(\theta_1)(\theta_1 - \rho_1)$ represents the marginal benefit of choosing $z_H = \theta_1$

rather than $z_H = \theta_2$. The marginal benefit is that the type θ_1 provides resources to cross-subsidize lower types. The term $(\rho_1 - \rho_2) \sum_{i=2}^3 p(\theta_i)$ represents the marginal cost; that is, the resources that types θ_2 and θ_3 require are determined by type θ_1 's reservation price rather than by type θ_2 's reservation price.

References

- Admati, Anat, Pfleiderer, Paul, 2000. Forcing firms to talk: disclosure regulation and externalities. *Rev. Financ. Stud.* 13, 479–519.
- Akerlof, George A., 1970. The market for “Lemons”: quality uncertainty and the market mechanism. *Q. J. Econ.* 84, 488–500.
- Allen, Franklin, Gale, Douglas, 2000. Financial contagion. *J. Polit. Econ.* 108, 1–33.
- Alvarez, Fernando, Barlevy, Gadi, 2015. Mandatory Disclosure and Financial Contagion. NBER Working Paper No. 21328.
- Andolfatto, David, Berentsen, Aleksander, Waller, Christopher, 2014. Optimal disclosure policy and undue diligence. *J. Econ. Theory* 149, 128–152.
- Angeletos, George-Marios, Pavan, Alessandro, 2007. Efficient use of information and social value of information. *Econometrica* 75, 1103–1142.
- Bergemann, Dirk, Morris, Stephen, 2017. Information Design: a Unified Perspective. Cowles Foundation Discussion Paper No. 2075R.
- Beyer, Anne, Cohen, Daniel A., Lys, Thomas Z., Walther, Beverly R., 2010. The financial reporting environment: review of the recent literature. *J. Account. Econ.* 50, 296–343.
- Bond, Philip, Goldstein, Itay, 2015. Government intervention and information aggregation by prices. *J. Finance* 70, 2777–2812.
- Bouvard, Matthieu, Chaigneau, Pierre, De Motta, Adolfo, 2015. Transparency in the financial system: rollover risk and crises. *J. Finance* 70, 1805–1837.
- Calzolari, Giacomo, Pavan, Alessandro, 2006a. On the optimality of privacy in sequential contracting. *J. Econ. Theory* 130, 168–204.
- Calzolari, Giacomo, Pavan, Alessandro, 2006b. Monopoly with resale. *Rand J. Econ.* 37, 362–375.
- Chakraborty, Archishman, Harbaugh, Rick, 2007. Comparative cheap talk. *J. Econ. Theory* 132, 70–94.
- Dang, Tri Vi, Gorton, Gary, Holmström, Bengt, Ordóñez, Guillermo, 2017. Banks as secret keepers. *Am. Econ. Rev.* 107, 1005–1029.
- Diamond, Douglas, 1985. Optimal release of information by firms. *J. Finance* 60, 1071–1094.
- Dworczak, Piotr, 2017. Mechanism Design with Aftermarkets: Cutoff Mechanisms. Manuscript.
- Dworczak, Piotr, Martini, Giorgio, 2017. The Simple Economics of Optimal Persuasion. Manuscript.
- Elliott, Matthew, Golub, Benjamin, Jackson, Matthew O., 2014. Financial networks and contagion. *Am. Econ. Rev.* 104, 3115–3153.
- Faria-e-Castro, Miguel, Martinez, Joseba, Philippon, Thomas, 2017. Runs versus lemons: information disclosure, fiscal capacity and financial stability. *Rev. Econ. Stud.* 84, 1683–1707.
- Fishman, Michael, Hagerty, Kathleen, 1990. The optimal amount of discretion to allow in disclosure. *Q. J. Econ.* 105, 427–444.
- Fishman, Michael, Hagerty, Kathleen, 2003. Mandatory versus voluntary disclosure in markets with informed and uninformed customers. *J. Law Econ. Organ.* 19, 45–63.
- Flannery, Mark, Hirtle, Beverly, Kovner, Anna, 2017. Evaluating the information in the federal reserve stress tests. *J. Financ. Intermed.* 29, 1–18.
- Garcia, Daniel, Teper, Roe, Tsur, Matan, 2017. Information Design in Insurance Markets: Selling Peaches in a Market for Lemons. Working Paper.
- Gentzkow, Matthew, Kamenica, Emir, 2016. A Rothschild–Stiglitz approach to Bayesian persuasion. *Am. Econ. Rev.* 106, 597–601.
- Gick, Wolfgang, Pausch, Thilo, 2014. Bayesian Persuasion by Stress Test Disclosure. Working Paper. Deutsche Bundesbank.
- Goel, Anand, Thakor, Anjan V., 2015. Information reliability and welfare: a theory of coarse credit ratings. *J. Financ. Econ.* 115, 541–557.
- Goldstein, Itay, Huang, Chong, 2016. Bayesian persuasion in coordination games. *Am. Econ. Rev.* 106, 592–596.

