Economic Quarterly—Volume 94, Number 1—Winter 2008—Pages 1–16

Should Bank Supervisors Disclose Information About Their Banks?

Edward Simpson Prescott

In order to preserve the safety and soundness of the banking system, bank supervisors collect a great deal of information about a bank. They examine its balance sheet, its operations, and its management. They observe the reports made by a bank's internal management reporting system. To gather this information, they have legal and regulatory powers that are not available to others. Collecting this information is expensive. In the United States, federal and state regulators spent nearly 3 billion dollars in 2005 and banks spent substantially more complying with bank regulations.[1]

Anyone who trades with a bank or buys one of its securities would find this information valuable. Indeed, anyone who even thinks about trading or buying a bank security would find this information valuable. But right now, people cannot view this information because the bank's supervisor does not disclose it.[2] Furthermore, once the supervisor collects the information and forms his assessment, a bank is not allowed to disclose the assessment without the supervisor's approval.

Given that this information is expensive to collect and the market would like to use it, why not require a supervisor to disclose it, or at least allow a bank to voluntarily disclose it? This would make it easier for potential investors to evaluate a bank and would avoid expensive duplication of information collection and analysis. Indeed, some have argued for precisely this course of action (Shadow Financial Regulatory Committee [1996]).

The author would like to thank Kevin Bryan, Mark Vaughan, John Weinberg, and Alex Wolman for their helpful comments. The views expressed in this article do not necessarily reflect those of the Federal Reserve Bank of Richmond or the Federal Reserve System. E-mail: Edward.Prescott@rich.frb.org.

[1] Author's calculations.

[2] Actually, some information is disclosed by bank supervisors, but only a subset of information and only in certain circumstances.

In this article, we argue that the above logic is incorrect, or more accurately, seriously incomplete in that it ignores an important cost to disclosure. Namely, supervisory disclosure would make it harder for the supervisor to collect the information in the first place. This argument is developed with a model that explicitly takes into account the incentives of a bank to accurately report information to the supervisor and the effect of broad disclosure on these incentives. The ability of a bank supervisor to accurately gauge the quality of a bank and the incentives of a bank to keep that information quiet in bad times are fundamental problems in bank regulation. This is why understanding how disclosure impacts the ability of supervisors to collect the information is necessary for evaluating proposals that require a supervisor to disclose it.

The bigger issue here is just how much disclosure should there be. Information in security markets is similar to a public good in that it is useful to everyone and if one person uses it, it does not reduce someone else's use of it. Nevertheless, this article argues that public dissemination of information can hurt the ability to collect it in the first place. The analysis demonstrates that it matters who receives the information and for what purpose.

1. BANK SUPERVISION

The purpose of bank supervision is to keep banks safe and sound. It protects taxpayers from liabilities resulting from deposit insurance and helps preserve financial stability.

Bank supervisors use a variety of tools to meet these objectives. The most important is direct examination.[3] All U.S. banks are examined periodically. According to federal law, all banks must have a formal, on-site exam conducted at least once every year, though under certain conditions banks with less than 250 million dollars in assets can be examined once every 18 months. The periodic on-site examinations are not the only source of direct supervision. Bank supervisors also monitor banks between exams by analyzing a variety of data. (In supervision, this is often called off-site surveillance.) This information can be used to determine if a targeted on-site exam is necessary. Furthermore, for a large bank, supervisors' offices are located at the bank throughout the year, which enables supervisors to generate a constant flow of information.

An exam is broad in its scope, but is based on the CAMELS system.[4] This system includes assessments of each of the following components of a bank:

[3] Information on supervision is from Spong (2000).

[4] The CAMELS system is used for examinations of commercial banks. There are similar systems used to assess a bank holding company. Note that combining the first letter of each component creates the acronym, CAMELS (see top of next page).

- Capital Adequacy
- Asset Quality
- Management and Administrative Ability
- Earnings Level and Quality
- Liquidity Level
- Sensitivity to Market Risk

Each component is assigned a rating of one to five, with one being the best and five being the worst. The components are then combined to create an overall CAMELS rating. The overall rating uses the same scale as the components. The exam report also contains more detailed assessments and comments about the condition of the bank. Finally, the exam report is confidential and cannot be disclosed by the bank without the permission of the supervisor.

