Stress Tests and Information Disclosure

Itay Goldstein (Wharton)   Yaron Leitner (Philadelphia Fed)[1]

September 2015

[1] The views expressed here are those of the authors and do not necessarily reflect those of the Federal Reserve Bank of Philadelphia or of the Federal Reserve System.

Motivation

- The paper is about whether a regulator should disclose information about banks.
- Very controversial. For example, with regards to disclosure of stress tests results:
  - Fed Governor Tarullo expresses support for wide disclosure as it "allows investors and other counterparties to better understand the profiles of each institution."
  - But the Clearing House Association is concerned of "unanticipated and potentially unwarranted and negative consequences to covered companies and U.S. financial markets." (WSJ, 2012)

This paper

- A new theory of (optimal) disclosure, focusing on the following tradeoff:
  - Disclosure harms risk sharing arrangements among banks. (Relates to Hirshleifer effect.)
  - But some disclosure may be necessary to prevent a market breakdown.
- We find that:
  - During normal times, no disclosure is optimal.
  - During bad times, some disclosure is necessary. We characterize its optimal form; e.g., under what conditions a simple cutoff rule is optimal.

Risk Sharing and Disclosure Policy

- In our model, risk sharing takes a simple form:
  - A bank has an asset that yields a random cashflow.
  - The bank can replace the random cash flow with a deterministic cashflow by selling the asset in a competitive market.
- The sale price – and hence the bank's ability to share risk – depends on the regulator's disclosure policy.
- The regulator does not inject money in our model. (We discuss extensions.)

Related Literature

- Bayesian persuasion games (e.g., Kamenica & Gentzkow, 2011)
- Disclosure
  - by regulator (e.g., Morris & Shin, 2002; Angeletos & Pavan, 2007; Prescott, 2008; Leitner, 2012; Bond & Goldstein, 2012; Bouvard, Chaigneau & de Motta, 2013; Shapiro & Skeie, 2013; Goldstein & Sapra, 2014; Gick and Pausch, 2014; Andolfatto, Berentsen, and Waller, 2014)
  - by firm (e.g., Diamond, 1985; Fishman & Hagerty, 1990, 2003; Adamati & Pflediderer, 2000)
  - by credit rating agencies (e.g., Lizzeri, 1999; Kartasheva & Yilmaz, 2012; Goel and Thakor, 2015)
- Market incompleteness based on Hirshleifer effect vs. adverse selection (Marin & Rahi, 2000)
- Financial networks (e.g., Allen & Gale, 2000; Leitner, 2005)

The model

- There is a bank, a regulator (planner), and a perfectly competitive market.
- The bank has an asset that yields $\tilde{\theta} + \tilde{\varepsilon}$. $\quad \tilde{\theta} \perp \tilde{\varepsilon}, \quad E(\tilde{\varepsilon}) = 0$
- The bank can sell its asset in the market for an amount x (derived endogenously).
- Everyone is risk neutral, and the risk-free rate is 0%.
- Hence, $x = E[\tilde{\theta} + \tilde{\varepsilon} \mid \text{market information}]$.
- Bank's final cash holding: $z = \begin{cases} x & \text{if bank sells asset} \\ \tilde{\theta} + \tilde{\varepsilon} & \text{if bank keeps asset} \end{cases}$

The model (cont'd)

- Bank's final payoff is

$$R(z) = \begin{cases} z & \text{if } z < 1 \\ z + r & \text{if } z \geq 1 \end{cases} \qquad\qquad (r > 0)$$

  - Several motivations: project, debt liability, bank run
  - Results hold for more general specifications.

- Bank maximizes $E\left[R(z)\middle|\text{ bank's information}\right]$.

The model (cont'd)

- $\tilde{\theta}$ is drawn from a finite set $\Theta \subset \mathbb{R}$ according to $p(\theta) = \Pr(\tilde{\theta} = \theta)$.
- $\tilde{\varepsilon}$ is drawn from a continuous cumulative distribution function F.
- Probability structure (i.e., functions p and F) is common knowledge.
- Assume: $\theta_{\max} \geq 1$, $F(1 - \theta_{\min}) < 1$, $F(1 - \theta_{\max}) > 0$.

