Designing Agile Banking Supervision

Jeong Ho (John) Kim
Emory University

Kyungmin (Teddy) Kim
Emory University

Victoria Liu
Federal Reserve Bank of Boston

Noam Tanner
Federal Reserve Bank of Boston

July 2023

Introduction

- Banking supervision faces three key challenges.
  1. there is incomplete information about various risks within the economy spread across the supervisor and the banks
     - banks have very detailed views of their own portfolios, but they cannot look into the business lines of their peers
     - supervisor, on the other hand, is able to probe into portfolios across all of his supervised institutions, despite a lack of finer details
  2. there is a conflict of interest between the banks and the supervisor
     - banks tend to have greater risk appetite than the supervisor
- To achieve socially desirable outcome, supervisory authorities design their public messages to guide and monitor bank behaviors.
  - while the cost of supervisory objection is rigid given the legal setup, supervisory communication should be an agile response to the fluid informational dynamics.

This Paper

- We model banking supervision as a game of strategic communication, and solve for the supervisor's optimal communication strategy.
  - incomplete information about the state of the economy
    - each of the bank and the supervisor receives a private signal
  - a conflict of interest between the banks and the supervisor
    - the bank prefers high risk endeavors to conservative risk taking in every state of the economy
    - the supervisor prefers high risk endeavors only if the state of the economy is good
  - before the bank takes its action, the supervisor recommends how to act
    - "be aggressive regardless of your signal"
    - "be aggressive only if your signal is good"
    - "be conservative regardless of your signal"
  - after the bank takes its action, the supervisor can object to the action
    - costly change from aggressive risk level to conservative risk level

This Paper, Cont.

- We find that an increase in the bank's informational advantage (γ) has two distinct effects.
  1. the information effect (dominates when γ is small)
     - an increase in γ enables the supervisor to make more informed supervisory decisions when he can induce the bank to reveal its information and, therefore, improves welfare
  2. control dilution (dominates when γ is large)
     - an increase in γ reduces the probability that the bank thinks the supervisor will object to its aggressive risk-taking
     - this implies that the bank reveals information less frequently and, therefore, welfare deteriorates
- The welfare effect of increased private information in the hands of the private sector is non-monotonic!
- The bank's cost in case of supervisory objection ought to be set high.
  - intuitively, what the bank cares about is the "cost-adjusted" probability that the supervisor will object to its aggressive risk-taking

The Model

- There is a bank and a supervisor.
- The bank decides whether to take high risks ("Aggressive") or take low risks ("Conservative").
  - The payoff from high risk endeavors is u_ω for the bank and v_ω for the supervisor.
    - $\omega \in \{G, B\}$ is the state of the economy
  - The payoff from conservative risk taking is normalized to zero for both the bank and the supervisor.
    - all that matters is the *relative gains* from taking on high risk endeavors
  - We focus on the case where there is a conflict of interest between the bank and the supervisor:

$$u_G = u_B > 0 \text{ and } v_G > 0 > v_B.$$

1. the bank prefers high risk endeavors to conservative risk taking in every state of the economy
2. the supervisor prefers high risk endeavors only if the state of the economy is good

The Model, Cont.

- In this case, the payoff from aggressive risk-taking when $\omega = G$ can be further normalized to one for both the bank and the supervisor:

$$u_G = v_G = 1.$$

- Given our assumptions on payoffs, $u_B = 1$ and

$$v_B = -d, \ d > 0.$$

The Model: Incomplete Information

- Both the bank and the supervisor do not observe ω, but each of them receives a private signal about the state of the economy.
  - the bank's signal s takes one of two values, g or b:

$$\gamma = \Pr\left(s = g \,|\omega = G\right) = \Pr\left(s = b \,|\omega = B\right)$$