Needless to say, this is precisely the sort of information that would be of considerable value to any potential investor or bank counterparty. Indeed, despite the threat of legal sanctions, bank counterparties have, at times, used the ratings in bank transactions. For example, Supervision and Regulation Letter 02-14 (2002) stated that during the time discussed in the letter, supervisors noticed that CAMELS ratings were being included in covenants for securitization transactions, which was forbidden without supervisory approval.[5]

There is also a great deal of statistical evidence that supervisory ratings contain useful information. There is substantial literature in banking that assesses the correlation between bank exam ratings and market prices of bank securities. This literature finds that bank exams predict changes in market prices, though this information tends to decay within approximately six months after an exam, e.g., Berger and Davies (1998).[6]

2. A SIMPLE MODEL

In our model, there is the bank, the bank supervisor, and investors. The bank randomly produces a gross return of either 0 or 1. The probability of producing the high return is $\theta \in (0, 1)$, which is a random shock. There is a finite number

[5] A supervisory letter is a letter written by the Federal Reserve Board Division of Bank Supervision and Regulation concerning policy and procedural matters related to Federal Reserve supervision of banks. It is used to disseminate information to the banks and to regional Federal Reserve supervision staff.

[6] Actually, because this literature is interested in the value of market data for supervisory purposes, most of it asks the opposite question of whether market prices can predict changes in exam ratings. They do, particularly when a substantial amount of time has passed since the last exam. See Flannery (1998) for a survey.

of possible shocks and the probability of a shock is $h(\theta)$. In our model, only the bank observes the shock.

After observing the shock, the bank raises capital from investors to finance its investment. Investors do not observe the shock. They require an expected rate of return of $\bar{r}$ and for this reason would like to know the shock. We assume that the bank has limited liability, so the market only receives a payment if the bank produces a high return. Even though the market does not observe the shock, the payment may depend on the shock to the extent that the market learns the value of θ. The interest rate is $r(\theta)$, which is also the payment in the high return state. For simplicity, we ignore deposits and treat all invested funds as uninsured debt. The absence of deposits is not necessary for our results; we could have assumed that the bank's gross return is a quantity in excess of its deposit liabilities and the results still would not change.

Like the investors, the bank supervisor does not observe the shock. He wants to know the shock to better target his supervisory resources. We do not explicitly model what the supervisor does, but instead assume that the supervisor takes an action a and gets utility $W(a, \theta)$. The utility function is such that if the supervisor knows the shock, $a(\theta)$ is decreasing. The idea is that the higher the probability of success, the less involved with the bank the supervisor needs to be.

To illustrate the best case for a supervisor not disclosing the bank's information, we assume that the bank does not care about the supervisor's action, but instead only cares about its expected profits. This is an extreme assumption. In practice, supervisors can take actions that will hurt a bank's profits, so banks do care what they report to the supervisors. At this point, however, we want to illustrate the simplest case for supervisors not disclosing information. Later, we will relax this assumption.

As we said earlier, we assume that the bank has limited liability, so that if it produces 0, there is nothing to pay the investors. Given shock θ, expected profits for the bank are $\theta(1 - r(\theta))$. For this reason, the bank prefers a low value of r. Finally, we assume that the bank would not even operate unless its expected return equals a reservation level of profits $\bar{U}$.

Reporting

The key element in determining the effect of supervisory disclosure is the bank's incentive to share information. We have assumed that the bank is the only entity that observes the shock θ. After observing the shock, the bank sends a costless, unverifiable report on it. By unverifiable, we mean that the bank's report need not be the same as the true value of the shock and there is no way to check its veracity.[7]

We will consider two different reporting models. In the first model, the bank sends a report to the supervisor who shares it with the market. In the second model, the bank sends separate reports to the market and to the supervisor and the supervisor does not share his information. Furthermore, the market does not observe the supervisor's action; otherwise, the market might be able to infer the report. The second model most closely resembles current practice.

To summarize, the timing of the problem is as follows:

1. The bank observes θ. The supervisor and the market do not observe it.
2. The bank reports the following:
   - **Model 1–**A single report is sent to the supervisor who then shares it with the market.
   - **Model 2–**Separate reports are sent to the market and to the supervisor and the supervisor does not share his information.
3. The market sets its interest rate and the supervisor takes his action.

