Federal Reserve Bank of New York
Staff Reports

The Economics of Bank Supervision

Thomas M. Eisenbach
David O. Lucca
Robert M. Townsend

Staff Report No. 769
March 2016
Revised January 2017

The Economics of Bank Supervision
Thomas M. Eisenbach, David O. Lucca, and Robert M. Townsend
Federal Reserve Bank of New York Staff Reports, no. 769
March 2016; revised January 2017
JEL classification: D82, G21, G28

Abstract

We study bank supervision by combining a theoretical model of asymmetric information and a novel dataset on work hours of Federal Reserve supervisors. We highlight the trade-offs between the benefits and costs of supervision and use the model to interpret the relationship between supervisory efforts and bank characteristics observed in the data. More supervisory hours are spent on larger, more complex, and riskier banks. However, hours increase less than proportionally with bank size, suggesting technological economies of scale in supervision. We provide evidence of constraints on supervisory resources, documenting reallocation of hours at times of stress and in the post-2008 period. Using variation implied by this resource reallocation, we find evidence that supervision lowers risk.

Key words: bank supervision, bank regulation, monitoring, time use

---

Eisenbach, Lucca: Federal Reserve Bank of New York (e-mails: thomas.eisenbach@ny.frb.org, david.lucca@ny.frb.org). Townsend: Massachusetts Institute of Technology (e-mail: rtownsend@mit.edu). For valuable comments, the authors thank Doug Diamond, Mark Flannery, Xavier Freixas, and Gyongyi Loranth, as well as audience members at the New York Fed Conference on Supervising Large and Complex Financial Institutions, the Wharton Conference on Liquidity and Financial Crises, the BAFFI CAREFIN Annual International Banking Conference, Hong Kong University, Hong Kong University of Science and Technology, Singapore Management University, Nanyang Technological University, and National University of Singapore. They also thank Emily Eisner, Eilidh Geddes, and Alice Liang for excellent research assistance. Any errors are the authors' own. The views expressed in the paper are those of the authors and do not necessarily reflect the positions of the Federal Reserve Bank of New York or the Federal Reserve System.

1 Introduction

Much debate in recent years has focused on bank supervision and regulation as a result of the significant changes in policies aimed at the banking sector post-financial crisis. The 2010 Dodd Frank Act (DFA) introduced a number of provisions tightening regulations for the largest and most complex banking organizations.[1] While harder to measure, bank supervision has also undergone a significant expansion.

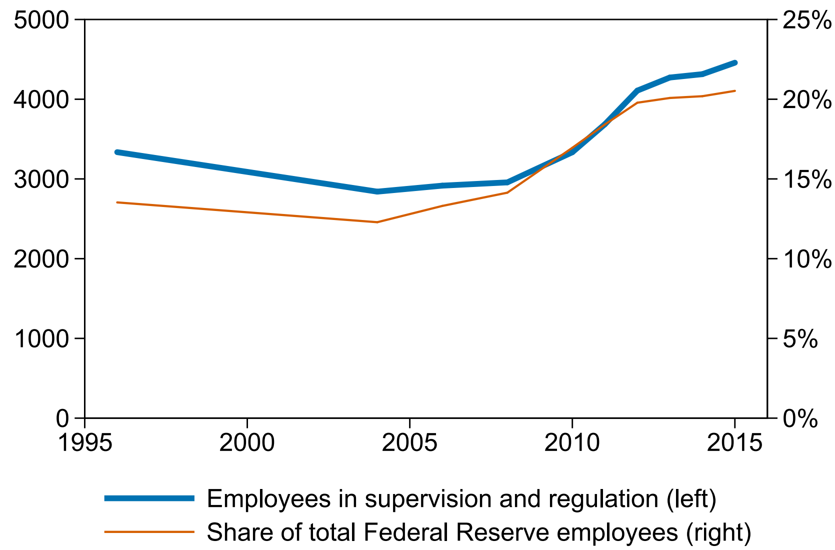

Figure 1: Federal Reserve supervisory staff. (Source: FRB Annual Reports)