    - $\gamma \in \left(\frac{1}{2}, 1\right)$ implies that the signal is indeed informative about the state and the bank does not perfectly observe ω
  - the supervisor observes the probability t that $\omega = G$:

$$t \sim F_{[0,1]}$$

    - t is the supervisor's "type"
    - this formulation is equivalent to the standard one in which we would specify the prior probability t_0 that $\omega = G$ and the supervisor's signal $s' \in [\underline{s}, \overline{s}]$ that has CDF F_ω conditional on ω

The Model: After the Bank Decides on Its Risk Level

- The supervisor assesses the bank's risk-management practices.
  1. he observes its risk level
  2. decides whether to allow or object to its risk-management practices
- If the supervisor allows the bank's risk-management practices, it keeps its risk level intact as it chose.
- If the supervisor objects to the bank's risk-management practices, it is forced to readjust its risk level to be low.
  - in this case, the bank incurs a cost $c > 0$
    - c can represent the fact that the bank may be forced to sell its high-risk assets at fire sale prices
    - c can reflect the bank's cost of reputation loss
- An implicit assumption here is that the supervisor *never* objects to a conservative bank.

The Model: Before the Bank Decides on Its Risk Level

- The supervisor discloses information about his type t.
- The supervisor's communication strategy is modeled following the recent literature on information design.
  1. an arbitrary finite set M of messages
  2. a function $\pi : [0, 1] \rightarrow M$
     - $\pi(t)$ denotes the message that the supervisor of type t picks to send
- We let $F(\cdot \mid m)$ represent the bank's posterior belief distribution about the supervisor's type t after observing $m \in M$
- We let $\delta(m) \in \{0, 1\}$ denote the bank's (observed) risk level following message $m \in M$
  - 1 stands for "Aggressive"
  - 0 stands for "Conservative"

The Model: The Timing of the Game

1. The supervisor publicly commits to his communication strategy (M, π).
2. Nature chooses ω, the bank observes s, and the regulator observes t.
3. The supervisor discloses information about t according to his communication strategy.
4. The bank decides whether to take high risks ("Aggressive") or take low risks ("Conservative").
5. The supervisor assesses the bank's risk-management practices.
    - if the bank is aggressive, he decides whether to accept or object to its risk-management practices
6. Finally, the payoffs are realized.

Preliminaries: Supervisor's Policy

- Let q denote the probability that the supervisor thinks the state of the economy is good ($\omega = G$).
  - If the bank is aggressive, the supervisor's expected payoff is

$$qv_G + (1-q)\,v_B = q - (1-q)\,d.$$

  - If the bank is conservative, the supervisor's payoff is zero.
  - Hence, he allows the bank's aggressive risk-taking if and only if

$$q \geq \widehat{t} := \frac{d}{1+d}.$$

    - based solely on his private information, the supervisor allows aggressive risk-taking if and only if $t \geq \widehat{t}$

Preliminaries: Supervisor's Policy, Cont.

- In equilibrium, the bank can be
  - aggressive regardless of its signal
    - in this case, the supervisor will allow the bank to be aggressive iff $t \geq \widehat{t}$
  - aggressive if its signal was good and conservative if its signal was bad
    - in this case, the supervisor will learn that $s = g$ $(s = b)$ from observing that the bank is aggressive (conservative)
    - based on his type & $s = g$, the supervisor will allow the bank to be aggressive iff

$$\Pr\left(\omega = G \left| s = g, t\right.\right) = \frac{\gamma t}{\gamma t + (1-\gamma)(1-t)} \geq \widehat{t}$$

$$t \geq \underline{t} = \frac{(1-\gamma)\, d}{\gamma + (1-\gamma)\, d} \; (< \widehat{t})$$

  - conservative regardless of its signal

Supervisor's Problem

$$\max_{T^{(1,0)}, T^{(1,1)} \subset [0,1]} \int_{T^{(1,0)} \cap [\underline{t},1]} \left[\gamma t - (1-\gamma)(1-t) d\right] dF(t) + \int_{T^{(1,1)} \cap [\widehat{t},1]} \left[t - (1-t) d\right] dF(t)$$