In both models, the bank's report may influence the payment demanded by the market and the supervisor's action. We will call the pair of functions $r(\theta)$ and $a(\theta)$ an *allocation*. Determining possible allocations is not straightforward because there is a wide variety of messages the bank can send. Fortunately, we can simplify the analysis considerably by using the revelation principle.[8] In our model, this principle states that allocations that are consistent with the bank's private information can be determined by only considering the class of direct mechanisms. A direct mechanism is one in which the reporting space consists only of the values of the shock, θ, and the allocation satisfies incentive constraints that guarantee that the bank reports truthfully.

The revelation principle is an extremely useful device for determining whether an allocation is consistent with incentives. Because the bank truthfully reports its shock in a direct mechanism, there may be confusion about what is meant in this article by sharing information. In a direct mechanism, it is true that the receiver of the information learns the true value of the shock. However, the only reason the receiver learns the value is because there are incentive constraints that the allocation must satisfy. As we will see in the first

[7] In practice, the examination process puts limits on what banks can report. Later, we will extend the model to capture some of these features.

[8] The revelation principle was developed in Harris and Townsend (1981) and Myerson (1979). For a textbook treatment, see Myerson (1997).

model, incentive constraints can be very restrictive. For example, they might not allow the interest rate to depend on the shock in any way. In this case, the receiver learns the value of the shock, but this is only because he is not going to do anything with this information. Rather than saying how little information is transmitted, the constraints determine to what extent we, in the model, limit the dependence of the interest rate on the shock so that the truth is reported. In the two models we are comparing, the different reporting assumptions lead to different incentive constraints and different sets of feasible allocations.

Model 1: The Supervisor Shares His Information

Proposals for information sharing require the supervisor to disclose the information he receives from the market. Under such a disclosure rule, the bank would know that any information it shared with the supervisor would also be seen by the market. For this reason, we model this proposal by allowing the bank to send a single report that is seen by both the supervisor and the market.

Let $r(\theta)$ be the interest rate if the bank reports θ. The report need not be the true value of the shock. The bank will select the report that maximizes its utility, that is, given $r(\theta)$ and the shock θ, the report has to solve

$$\max_{\tilde{\theta}} \theta(1 - r(\tilde{\theta})). \tag{1}$$

According to the revelation principle, we can consider interest rate schedule $r(\theta)$ if a solution to (1) is for the bank to truthfully report θ. This is called an incentive compatibility constraint. An alternative representation of it is

$$\theta(1 - r(\theta)) \geq \theta(1 - r(\theta')), \quad \forall \theta, \theta'. \tag{2}$$

The left-hand side of (2) is the profits received by the bank if the shock is θ and it tells the truth. The right-hand side is the bank's profits if it, instead, reports θ'.

This constraint strongly restricts the form that $r(\theta)$ can take. For example, if $r(\theta)$ decreases with θ then the bank can always report a higher value of θ and receive the benefit of a lower interest rate. Indeed, the only function $r(\theta)$ that satisfies this constraint is $r(\theta) = r$, that is, the interest rate cannot depend on θ. In this case, it does not hurt the bank to lie, so it might as well claim that its probability of success is as high as possible. Of course, the market recognizes this incentive, so it completely discounts the bank's report and just demands a fixed interest rate.

The market demands a return of $\bar{r}$. We assume that the market cannot commit to not using the information it receives from the supervisor. Because of this limited commitment, if the information disclosed by the supervisor perfectly informed the market what the shock was, the constraint would take the form $\theta r(\theta) = \bar{r}$. Alternatively, if the supervisor had no information to disclose about the shock—as is the case here—the constraint would take the form $\sum_\theta h(\theta)\theta r = \bar{r}$.

In this simple example, it is very easy to determine what these reporting incentives mean for the supervisor's action. The bank's report contains no useful information, so the supervisor might as well ignore it and choose his action as if he knew nothing about the bank, other than the distribution of its probability of success $h(\theta)$. This means that he will have to choose a constant action, that is, $a(\theta) = a$.[9]

There are a variety of (r, a) pairs that could occur in equilibrium. We use a constrained maximization problem to pick a particular pair. This pair will maximize the supervisor's utility subject to the bank receiving a minimal guaranteed level of utility and the market receiving its required return.