As shown in Figure 1, Federal Reserve supervisory staff, which had shrunk in the decade preceding the financial crisis, has risen about 50% since then to about 4,500 employees in 2015, and now accounts for about 20% of all Fed employees. Academic work on bank supervision remains limited, partly reflecting the opacity of supervisory activities that stems from supervisors' reliance on confidential information. Limited work and supervisory opacity make it hard to assess the post-financial crisis supervisory expansion.

This paper first develops an abstract framework founded on institutional details, which helps crystallize what bank supervision is and how it differs from regulation.

[1] The DFA contains a large number of regulatory changes both in the context of financial stability and consumer protection. In terms of financial stability, the DFA established the Financial Stability Oversight Committee (FSOC) with the authority to designate systemically important institutions. The DFA also mandated the Federal Reserve Board and other federal banking regulatory authorities to issue tighter regulations for the largest institutions related to the level of regulatory capital, proprietary trading and stress testing, among others.

Second, we use a unique dataset of hours spent by Federal Reserve bank examiners overseeing individual U.S. bank holding companies (BHCs) to illustrate details of the supervisory process, quantify key empirical relations, and study effects of bank supervision both pre- and post-financial crisis.

Supervision is closely related to, but distinct from, regulation of banking organizations. Regulation involves the development and promulgation of the rules under which banking organizations operate, as well as their enforcement in the court of law. Supervision is closely related to regulation to the extent that it is often entrusted with compliance with regulation. But a key feature of supervision is ensuring that banks don't engage in "unsafe and unsound" practices. "Safety and soundness" is not hard-coded into law, reaches far beyond written rules, and crucially involves judgment in assessing whether a bank may be engaging in excessive risk. This flexibility paired with reliance on confidential information makes supervision opaque and hard to assess for the public and researchers. In practice, supervisory activities involve monitoring banks and using this information to request corrective actions from banks should their conditions or practices be deemed unsafe or unsound.

In our framework, the need for regulation and supervision results from a bank's limited liability and from macro-prudential spillovers of bank failures. In differentiating supervision from regulation, we assume that regulation is "coarse" and can only be contingent on verifiable information, whereas supervision is "discretionary" and can be contingent on non-verifiable information.[2]

We assume that supervision can collect information about a bank's actions through costly monitoring, leading to an interim signal, which in the data we proxy with the bank's supervisory rating. After observing the signal, we assume that supervision can respond and influence asset payoffs before they are realized, allowing for improved outcomes compared to when only regulation is present. Although the supervisory signal is informative on average, we assume that it is imperfect and can point to the wrong bank action on occasions, in which case the corrective action would also prove incorrect. In practice, these supervisory interventions are taken through corrective supervisory actions, which are requests to BHCs to address unsafe or unsound practices and violations of law and can result in restrictions on bank activities and growth.

[2] In practice, a gray area can exist between supervision and regulation, as exemplified by stress testing and the fact that some supervisory activities involve compliance with regulations as previously noted (see, for example Goldsmith-Pinkham, Hirtle, and Lucca, 2016). In its purest form, regulation is written into law and enforced through courts, and thus requires verifiability (Maskin, Laffont, and Hildenbrand, 1982).

We assume that both supervisory monitoring and correction are costly for the banking authority. The optimally chosen supervisory strategy therefore trades off benefits and costs of supervision. Even though more supervision may always be better when only considering its benefits, optimal supervision will be limited overall with higher intensity warranted only if it results in higher marginal benefits than marginal costs.

Given these trade-offs, our model predicts that correction efforts are higher for riskier banks, which are those with worse observed signals. Consistent with this prediction, we find empirically that Fed supervisors' efforts increase between 70 and 140 percent for banking organizations whose confidential supervisory ratings indicate moderate to extreme levels of concern (ratings of 3, 4, or 5).