The model (cont'd)

- Planner observes the realization of $\tilde{\theta}$ (denoted by θ).
- Market does not observe θ.
- As for the bank, we focus on 2 cases:
  1. Bank does not observe θ.
  2. Bank observes θ.
- In both cases, no one observes the realization of $\tilde{\varepsilon}$.

Disclosure rules

- Before observing θ, the planner chooses (and publicly announces) a disclosure rule.
- A disclosure rule is a set of "scores" S, and a function that maps each type to a distribution over scores. (Without loss, S is finite.)
- Denote

$$g(s|\theta) = \Pr(\tilde{s} = s|\tilde{\theta} = \theta)$$

$$\mu(s) = E[\tilde{\theta} + \tilde{\varepsilon}|\tilde{s} = s)] = \frac{\sum_{\theta \in \Theta} \theta p(\theta) g(s|\theta)}{\sum_{\theta \in \Theta} p(\theta) g(s|\theta)}$$

Disclosure rules (cont'd)

- The planner can commit to the chosen disclosure rule.

- Planner's objective: maximize expected total surplus.
- Same as maximizing bank's expected payoff across all types.

Sequence of events

1. The planner chooses a disclosure rule and publicly announces it.
2. The bank's type θ is realized and observed by the planner. (In case 2, θ is also observed by the bank.)
3. The planner assigns the bank a score s and publicly announces it.
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. So, z and $R(z)$ are determined.

- Essentially, a score is a price recommendation to the market.

Case 1: Bank does not observe its type

- Bank's action depends only on s, and so does not convey additional information to the market.
- Hence, the market sets a price $x(s) = \mu(s)$.
- Hence, in equilibrium the bank sells if and only if $\mu(s) \geq 1$. (Explain.)

Case 1: Bank does not observe its type

- Expected payoff for type θ, given disclosure rule (S, g):

$$u(\theta) = \sum_{s:\mu(s)<1} [\underbrace{\theta + r\,\Pr(\tilde{\varepsilon} \geq 1-\theta)}_{\text{bank keeps asset}}]g(s|\theta) + \sum_{s:\mu(s)\geq 1} [\underbrace{\mu(s) + r}_{\text{bank sells}}]g(s|\theta)$$

- The planner chooses (S, g) to maximize $\sum_{\theta\in\Theta} p(\theta)u(\theta)$.
- Same as maximizing

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

Case 1: Bank does not observe its type

- We can focus (without loss) on disclosure rules that assign at most two scores, s_1 and s_0, such that $\mu(s_1) \geq 1$ and $\mu(s_0) < 1$.

- $h(\theta)$: probability of obtaining the "high" score s_1.

Planner's problem

Lemma

The planner's problem reduces to choosing $h : \Theta \to [0,1]$ *to maximize*

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

subject to

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

- Constraint follows since $\mu(s_1) \geq 1$.

Solution to planner's problem

- If $E(\tilde{\theta}) \geq 1$, set $h(\theta) = 1$ for every $\theta \in \Theta$. ("normal" times)
- If $E(\tilde{\theta}) < 1$ ("bad" times), the solution depends on the gain-to-cost ratio:

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

  - For $\theta \geq 1$: set $h(\theta) = 1$
  - For $\theta < 1$: set $h(\theta) = 1$ to types with high $G(\theta)$, and $h(\theta) = 0$ to types with low $G(\theta)$

Solution to planner's problem

- If $E(\tilde{\theta}) \geq 1$, set $h(\theta) = 1$ for every $\theta \in \Theta$. ("normal" times)
- If $E(\tilde{\theta}) < 1$ ("bad" times), the solution depends on the gain-to-cost ratio:

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

  - For $\theta \geq 1$: set $h(\theta) = 1$
  - For $\theta < 1$: set $h(\theta) = 1$ to types with high $G(\theta)$, and $h(\theta) = 0$ to types with low $G(\theta)$

- Types that obtain the low score are not necessarily the lowest.

Implementation

- If $E(\tilde{\theta}) \geq 1$, the planner can give every type the same score (i.e., no disclosure)
  - It is also possible to give multiple scores, such that $\mu(s) \geq 1$ for every score.
  - If $\theta_{\text{min}} \geq 1$, we can even have full disclosure.
- If $E(\tilde{\theta}) < 1$, the planner must assign at least two scores. Yet, full disclosure is suboptimal.