The constrained maximization problem is

$$\max_{r,a} \sum_\theta h(\theta) W(a, \theta),$$

subject to a constraint that the bank receives its participation utility

$$\sum_\theta h(\theta)\theta(1 - r) \geq \bar{U}, \tag{3}$$

and the constraint that the market receives its required return

$$\forall\theta, \quad \sum_\theta h(\theta)\theta r = \bar{r}. \tag{4}$$

Any connection between $a(\theta)$ and $r(\theta)$ was already incorporated through the incentive constraints, which restricted these functions to take on constant values. The resulting maximization problem is extremely simple in that r is determined solely by the constraints and a is determined solely by maximizing the objective function.

Model 2: The Supervisor Does Not Share His Information

In this section, we allow the bank to send a separate report to the supervisor, one that the market does not see. This assumption resembles existing practices. Supervisors engage in a great deal of direct communication with banks and, except in certain extreme cases, these communications are not shared with the public.

Allowing two separate reports only affects the analysis of the supervisor's action. For the report to the market, the incentives are exactly as before. When

[9] If we did not make the limited commitment assumption on the market, then the market (or more likely another bank) could offer a line of credit with fixed interest rate r that was not contingent on the supervisor's disclosure. It would then be incentive compatible for the bank to report truthfully to the supervisor and the report would then vary with θ.

the bank sends its report to the market, its incentive is to claim that the bank is as profitable as possible. This means that, as before, $r(\theta) = r$. Now, however, the report to the supervisor can say something different. When the bank sends its message to the supervisor, it considers what effect this will have on itself. In the simple example here, the supervisor's action has no effect on the bank. The bank does not gain anything by lying to the supervisor, so it might as well tell him the truth.[10]

What this means for the problem is that now the supervisor's decisions can be made to explicitly depend on θ, that is, $a(\theta)$ need not be a constant. This means that the maximization problem is now

$$\max_{r,a(\theta)} \sum_{\theta} h(\theta) W(a(\theta), \theta)$$

subject to the participation constraints (3) and (4), which are the same as before.

The only difference between the two programs is that now the supervisor can make his action depend on the shock, which is much better for the supervisor. The bank is willing to share information with the supervisor because the information is not then passed along to the market. The market's knowledge of θ would have a larger impact on the bank than the supervisor's knowledge of θ. For this reason, the bank is less willing to share this information with the market.

This example starkly illustrates the potential cost of supervisory disclosure. It gives the bank an incentive to keep information hidden, not because it cares if the supervisor receives the information, but because it cares if the market receives it.

In practice, there are plenty of situations in which the supervisor will use his information about the bank to take actions that the bank does not want. It is in these situations that the analysis will become more complicated. We will discuss this situation later when we more explicitly model the verification technology and the punishments available for lying. Before discussing these issues, there is a related and interesting analysis of actions that the bank can take that is worth examining because it ties into regulatory rules that forbid banks from disclosing their CAMELS rating.

[10] The bank does not gain anything here by telling the truth either. In these models, it is customary to assume that if the information sender is indifferent between two options, he does what is best for the principal (here the supervisor). Furthermore, in most of these models an arbitrarily small change in the contract will make the agent strictly prefer to tell the truth, while only marginally impacting the principal. That is not the case here, though it is in the extension discussed later.

3. WHY FORBID A BANK FROM DISCLOSING ITS EXAMINATION RESULTS?

Regulatory policy explicitly forbids a bank from disclosing its examination results. As we mentioned earlier, the market values this information and has tried to use it in the past. This raises the question that since banks sometimes want to disclose their examination results, why not let them? If it is a voluntary decision, why could not they release it only when it helps the market?

In this section, we demonstrate that allowing a bank to voluntarily disclose its examination result is, in equilibrium, exactly the same as requiring the supervisor to share his information with the public. For this reason, the supervisor must forbid banks from disclosing the results if he does not want to disclose the information to the market.

When the supervisor examines a bank, the exam's results are put into writing and delivered to the bank. This introduces an important departure from the previous analysis about communication. If the bank discloses the examination results to the market, it can also provide the market with verifiable evidence in the form of a copy of the formal report.

Now consider the following variation of the last model. As before, the bank sends its report to the supervisor, but now the supervisor sends back a document stating what the bank reported to the supervisor. We assume that this report is a legal document that cannot be falsified by the bank. After receiving the supervisory document with the CAMELS rating, the bank has the choice of whether to disclose the report to the market. Because the document is non-falsifiable, if the market sees it, then the market knows that this information was reported to the supervisor. Finally, this disclosure is considered to occur before the market sets the interest rate and the supervisor takes his action.