- $t \in T^{(1,0)}$: guide the bank to be aggressive only if its signal is good
  - BUT object to aggressive risk-taking if $t \in T^{(1,0)} \cap [0, \underline{t})$
    - this *does* happen in equilibrium
- $t \in T^{(1,1)}$: guide the bank to be aggressive regardless of its signal
  - BUT object to aggressive risk-taking if $t \in T^{(1,1)} \cap [0, \widehat{t})$
    - this will *not* happen in equilibrium
- $t \in T^{(0,0)} := [0,1] \setminus \left(T^{(1,0)} \cup T^{(1,1)}\right)$: guide the bank to be conservative regardless of its signal

Supervisor's Problem, Cont.

subject to a set of incentive compatibility constraints for the bank:

$$\Pr\left(t \geq \underline{t} \,|\, s = g, m = (1,0)\right) \geq c/\left(1+c\right) \qquad (IC_g^{(1,0)})$$

$$\Pr\left(t \geq \underline{t} \,|\, s = b, m = (1,0)\right) \leq c/\left(1+c\right) \qquad (IC_b^{(1,0)})$$

$$\Pr(t \geq \widehat{t} \,|\, s = g, m = (1,1)) \geq c/\left(1+c\right) \qquad (IC_g^{(1,1)})$$

$$\Pr(t \geq \widehat{t} \,|\, s = b, m = (1,1)) \geq c/\left(1+c\right) \qquad (IC_b^{(1,1)})$$

$$\Pr(t \geq \widehat{t} \,|\, s = g, m = (0,0)) \leq c/\left(1+c\right) \qquad (IC_g^{(0,0)})$$

$$\Pr(t \geq \widehat{t} \,|\, s = b, m = (0,0)) \leq c/\left(1+c\right) \qquad (IC_b^{(0,0)})$$

$$\Pr\left(t \geq \underline{t} \,|\, s = b, m = (1,0)\right) = \frac{\int_{T^{(1,0)} \cap [\underline{t},1]} \left[(1-\gamma)\, t + \gamma\,(1-t)\right] dF\left(t\right)}{\int_{T^{(1,0)}} \left[(1-\gamma)\, t + \gamma\,(1-t)\right] dF\left(t\right)} \leq \widehat{p} = \frac{c}{1+c}$$

Solution to the Supervisor's Problem

- $t \in T^{(1,0)} = [0, \tau) \cup [\tau_* (\lambda^*), \tau^* (\lambda^*))$, where $\tau \leq \underline{t}$ and $[\tau_* (\lambda^*), \tau^* (\lambda^*)) \subset [\underline{t}, \overline{t})$: guide the bank to be aggressive only if its signal is good
  - BUT object to aggressive risk-taking if $t \in [0, \tau)$
    - the bank thinks $t \in [0, \tau)$ if it has $s = g$ and receives the message $(1, 0)$
    - the bank thinks $t \in [\tau_* (\lambda^*), \tau^* (\lambda^*))$ if it has $s = b$ and receives the message $(1, 0)$
  - $[\tau_* (\lambda^*), \tau^* (\lambda^*)) \nearrow [\underline{t}, \overline{t})$ as $\lambda^* \searrow 0$ and $[\tau_* (\lambda^*), \tau^* (\lambda^*)) \searrow \widehat{t}$ as $\lambda^* \nearrow \overline{\lambda}$
- $t \in T^{(1,1)} = [\tau^* (\lambda^*), 1]$: guide the bank to be aggressive regardless of its signal
  - AND *do* allow aggressive risk-taking
- $t \in T^{(0,0)} = [\tau, \tau_* (\lambda^*))$: guide the bank to be conservative regardless of its signal

Solution to the Supervisor's Problem, Cont.