A key determinant of the absolute level of hours is bank size which varies widely among U.S. BHCs. Although we do not necessarily measure hours of all supervisory staff, the data suggest that institutions with assets below $10 billion of consolidated assets (trend-adjusted) are allocated about 100 hours per calendar quarter compared to about 1,500 for institutions above that asset threshold. Our model reveals that two opposite effects are at play with respect to bank size. If the optimal intensity of supervision is independent of bank size, then technological scale economies of supervision would suggest that supervisory hours increase less than proportionally with bank size. However, the optimal intensity of supervision may increase with bank size, for example if spillovers are sufficiently large for large banks. In the data, we confirm that supervisory hours are increasing in bank size.

We find that the elasticity of supervisory hours with respect to asset size, while clearly positive, is always close to or less than 1 across a number of different specifications—hours increasing less than proportional with bank size. This includes the post-2008 subsample, which has been characterized by heightened concerns about spillovers from stress at large and complex banking institutions. This evidence points to the existence of economies of scale in bank supervision that are sufficiently strong to outweigh the effect of enhanced supervision for larger banks. This result also suggests that, in terms of realized hour allocations, banks in our sample do not appear to have grown to be "too large to be supervised." A direct implication of this finding is that aggregate supervisory costs would increase if large banks were broken up, although these costs are clearly not the only factor to be considered when assessing bank size from a financial stability perspective.

As a check to our results, we compare estimates on Fed supervisory hours to those implied by the schedule of assessment fees that the Office of the Comptroller of the

Currency (OCC) charges banks it supervises. We find similar effects with respect to size and risk. We also study how hours may be related to other bank characteristics, such as headcounts, asset and institutional complexity.

We then extend the theoretical framework to multiple banks. Supervisory efforts are reallocated from banks with good signals to those with bad signals when resources are fixed. Resources also shift towards certain banks when their (actual or perceived) failure spillovers increase. Accordingly, we find in the data that when the fraction of bank assets in distress (rating of 4 or 5) in a Fed district increases, supervisory hours are allocated away from other institutions in that district. Similarly, while resources increase for the largest banks post-2008, they decline for smaller ones. We use variation implied by this reallocation as instruments to study the effect of supervision on bank risk. We find that the year-ahead standard deviation of a bank's return on assets (ROA) and the probability of bad outcomes (rating of 4 or 5, or failure) increase as a result of these shifts in attention. We finally use the model to interpret the empirical relation between ROA volatility, failures and supervisory attention to assess the effectiveness of supervision.

Related literature: The existing literature typically derives the need for supervision and regulation of banks from the special nature of banks as opposed to other firms and the resulting frictions. As for any debt-financed firm, a bank's limited liability raises basic moral hazard issues (Jensen and Meckling, 1976). However, there are limits to the ability of markets to provide the necessary discipline (Flannery, 1998; Rochet, 2004). The main debt of banks is in the form of deposits which impedes the functioning of market discipline through various channels. For example, the sequential service constraint inherent in deposit contracts creates strategic complementarities among depositors that can lead to inefficient runs (Diamond and Dybvig, 1983). To address the problem of runs, depositors are insured which eliminates their incentive to monitor the bank and creates a need for supervision/regulation (for example Mishkin, 2001). Alternative approaches view supervision/regulation as more effective monitoring than that of small, dispersed and uninformed depositors (Dewatripont and Tirole, 1994) or point to market incompleteness specific to banks' business models as a rationale for supervision/regulation (Merton, 1995). Since the financial crisis of 2008-09, the literature has focused more on macro-prudential as opposed to micro-prudential regulation, meaning not only on the default risk of an individual bank in isolation but also taking into account systemic effects such as spillovers of defaults on other banks and the economy more widely (for example Acharya, 2009; Brunnermeier et al., 2009; Tirole, 2013).