This seemingly small modification to the last model is actually a significant change. Being able to disclose information in a verifiable manner will affect incentives and the quality of the information transmitted. The literature on this type of communication is referred to as the disclosure literature. Early articles in this literature include Grossman (1981) and Milgrom (1981). For a survey of this literature, see Fishman and Hagerty (1998).

To see the effects of allowing banks to disclose this verifiable information, we will start with the allocation that was the solution to the second model, the case in which the supervisor does not disclose his information. That allocation included an interest rate function $r(\theta) = r$ that did not depend on θ and a supervisory action $a(\theta)$ that could depend on θ. The report to the market contained no information because the bank had the ability to lie. For simplicity, we will assume that the only relevant information in the examination is the CAMELS rating, which we will assume corresponds exactly to the shock θ.

Now, consider the stage at which the bank can disclose its report to the supervisor if it chooses. Imagine that the bank received the highest profitability shock θ_1, which corresponds to receiving the best CAMELS rating of 1. The bank has a choice that it did not have in the earlier analysis. Instead of sending an unverifiable report to the market, it can display the supervisory document with the CAMELS rating. This kind of communication is very different than the unverifiable kind analyzed previously. In particular, if the bank displays its report, the market knows the shock is θ_1 (we are assuming for the moment that the bank reported truthfully to the supervisor) and will set $r(\theta_1)$ to satisfy

$$r(\theta_1) = \bar{r}/\theta_1. \tag{5}$$

Compare this interest rate with that of our starting allocation. In that allocation, $r = \bar{r}/(\sum_\theta h(\theta)\theta)$, and because $\theta_1 > \sum_\theta h(\theta)\theta$, we have $r(\theta_1) < r$. Therefore, if the bank receives shock θ_1, it will prefer to display this value rather than sending the unverifiable report to the market.[11]

Therefore, it is only possible for the bank to send the unverifiable report for lower values of θ. But the market would realize this and demand an interest rate of r based on θ being less than θ_1 (remember θ_1 is the *high* value of θ here), that is, r satisfies

$$r = \frac{\bar{r}}{\sum_{\theta<\theta_1} h(\theta)\theta}. \tag{6}$$

Under this interest rate, θ_2 is the highest value for which the bank does not disclose. As before, however, if the bank received shock θ_2, it would want to disclose this to the market and receive interest rate

$$r(\theta_2) = \bar{r}/\theta_2,$$

because it is less than the interest rate calculated in (6).

We can repeat this analysis, inductively, for each value of θ and end up with the result that the bank discloses for all values of θ.[12] The market is thinking that if the bank is not willing to disclose its rating, then it must be a risky bank. Consequently, a bank that receives anything other than the worst shock feels obligated to prove that it did not receive the worst shock.

This analysis demonstrates that if the bank reports truthfully to the supervisor and the supervisor gives the bank its CAMELS rating, then in equilibrium, the bank will be forced to disclose its rating. Now, we need to take a step back and ask what these disclosure incentives mean for the bank's incentives to report to the supervisor. Because the bank prefers the low interest rate, it will always tell the supervisor that its shock is θ_1, so that once it receives the CAMELS rating from the supervisor, it can display it to the market. Of course, the market understands the bank's incentive to lie to the supervisor (as does the supervisor), so the market ignores the displayed CAMELS rating and we are back in the public reporting model. Therefore, with voluntary disclosure,

[11] We are assuming that the bank cannot commit to not disclosing.

[12] Actually, for the lowest value of θ, the bank is indifferent to disclosing or sending an unverifiable report, but that is not important here.

no information is transmitted and both the interest rate and the supervisor's action do not depend on θ in any way.

The analysis demonstrates that the pressure for a bank with the good shock to disclose its CAMELS rating is very strong. This is probably why this information gradually started appearing in bank contracts, despite the rules against such actions, and why a strong reminder was necessary for why they should not be disclosed. A very good question is whether this information still makes its way around the market more informally.

4. A MORE COMPLEX MODEL

As we described earlier, a CAMELS rating is a summary number generated by supervisors. We modeled this information as coming from an unverifiable report sent by the bank to the supervisor. In practice, a CAMELS rating is not only based on information reported by the bank to supervisors, but is also based on the assessment the supervisor makes from detailed examination of the bank. Furthermore, this information is costly to collect. One of the arguments for supervisory disclosure is that information collection and assessment is costly so why duplicate this effort?[13] Furthermore, supervisors have special legal powers that allow them to gather information more cheaply than markets. For these two reasons, it would be efficient for supervisors to collect and then share the information.