- If the IC constraint is *not* binding, $\lambda^* = 0, \tau_* = \underline{t}, \tau^* = \overline{t}$ and solve

$$\int_{\underline{t}}^{\overline{t}} \left[(1-\gamma)\, t + \gamma\, (1-t)\right] dF\,(t) = c \int_0^{\tau} \left[(1-\gamma)\, t + \gamma\, (1-t)\right] dF\,(t)$$

for τ.

- If the IC constraint is binding, $\tau = \underline{t}$ and solve

$$\int_{\tau_*(\lambda^*)}^{\tau^*(\lambda^*)} \left[(1-\gamma)\, t + \gamma\, (1-t)\right] dF\,(t) = c \int_0^{\underline{t}} \left[(1-\gamma)\, t + \gamma\, (1-t)\right] dF\,(t) \tag{1}$$

for λ^*, which in turn implies τ^* and τ_*.

Proposition

If γ is sufficiently close to $\frac{1}{2}$, then the supervisor's expected payoff increases in γ. For γ sufficiently close to 1, the supervisor's expected payoff decreases in γ.

An Example

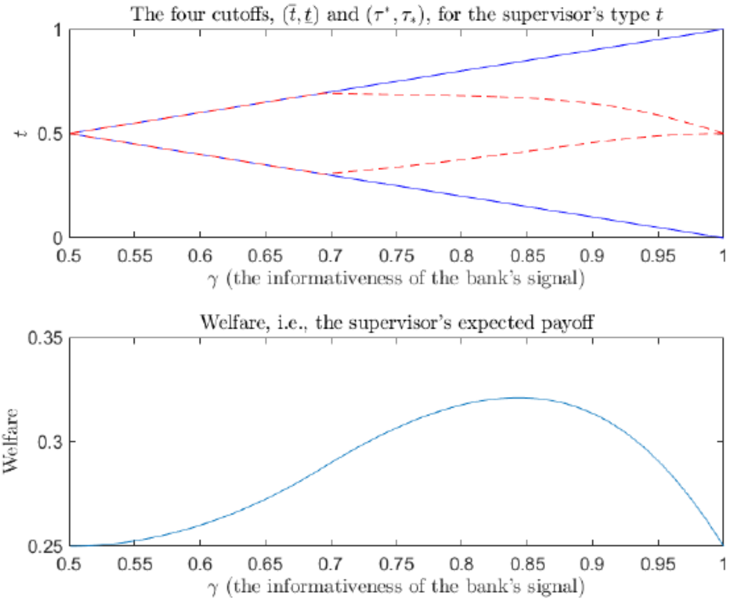

The four cutoffs, $(\bar{t}, \underline{t})$ and (τ^*, τ_*), for the supervisor's type t
1
0.5
0
t
0.5
0.55
0.6
0.65
0.7
0.75
0.8
0.85
0.9
0.95
1
γ (the informativeness of the bank's signal)
Welfare, i.e., the supervisor's expected payoff
0.35
0.3
0.25
Welfare
0.5
0.55
0.6
0.65
0.7
0.75
0.8
0.85
0.9
0.95
1
γ (the informativeness of the bank's signal)

An Example, Cont.

- $\bar{t}$ is upward-sloping, while $\underline{t}$ is downward-sloping
  - an increase in γ expanding the *ideal* information-acquision region $[\underline{t}, \bar{t})$.
- τ^* is initially $\nearrow$ but eventually $\searrow$, τ_* is initially $\searrow$ but eventually $\nearrow$
  - an increase in γ is initially expanding but eventually shrinking the *actual* information-acquision region $[\tau_*, \tau^*)$.
  - As γ increases, notice that the default-objection region $[0, \underline{t})$ contracts, which makes it more challenging to satisfy the incentive constraint (1).
    - This forces the supervisor to eventually shrinking the actual IA region despite the ever expanding ideal IA region.
- Welfare
  1. $\uparrow$ in γ initially expands (eventually shrinks) the actual IA region $[\tau_*, \tau^*)$
     - the bank reveals information more (less) frequently
     - therefore, welfare improves (deteriorates).
  2. $\uparrow$ in γ has another effect of enabling the supervisor to make more informed decisions within the actual IA region
     - therefore, welfare continues to improve beyond the point at which $[\tau_*, \tau^*)$ starts to shrink.