Few papers in the literature explicitly distinguish between supervision and regulation or focus on supervision. Supervision then takes on a number of different roles: auditing bank asset values to detect breaches of capital requirements (Rochet, 2007); preventing banks from taking observable but non-verifiable actions (Dewatripont and Tirole, 1994); incentivizing banks to take the right non-observable actions through punitive interference after bad, verifiable outcomes (Marshall and Prescott, 2001, 2006); incentivizing proper risk-taking and information disclosure through ex-post interventions (Harris and Raviv, 2012); and screening different risk-types of banks by offering a menu of admissible combinations of risk-taking and bank capital requirements (Prescott, 2004). Motivated by the inclusion of supervision as a pillar in the Basel framework, some papers study the interaction of supervision and regulation. For example, Bhattacharya et al. (2002) study the optimal combination of random supervisory exams and closure rules conditional on exam results; Decamps et al. (2004) have supervisors choosing intervention thresholds to maintain adequate incentives for bank risk taking and study the effects of ex-post liquidity assistance and forbearance.

Some empirical work studies the information produced by bank examinations when compared to off-site monitoring (Cole and Gunther, 1995), conditional on positive/negative findings (Berger and Davies, 1998) or relative to publicly available information (Hirtle and Lopez, 1999). In addition, there is a small literature taking on the challenge of identifying effects of supervision on bank outcomes (for example Hirtle, Kovner, and Plosser, 2016, and references therein).

The rest of the paper is organized as follows. Section 2 provides institutional detail used as a foundation to the model presented in Section 3. Section 4 describes the data. Section 5 studies the main determinants of supervisory attention at an individual bank. Section 6 focuses on the issues of supervising multiple banks with limited resources and Section 7 studies the impact of supervision. Section 8 concludes.

2 Institutional detail: What do supervisors do?

This section, which draws heavily on Eisenbach et al. (2015), provides a brief institutional overview of the list of activities in which bank supervisors engage. While some of the information discussed is pertinent to supervision of depository institutions, we focus on supervision of BHCs, which is the subject of the empirical investigation. We use this institutional detail as the foundation of the theoretical framework in the next section. Supervisory activities can be broadly classified into two main categories: monitoring and intervention through corrective actions.

In their monitoring activities, bank examiners meet with senior and business line management, with risk and control departments, and with members of a BHC's board of directors. Supervisors conduct independent reviews of regular or ad-hoc internal reports, such as those related to a bank's risk position, performance, budget, and strategy. These independent assessments often involve an analysis of inconsistencies in reports received across the firm. The largest BHCs are typically monitored by a dedicated staff that is assigned to the same institution on an ongoing basis. Assessments of these institutions, which typically hold consolidated assets above $10 billion, take place under "continuous monitoring," and supervisors take stock of these continuous activities once a year in so-called "roll-ups." Smaller institutions are monitored with a staff that rotates over a portfolio of banks, and comprehensive periodic assessments (typically yearly) are based on full-scope examinations.[3] Full-scope examinations and annual roll-ups for the largest institutions culminate in the assignment of a confidential supervisory rating.[4]

Bank holding companies are assigned a 1-to-5 rating under the "RFI/C(D)" rating system, with lower numbers indicating fewer issues and thus a better rating. Banks with a rating of 1 or 2 are considered in satisfactory condition, presenting few significant supervisory concerns. Banks with a 3, 4, or 5 rating present moderate to extreme levels of regulatory concerns.[5] Prior to 2004, BHCs received supervisory ratings known as BOPECs, an acronym which stood for five areas of supervisory concern: conditions of the BHC's bank subsidiaries, other nonbank subsidiaries, parent company, earnings and capital adequacy. BOPECs were also assigned on a 1-to-5 scale, and BOPECs and RFI/C(D) rating levels have similar supervisory interpretations. As a result, we splice these measures together in our analysis.