In this section, we describe an extension of the model that will allow us to better discuss this additional tradeoff. The extension has two additional features. The first is that we give the bank a distaste for supervisory actions. The second is that we provide the supervisor with a technology for verifying the information he receives from the bank. This technology is costly to operate. It is our interpretation of the examination process.

The supervisory technology is an audit that detects a lie a positive fraction of the time. If a lie is detected, then the supervisor can impose a penalty. We assume that the audit never generates a false positive, that is, it never concludes that the bank lied when it actually did not. It can only detect a lie if one was actually made. For simplicity, we also assume that the detection probability does not depend on the shock. In both theory and practice, it would be desirable to allow the supervisor to vary the audit intensity with the report. Still, several components of a supervisory auditing system are fixed and planned far in advance, and this assumption captures these features fairly well.

---

[13] This problem also arises in private markets for information. Two prominent examples are financial accounting information and rating agencies' ratings. To avoid free riding from information sharing, the evaluated firm pays the accountants and the ratings agencies for a self-evaluation rather than each potential investor paying for his own evaluation. While this solves the free-riding problem, it can create some rather severe incentive problems.

The other feature we add to this model is to let the bank care about which action the supervisor takes. We will model supervisory action a as imposing a pecuniary cost to the bank of a. Furthermore, to simplify the analysis, we assume that if the supervisor detects a lie, the supervisor responds to the actual results in the same manner that he would have if the bank had reported truthfully. Thus, the supervisor will not use the action as an additional punishment for lying.

Finally, we consider the problem of implementing supervisory action schedules in which $a(\theta)$ is decreasing in θ. The advantage of this approach is that we do not have to solve the programming problem. Furthermore, this class of supervisory actions is intuitively appealing and characteristic of solutions to many parameterizations of the problem.

We start with the case in which the information is not shared (Model 2). In this setup, the bank sends separate reports to the supervisor and to the market and then the supervisor audits the quality of his report with probability π. If there is an audit and it detects a lie, the supervisor imposes a penalty P.

As in the earlier analysis, the interest rate is a constant r, but now because of the audit, the supervisor's action can depend on the shock. The limitation of this dependence is described by the incentive constraint on the supervisor's report. It is

$$-a(\theta) \geq -(1-\pi)a(\theta') - \pi(a(\theta) + P), \quad \forall \theta, \theta'. \tag{7}$$

The interest rate r is not in (7) because it drops out of both sides of the equation. The left-hand side of (7) is the utility the bank receives from the supervisory action if it reports truthfully. It is negative because it is a cost imposed on the bank. The right-hand side is the utility from lying. The first term is the probability of not being audited times the utility from receiving the $a(\theta')$ supervisory action. The second term is the probability of being audited times the utility from being caught lying. The utility from lying consists of the utility from the supervisor taking the action he is supposed to take, $a(\theta)$, plus the utility cost of the penalty.

Let $a(\theta_l)$ be the supervisory action taken by a bank receiving the lowest shock, θ_l. The only binding incentive constraints will be those for shocks in which a bank that is supposed to take $a(\theta_l)$ claims to have received the highest shock, that is, θ_h. The intuition for this result is that if the bank is going to lie, it might as well report θ_h and receive the least onerous supervisory action $(1-\pi)$ of the time. Furthermore, if the incentive constraints prevent these banks from claiming to have received the θ_h shock, then the constraints will prevent any bank with a higher value of θ from claiming to be θ_h, as well.

From the arguments above and letting $\Delta A = a(\theta_l) - a(\theta_h)$, we can simplify (7) to the single incentive constraint.[14]

$$(1-\pi)\Delta A \leq \pi P. \tag{8}$$

The left-hand side is the gain from lying and the right-hand side is the cost. To guarantee truth telling, the former cannot exceed the latter.

We now examine how much stronger the penalty would have to be to implement the same regulatory schedule, $a(\theta)$, along with an interest rate schedule $r(\theta)$ that varies with θ, as would happen if the supervisor shares his information. Let this new penalty be $\tilde{P}$. In this case, we assume that if the supervisor conducts an audit and finds that the bank lied, then he shares the correct number with the market.