Example 1 (“normal” times)

- $\tilde{\theta} \in \{0.8, \; 1.0, \; 1.2\}$, equal probabilities.
- With no disclosure, every type sells (for \$1) $->$ optimal.
- With full disclosure, only types 1 and 1.2 sell $->$ suboptimal.

Example 2 (“bad” times)

- $\tilde{\theta} \in \{0.6, \;\; 0.8, \;\; 1.0, \;\; 1.2\}$, equal probabilities.
- With no disclosure, no one sells (since average is 0.9).
- With full disclosure, only types 1 and 1.2 sell.

- Partial disclosure can do better (since more types sell).
  - $G(\theta)$ increasing –> high score to 0.8 1.0 1.2
  - $G(\theta)$ decreasing –> high score to 0.6 (with probability 0.5) 1.0 1.2

Case 2: Bank observes its type

- The solution so far (when bank does not observe its type) is close to Kamenica & Gentzkow (2011); but since we put more structure on the planner's objective, we can say more.
- The case in which the bank observes its type is harder (and new).
  - Now each type has its own "reservation price," i.e., a minimum price at which it is willing to sell.
- The planner may need to assign more than 2 scores to distinguish among types with different reservation prices.

Optimal disclosure rules

- ρ_1: reservation price of highest type
- If $E(\tilde{\theta}) \geq \rho_1$, no disclosure achieves the optimal outcome.
- If $E(\tilde{\theta}) < \rho_1$, some disclosure is necessary.

- Next, we focus on the case in which resources are scarce
  - I.e., it is impossible to implement an outcome in which every type sells with probability 1.
- In this case, if the highest type that obtains score s is $\theta_i > 1$, then $x(s) = \rho(\theta_i)$

Optimal disclosure rules (an example)

- Consider 2 types above 1 $(\theta_1 > \theta_2 > 1)$ with different reservation prices $(\rho_1 > \rho_2 \geq 1)$.

- First result: θ_1 and θ_2 must obtain different scores.
- "Proof":
  - If θ_1 and θ_2 obtain the same score, type θ_2 ends up with ρ_1.
  - This is a waste of resources, but without any gain.
  - Better to give type θ_2 its own score, so that it ends up with only ρ_2.

- Second result: Among the types below 1 that are pooled with types above 1, the lowest types below 1 are pooled with the highest types above 1.

Optimal disclosure rules (an example)

type	0.5	0.7	0.8	1.4	1.7
Sale price	n/a	1.2			1.2
			1.1	1.1	

Intuition:

- As before, the planner uses a gain-to-cost ratio to assign scores, but now the cost depends on the assigned score.

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

- Nonmonotonicity follows because it is relatively more costly to assign a high score to a high type. (That is, when $\rho_1 > \rho_2$, $\frac{\rho_1 - \theta}{\rho_2 - \theta}$ is increasing in θ.)

Will nonmonotonicity prevail if we enrich our model?

- Add a constraint that higher types must end up with higher expected equilibrium payoff
  - E.g., banks can freely dispose assets (Innes, 1990).
- If planner *can* randomize:
  - Lower types may continue to sell for higher prices, but they sell with probability that is less than 1.
  - Types above 1 may sell above their reservation prices.
- If planner *cannot* randomize:
  - Optimal rule becomes monotone and generally involves two cutoffs.
  - For some parameter values, full disclosure is uniquely optimal.

Comments

- Risk sharing can take a more complicated form.
- Model can capture externalities imposed by banks on the rest of society. (Hence, regulation is necessary.)
- In many cases, regulator's commitment would arise endogenously.
- Model can be used as benchmark to think of credit rating agencies.
- An interesting extension: regulator can 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.
- The results could be applied to other settings of Bayesian persuasion

Conclusion/ Implications

- If $E(\theta)$ is sufficiently high, no disclosure is necessary.
- Otherwise, some disclosure is needed to enable trade.
    - True even if banks do not have private information.
- In many cases, the weakest banks receive the lowest possible score and are out of the market. But more generally, use “gain-to-cost” ratio.
- When banks observe their types, more disclosure is needed.
- Low types receiving high scores can emerge as a socially optimal outcome.

Thank you!