- Goldstein, Itay, Sapra, Haresh, 2013. Should banks' stress test results be disclosed? An analysis of the costs and benefits. *Found. Trends Finance* 8, 1–54.
- Goldstein, Itay, Yang, Liyan, 2017. Information disclosure in financial markets. *Annu. Rev. Financ. Econ.* 9, 101–125.
- Gorton, Gary, 2015. Stress for success: a review of timothy Geithner's financial crisis memoir. *J. Econ. Lit.* 53, 975–995.
- Hirshleifer, Jack, 1971. The private and social value of information and the reward to inventive activity. *Am. Econ. Rev.* 61, 561–574.
- Innes, Robert D., 1990. Limited liability and incentive contracting with ex-ante action choices. *J. Econ. Theory* 52, 45–67.
- Inostroza, Nicolas, Pavan, Alessandro, 2017. Persuasion in Global Games with Application to Stress Testing. Manuscript.
- Kamenica, Emir, Gentzkow, Matthew, 2011. Bayesian persuasion. *Am. Econ. Rev.* 101, 2590–2615.
- Kartasheva, Anastasia, Yilmaz, Bilge, 2013. Precision of Ratings. Working Paper.
- Kolotilin, Anton, 2018. Optimal information disclosure: a linear programming approach. *Theor. Econ.* 13, 607–635.
- Leitner, Yaron, 2005. Financial networks: contagion, commitment, and private sector bailouts. *J. Finance* 60, 2925–2953.
- Leitner, Yaron, 2012. Inducing agents to report hidden trades: a theory of an intermediary. *Rev. Finance* 16, 1013–1042.
- Leitner, Yaron, 2014. Should Regulators Reveal Information About Banks? Federal Reserve Bank of Philadelphia *Business Review*, Third Quarter.
- Leitner, Yaron, Williams, Basil, 2017. Model Secrecy and Stress Tests. Federal Reserve Bank of Philadelphia Working Paper No. 17-41.
- Leitner, Yaron, Yilmaz, Bilge, 2018. Regulating a model. *J. Financ. Econ.* Accepted for publication.
- Lizzeri, Alessandro, 1999. Information revelation and certification intermediaries. *Rand J. Econ.* 30, 214–231.
- Marin, Jose M., Rahi, Rohit, 2000. Limited information revelation and market incompleteness. *Rev. Econ. Stud.* 67, 455–481.
- Mathevet, Laurent A., Perego, Jacopo, Taneva, Ina, 2017. On Information Design in Games. Manuscript.
- Mensch, Jeffrey, 2018. Monotone Persuasion. Manuscript.
- Monnet, Cyril, Quintin, Erwan, 2017. Rational opacity. *Rev. Financ. Stud.* 30, 4317–4348.
- Morgan, Donald P., Peristiani, Stavros, Savino, Vanessa, 2014. The information value of the stress test. *J. Money Credit Bank.* 46, 1479–1500.
- Morris, Stephen, Shin, Hyun Song, 2002. Social value of public information. *Am. Econ. Rev.* 92, 1521–1534.
- Morrison, Alan, White, Lucy, 2013. Reputational contagion and optimal regulatory forbearance. *J. Financ. Econ.* 110, 642–658.
- Nadarajah, Saralees, Kotz, Samuel, 2006. A truncated Cauchy distribution. *Int. J. Math. Educ. Sci. Technol.* 37, 605–608.
- Orlov, Dmitry, Zryumov, Pavel, Skrzypacz, Andrzej, 2017. Design of Macro-Prudential Stress Tests. Stanford University Graduate School of Business Research Paper No. 17-41.
- Ostrovsky, Michael, Schwarz, Michael, 2010. Information disclosure and unraveling in matching markets. *Am. Econ. J. Microecon.* 2, 34–63.
- Petrella, Giovanni, Resti, Andrea, 2013. Supervisors as information producers: do stress tests reduce bank opaqueness? *J. Bank. Finance* 37, 5406–5420.
- Prescott, Edward S., 2008. Should bank supervisors disclose information about their banks? *Fed. Reserv. Bank. Richmond Econ. Rev.* 94, 1–16.
- Shapiro, Joel, Skeie, David, 2015. Information management in banking crises. *Rev. Financ. Stud.* 28, 2322–2363.
- Verrecchia, Robert E., 2001. Essays on disclosure. *J. Account. Econ.* 32, 97–180.
- Williams, Basil, 2017. Stress Tests and Bank Portfolio Choice. Manuscript.