The incentive constraint for the bank is

$$\begin{aligned} &\theta(1-r(\theta)) - a(\theta) \\ &\quad \geq \quad (1-\pi)\left[\theta(1-r(\theta')) - a(\theta')\right] \\ &\qquad\quad +\pi\left[\theta(1-r(\theta)) - a(\theta) - \tilde{P}\right], \quad \forall\theta,\theta'. \end{aligned} \tag{9}$$

The left-hand side of (9) is the profit the bank receives if the shock is θ and it reports truthfully. Notice that no penalty is ever imposed if the bank takes this strategy. The right-hand side of (9) calculates the profit the bank receives if the shock is θ and it lies by reporting θ'. The term that starts with $(1-\pi)$ is the probability of the bank not being audited times the profit it gets in that event. If it is not audited, it only pays $r(\theta')$ and does not pay the penalty. The term that starts with π is the probability of the bank being audited times its profit. When the bank is audited, the supervisor finds out the true shock, takes the action $a(\theta)$, and imposes the penalty P. Furthermore, the market charges $r(\theta)$.

It is convenient to rearrange the terms in (9) to obtain

$$(1-\pi)\left[a(\theta) - a(\theta') + \theta(r(\theta) - r(\theta'))\right] \leq \pi\tilde{P}, \quad \forall\theta,\theta'. \tag{10}$$

As in the no-information sharing case, (10), can be considerably simplified. Remember, we are interested in implementing the same $a(\theta)$ contract, which was decreasing. This means that $r(\theta)$ is also decreasing because the market is using the same information that the supervisor uses to distinguish between shocks. Therefore, given θ, the left-hand side of (10) is maximized for $\theta' = \theta_h$, which means the only binding incentive constraint is the one in which the bank claims to be the lowest risk possible, θ_h.

The constraints can be further simplified by recognizing that if the single incentive constraint for a θ_l bank holds, then the single incentive constraint

---

[14] Technically, there need not be a single incentive constraint. As discussed above, higher values of the shock could also be assigned $a(\theta_l)$, but these incentive constraints will look the same.

for all other θ holds as well. To see this, first note that θ_l maximizes $a(\theta)$. Second, note that $\theta(r(\theta) - r(\theta_h)) = \bar{r} - \frac{\theta}{\theta_h}\bar{r}$. This term is also maximized at θ_l. Therefore, the highest value of the left-hand side of (10) occurs for the incentive constraint on θ_l. For these reasons, if we let $\Delta R = r(\theta_l) - r(\theta_h)$, we can represent the incentive constraints with the single incentive constraint

$$(1 - \pi)\,(\theta_l \Delta R + \Delta A) \leq \pi \tilde{P}. \tag{11}$$

The interpretation of (11) is very intuitive. The right-hand side is the expected loss from lying. The left-hand side is the expected gain. It has two components: the gain from lower expected interest payments and the gain from weaker supervisory actions. The inequality ensures that the penalty from lying exceeds the gain.

As in the earlier analysis, when compared with (8), incentive constraint (11) also takes into account the gain from lower interest payments to the market. This requires the penalty to be higher in the amount of $(1 - \pi)\theta_l \Delta R$. Alternatively, the supervisor could change the audit probability in order to implement the desired allocation. In either case, supervisory disclosure makes it harder for the supervisor to receive the information he wants.[15]

In this model, as in the earlier section, supervisory disclosure is unambiguously bad because it impedes the supervisor's collection of information. Ex ante, there is no value to the market and to the bank from disclosure because with both being risk-neutral, the shock only affects in what form they receive their payoffs. However, if the model was extended to include a feature in which it matters how r varies with the shock, for example, if the size of the investment was endogenous, then supervisory disclosure could have some real benefits. Furthermore, the case for information sharing would be unassailable if auditing was costless because then the supervisor would learn the shock at no cost and could share his information with the market without hurting his ability to gather the information. Of course, the empirically relevant case is somewhere between these two extremes.

5. DISCUSSION

The analysis above was designed to describe the tradeoffs to supervisory disclosure of information. The main argument is that if supervisors need the cooperation of a bank to receive information, then disclosure will increase the cost of cooperation to the bank . This increased cost either reduces the quality of information the supervisor receives or it requires the supervisor to spend more of his resources collecting the information.

---

[15] More generally, the supervisory action that the supervisor will try to implement would change as well.