[3]Examination frequency for smaller BHCs is detailed in the Federal Reserve Supervision and Regulation Letter (SR) 13-21.

[4]Sometimes examination activities only involve limited aspects of bank activities and either one or multiple banks, in so-called "horizontals," where a portfolio of banks are examined along the same dimensions. Post-2008, these horizontals have developed in a number of recurrent assessment programs for the largest and most complex BHCs, such as the Comprehensive Capital/Liquidity Analysis and Review (CCAR/CLAR), also known as stress tests.

[5]The letters in the rating system indicate different components considered in the rating assignment—"R" is for risk management, "F" is for financial condition, "I" is for potential impact of the non-depository entities in the holding company on the depository institution(s) in the holding company, "C" is for the composite rating (that is, the overall rating considering and weighing the ratings on "R", "F", and "I"), and "D" is the rating assigned to the depositories (for example commercial banks or thrifts) owned by the holding company.

Supervisory interventions aim at modifying BHCs' behaviors and typically reduce their risk positions. Interventions take place through corrective supervisory actions, which are requests to BHCs to address unsafe or unsound practices and violations of law. These actions impose timelines and, as a firm remediates the issue, restrictions on a BHC's asset growth and set of activities as well as mandated divestitures of certain assets. Actions can be either public (formal, more severe) or confidential (informal, less severe). Confidential actions vary in their severity and scope and range from matters requiring attention and immediate attention (MRAs and MRIAs) to memorandums of understanding (MOUs) and 4(m) agreements ("4Ms").[6] Formal actions include suspensions/removals, cease and desist orders, prompt corrective actions (restore capital) and deposit insurance threats.

Section 3 formalizes the interaction of monitoring and intervention in a simple model. In the empirical section of the paper we study the model's predictions with data on supervisory hours, ratings, and actions. We use information on all formal actions and the most severe confidential ones, MOUs and 4Ms.

3 Economic Model of Supervision

We lay out a general framework that contains a range of elements of supervision and regulation as they can be interpreted through an economic lens. The purpose of explicitly specifying a model is to translate the features of supervision and regulation in practice into concepts familiar to economists from theories of asymmetric information and contracting. After meaningfully distinguishing regulation from supervision, our analysis focuses on the elements of supervision.

Our model is essentially static and is meant to capture the interaction of banks and supervisors over a certain period, for example one year. We therefore focus on a sequence of events within a single period and leave the analysis of dynamics over multiple periods for future research. We first describe the primitives of the model without supervision/regulation.

[6]MOUs typically encompass multiple deficiencies and often incorporate restrictions on a firm as a firm remediates the concerns raised in the MOU. 4M agreements are issued when a BHC is either engaged in non-permissible activities, or when the holding company or one of its depository institution subsidiaries is either inadequately capitalized or not well managed. Non-compliance with these actions result in fines, restrictions on growth, or divestment of certain assets.

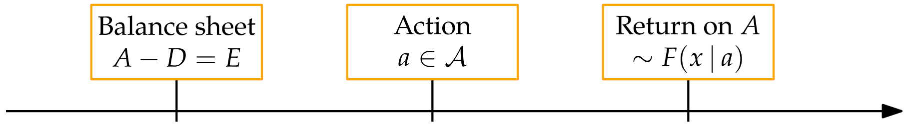

Figure 2: Model timeline without supervision or regulation

Bank: There is a set of banks $i = 1, \ldots, I$ that start the period with assets A_i, debt D_i, and equity E_i. Bank i has available a range of activities $F_i \in \mathcal{F}_i$ that it can engage in, where F_i represents the c.d.f. of the gross return x_i on the bank's assets in activity F_i.[7] We think of activities very broadly, such as making loans, trading securities, writing derivatives, and subdivisions and combinations thereof. Once the activity F is in place, the bank has to take an action $a \in \mathcal{A}$ that influences the distribution of its asset return x, that is $F(x \,|\, a)$. We think of the action a as risk taking, that is increasing the variance of x, and the relevant set of actions $\mathcal{A}$ could depend on the activity F.