Discussion

$$\int_{\tau_*(\lambda^*)}^{\tau^*(\lambda^*)} [\gamma - (2\gamma - 1)\, t]\, dF(t) = c \int_0^{\underline{t}} [\gamma - (2\gamma - 1)\, t]\, dF(t)$$

$$\Pr\left(t \geq \underline{t}, s = b, m = (1,0)\right) = c \Pr\left(t < \underline{t}, s = b, m = (1,0)\right)$$

- We find that an increase in the bank's informational advantage (γ) has two distinct effects.
  1. the information effect (dominates when γ is small)
     - an increase in γ enables the supervisor to make more informed supervisory decisions when he can induce the bank to reveal its information and, therefore, improves welfare
  2. control dilution (dominates when γ is large)
     - an increase in γ reduces the probability that the bank thinks the supervisor will object to its aggressive risk-taking
     - this implies that the bank reveals information less frequently and, therefore, welfare deteriorates

Comparison to the Cheap-Talk Equilibrium

Proposition

There always exists a cheap-talk equilibrium in which the supervisor sends the message $m = (1, 1)$ if he is of type $t \geq \widehat{t}$ and $m = (0, 0)$ otherwise; the bank is aggressive if and only if it receives message $m = (1, 1)$, in which case the supervisor will allow it to be aggressive.

- Notice that in the limit as $\gamma \longrightarrow \frac{1}{2}$ or $\gamma \longrightarrow 1$, the supervisor's expected payoff shrinks to his expected payoff in the cheap-talk equilibrium presented in the above proposition.
  - in the former case, the supervisor chooses not to induce the bank to act on its own information
  - in the latter case, he cannot induce the bank to act on its own information
- Compared to this cheap-talk equilibrium, commitment power on the supervisor's side improves welfare as long as the bank has some but not perfect information about the state.

Comparison to the Cheap-Talk Equilibrium, Cont.

Proposition

Suppose that the IC constraint is not binding. Then the supervisor's unconstrained optimum is a cheap-talk equilibrium: the supervisor sends the message $m = (1, 1)$ if he is of type $t \geq \overline{t}$, $m = (1, 0)$ if he is of type $t \in [0, \tau) \cup [\underline{t}, \overline{t})$ for some $\tau \in [0, \underline{t}]$, and $m = (0, 0)$ otherwise.

- In light of this proposition, we conclude that a sufficient condition for the supervisor's commitment power to improve welfare is that $\gamma \in (\gamma_*, 1)$, where $\gamma_* \in \left(\frac{1}{2}, 1\right)$ such that $\gamma > \gamma_*$ implies the IC constraint is binding:
  1. $\gamma > \gamma_*$ ensures that the supervisor is actually using the commitment power vested in him
  2. $\gamma < 1$ ensures that the supervisor has plausible deniability he will not always allow aggressive risk-taking even if he is of type $t \in [\underline{t}, \overline{t})$
- Note that welfare improvements from commitment power is still non-monotonic in γ.

Bank's Cost in Case of Supervisory Objection

- One crucial lesson from our analysis is that whether the IC constraint for the bank is binding or not plays a key role in determining the welfare implications of more information on the bank's side.
- Looking at

$$\int_{T_{\text{Allow}}^{(1,0)}} [(1-\gamma)\, t + \gamma\, (1-t)]\, dF\,(t) \leq c \int_0^{\underline{t}} [(1-\gamma)\, t + \gamma\, (1-t)]\, dF\,(t)\,, \tag{2}$$

it is immediate that increasing c relaxes it, ergo improving welfare until the IC constraint is no longer binding.