One interesting feature of the supervisory process that was omitted from the analysis is the nature of the information collected. In the model, information was represented as a single dimensional variable that summarized all of the relevant information for the market and the supervisor. Furthermore, the supervisor's examination technology detected inaccurate information, independent of the absolute value of the information. In practice, the ability of supervisors to detect inaccurate information should depend on the absolute value of the difference between the true value and the reported value.

A second important difference is the nature of the information the examination process is designed to capture. Bank supervisors are mainly concerned with risks to the safety net, that is, situations in which a bank could become insolvent. Supervisors care much less whether a bank is going to have average, good, or excellent profits. Markets care a great deal, however, about these distinctions. For this reason alone, markets will continue to monitor a bank whether supervisors disclose or not.

A third important feature that was not in the model is the incentives of the supervisors. In fact, supervisory forbearance worsened the Savings and Loans crises of the 1980s. One argument for disclosure is that the public release of this information may force the supervisor to act early, thus reducing the size of the deposit insurers' liability.[16] The prompt and corrective action provisions of the Federal Deposit Insurance Corporation Act of 1991 have this flavor. Supervisors have to take certain actions, some of which are publicly disclosed, if the amount of bank capital levels falls below certain levels. Of course, any analysis along this line of thought must take into account the incentives of supervisors to accurately disclose information. This suggests a need to audit the supervisor after a bank failure, though such an audit would never identify cases of successful supervisory forbearance.[17]

Finally, it should be pointed out that information sharing can go the other direction as well. A variety of proposals recommend that bank supervisors gather information conveyed from market prices. Furthermore, regulatory practice already uses the information generated by rating agencies' ratings of securities. Capital requirements for some bank holdings of securities are tied to these ratings. Reversing our model to examine the incentives for the ratings agency to accurately rate these securities suggests that regulatory use of the ratings increases the incentive for banks to encourage these agencies to inflate the ratings. In this case, the general principle at stake is that the diffusion of information can negatively affect the ability to collect it.

---

[16] For an argument along this line, see Shadow Financial Regulatory Committee (1996).

[17] When the Federal Deposit Insurance Corporation loses an amount equal to the greater of $25 million or 2 percent of a bank's assets from a bank's failure, the inspector general of the failed bank's federal supervisor is required to prepare a public report on the failure (Walter 2004).

REFERENCES

Berger, Allen N., and Sally M. Davies. 1998. “The Information Content of Bank Examinations.” *Journal of Financial Services Research* 14 (2): 117–44.

Fishman, Michael, and Kathleen Hagerty. 1998. “Mandatory Disclosure.” In *The New Palgrave Dictionary of Economics and the Law,* ed. Peter Newman, London: Macmillan Press.

Flannery, Mark. 1998. “Using Market Information in Prudential Bank Supervision: A Review of the U.S. Empirical Evidence.” *Journal of Money, Credit and Banking* 30 (3): 273–305.

Grossman, Sanford J. 1981. “The Informational Role of Warranties and Private Disclosure about Product Quality.” *Journal of Law and Economics* 24 (3): 461–84.

Harris, Milton, and Robert M. Townsend. 1981. “Resource Allocation under Asymmetric Information.” *Econometrica* 49 (1): 33–64.

Milgrom, Paul R. 1981. “Good News and Bad News: Representation Theorems and Applications.” *Bell Journal of Economics* 12 (2): 380–91.

Myerson, Roger B. 1979. “Incentive Compatibility and the Bargaining Problem.” *Econometrica* 47 (1): 61–73.

Myerson, Roger B. 1997. *Game Theory: Analysis of Conflict.* Cambridge, MA: Harvard University Press.

Shadow Financial Regulatory Committee. 1996. “Statement 132: Disclosure of Examination Reports and Ratings.” *Journal of Financial Services Research* 10 (4): 302–03.

Spong, Kenneth. 2000. *Banking Regulation: Its Purposes, Implementation, and Effects.* Federal Reserve Bank of Kansas City, 5th edition, Division of Supervision and Risk Management.

Supervision and Regulation Letter 02-14. 2002. ”Covenants in Securitization Documents Linked to Supervisory Action or Thresholds.” Board of Governors’ Division of Banking Supervision and Regulation.

Walter, John R. 2004. “Closing Troubled Banks: How the Process Works.” Federal Reserve Bank of Richmond *Economic Quarterly* 90 (1): 51–68.