Proposition

Let $v\,(c)$ denote the supervisor's maximal attainable payoff when the bank is faced with a cost c in case of supervisory objection. Then $v\,(c)$ is strictly increasing for $c \in [0, c^*)$ and is equal to $v\,(c^*)$ for $c \geq c^*$, where $c^* > 0$ is the value of c such that (2) holds with equality.

Bank's Cost in Case of Supervisory Objection, Cont.

- Intuitively, the bank is worried about
  - not only how frequently the supervisor will object to its aggressive risk taking ($\int_0^{t} [(1-\gamma)\, t + \gamma\, (1-t)]\, dF\, (t)$)
  - but also how costly those supervisory objections will be (c)
    - so increasing c can offset the control-dilution effect of increased γ
- Our analysis taking c as given reflects the fact that the supervisor can be agile in his communication strategy, but he cannot freely adjust the bank's cost in case of supervisory objection.
  - Yet the supervisor does have the power to occasionally change such costs for the bank by passing legislation to promote financial stability.
    - e.g., the Dodd-Frank Act made all banks with assets above $50 billion subject to a much more aggressive supervisory regime, effectively raising c for mid-sized banks;
    - in 2018, Congress scaled back Dodd-Frank, raising the threshold for increasing scrutiny of banks from $50 billion to $240 billion, effectively reducing c for mid-sized banks.

Bank's Cost in Case of Supervisory Objection, Cont.

- To the extent that the supervisor has some control over the bank's cost, the proposition has an important policy implication.
  - It is optimal to err on the side of giving the supervisor too much power *in case he finds that the bank does not meet supervisory expectations.*
    - if c is too high, the supervisor could simply scale back how frequently he will object to aggressive risk-taking after having sent $m = (1, 0)$
    - if c is too low, not only is the supervisor's unconstrained optimum infeasible (leaving welfare on the table), but the economy is exposed to experiencing a welfare loss in case the bank experiences a sudden boost in its private information

Additional Commitment to the Supervisory Ruling

- Our baseline model does not give the supervisor commitment power over his follow-up supervisory ruling.
  - the supervisor allows aggressive risk-taking only when it is ex post efficient:
    - he will allow the bank's aggressive risk-taking if and only if he is of type $t \geq \underline{t}$ ($t \geq \widehat{t}$) after having sent $m = (1, 0)$ ($m = (1, 1)$)
- We now turn attention to the case where the supervisor also has commitment power over his follow-up supervisory ruling.
  - the supervisor can commit a priori to allowing (objecting to) aggressive risk-taking even if it is ex post inefficient
    - e.g., he will object to the bank's aggressive risk-taking if he is of type $t \in T^{(1,0)}_{\text{Object}} \cap [\underline{t}, \overline{t})$ although he prefers ex post to allow it

Additional Commitment to the Supervisory Ruling, Cont.

- As in the baseline model, it continues to hold that $T^{(1,0)}_{\text{Allow}} = [\tau_{**}, \tau^{**})$ for some $\tau_{**} \in (\underline{t}, \widehat{t})$ and $\tau^{**} \in (\widehat{t}, \overline{t})$.
- In contrast to our baseline model, it is straightforward to prove that $\tau = \tau_{**}$—$T^{(1,0)}_{\text{Object}} = [0, \tau_{**})$ and $T^{(1,0)} = [0, \tau^{**})$.
  - in the baseline model, $T^{(1,0)}_{\text{Object}} = [0, \underline{t})$ and $T^{(0,0)} = [\underline{t}, \tau_*)$.
  - intuitively, the supervisor of type $t \in [\underline{t}, \tau_*)$ is tempted to respect the bank's decisions if they were reflective of its signal
    - sending $m = (1, 0)$ in this region would make the IC constraint even more binding, so he resorted to sending $m = (0, 0)$ instead
  - now, the supervisor is able to put this region to good use with the additional commitment power vested in him
    - he can overcome the temptation to respect the bank's risk-taking decision if t turns out to be in $[\underline{t}, \tau_*)$ by committing to object to aggressive risk-taking in this region even after having sent $m = (1, 0)$

Additional Commitment to the Supervisory Ruling, Cont.

Proposition

The supervisor's expected payoff with additional commitment is strictly monotone-increasing in γ on the interval $\left(\frac{1}{2}, 1\right)$.

- The above proposition shows that, unlike in our baseline model, the supervisor can do strictly better than with cheap talk even as $\gamma \longrightarrow 1$.
- The proposition shows that, with additional commitment to the supervisory ruling, more information does result in higher welfare.
  - The policy implication is that it is important to give the supervisor enough commitment power, particularly in the form of supervisory ruling.
    - Commitment power over how much he discloses about his own information alone can be impotent

An Example

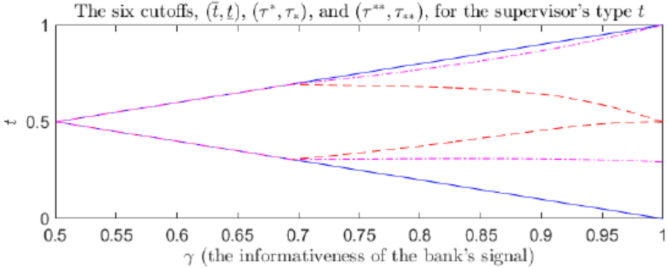

The six cutoffs, $(\bar{t}, \underline{t})$, (τ^*, τ_*), and (τ^{**}, τ_{**}), for the supervisor's type t
1
0.5
0
t
0.5
0.55
0.6
0.65
0.7
0.75
0.8
0.85
0.9
0.95
1
γ (the informativeness of the bank's signal)

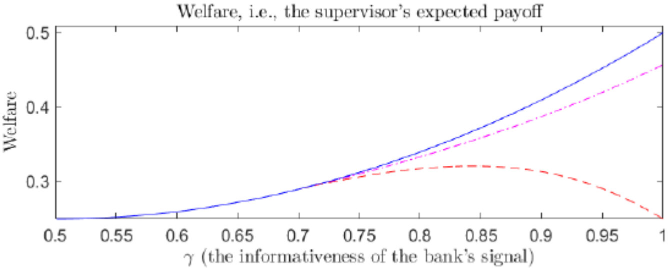

Welfare, i.e., the supervisor's expected payoff
0.5
0.4
0.3
Welfare
0.5
0.55
0.6
0.65
0.7
0.75
0.8
0.85
0.9
0.95
1
γ (the informativeness of the bank's signal)

An Example, Cont.

- As discussed above, the figure confirms that, unlike in the baseline model, the supervisor can do strictly better than with cheap talk even in the limit as $\gamma \to 1$.
- While it still is the case that the supervisor cannot attain his unconstrained optimum, the figure shows that the supervisor can do surprisingly well even in the limit as $\gamma \to 1$
  - the welfare gap from the unconstrained optimum is visibly small
  - it is easy to check that, as shown in the figure, $\tau^{**} \to 1$ in this limit
    - so the supervisor can induce the bank to reveal its information whenever he is optimistic enough
- Thus, it vividly reinforces the policy implication of the last proposition
  - it is important to give the supervisor enough commitment power, particularly in the form of supervisory ruling

Conclusion

- The welfare effect of increased private information in the hands of the private sector is non-monotonic!
- The bank's cost in case of supervisory objection ought to be set high.
  - intuitively, what the bank cares about is the "cost-adjusted" probability that the supervisor will object to its aggressive risk-taking
  - some criticism around stress testing is that capping dividend and suspending share repurchases are too severe as disciplinary measures
  - however, it is not at discretion of the supervisor to reset this cost from one period to the next
  - we show that, if the cost of rejection is too low, it can hamstrung the supervisor
  - we also show that, if the cost of rejection is too high, the supervisor can always achieve an unconstrained optimum by introducing strategic ambiguity into his communication strategy