The asset return x is realized at the end of the period and determines whether the bank is solvent or not. Figure 2 illustrates the timeline without supervision or regulation. Given assets A and debt with face value D, a bank is solvent at the end of the period if its equity is positive, $Ax - D \geq 0$, and insolvent otherwise. We denote by $U(z, a)$ the bank's utility which depends on its monetary payoff z at the end of the period as well as its action a, for example through effort costs. The bank's expected payoff in the absence of supervision or regulation is therefore

$$\int_0^{\frac{D}{A}} U(0, a)\, dF(x \,|\, a) + \int_{\frac{D}{A}}^{\infty} U(Ax - D, a)\, dF(x \,|\, a). \tag{1}$$

Note that we assume that while the bank's monetary payoff is 0 in case of default, the effort cost has to be borne irrespective of the bank's solvency since the action a is taken before the realization of x. By not distinguishing between the bank's shareholders and managers, we abstract away from potential incentive problems within the bank.

Banking authority: The banking authority can differ from the banker in three ways. First, while the bank only values its cash flow conditional on being solvent, the banking authority values payoffs to the bank *and* its creditors, $(Ax - D) + D = Ax$. This represents micro-prudential objectives, in the sense that the authority cares about the

[7] To simplify notation, we drop the index i for a specific bank unless explicitly analyzing multiple banks.

financial health, or safety and soundness, of a single bank.[8] Second, the banking authority also cares financial stability per se, representing macro-prudential objectives, in that it takes into account spillovers from banks to the rest of the economy. We capture this in reduced form with negative spillovers $-N(A) < 0$ of a bank failing. Note that spillovers depend on the bank's size A since a larger bank with more assets causes more spillovers in default. Finally, the banking authority could have a different utility $V \neq U$ over payoffs, for example with more risk aversion or with more weight on tail outcomes. The banking authority's expected payoff in the absence of supervision or regulation is therefore

$$\int_0^{\frac{D}{A}} V\big(Ax - N(A)\big)\, dF(x\,|\,a) + \int_{\frac{D}{A}}^{\infty} V\big(Ax\big)\, dF(x\,|\,a). \tag{2}$$

Regulation and supervision: In practice, there is significant overlap between regulation and supervision. In this paper, we want to focus on the parts of supervisions that are distinct from regulation. We therefore draw a stark distinction between regulation and supervision using the difference between verifiability and observability.[9]

In its purest form, regulation is written into law and enforced through courts, and therefore requires *verifiability*. In our model, the bank's initial balance sheet variables A, D, E can be thought of as verifiable so regulation can impose ex-ante constraints on the balance sheet, such as a capital requirement ρ requiring a minimum ratio between equity capital and assets, $E/A \geq \rho$. In economic terms, this provides both loss absorption capacity and incentives through "skin in the game." Further, the activities a bank is engaged in can be thought of as verifiable so regulation can limit ex-ante the available activities to a subset $\widetilde{\mathcal{F}} \subset \mathcal{F}$, for example allow lending and trading but not insurance. The 1933 Glass-Steagall Act and the Volcker Rule in the 2010 DFA are examples of such broad regulatory restrictions. In economic terms, this imposes some bounds on the risks available but cannot rule out risk entirely. The realization of the bank's asset return x at the end of the period is verifiable so regulation can be active ex-post, contingent on the realization of x. Since the capital requirement ρ has to be satisfied at the end of the period as well, or $Ax - D \geq \rho Ax$, it is possible for the bank to violate the capital requirement without defaulting, $Ax - D \in (0, \rho Ax)$. In this case, regulation can

[8] In general equilibrium and with complete markets, the pricing of the bank's debt would force it to take into account the creditor's payoff. We follow the banking literature in assuming that markets are incomplete, for example due to deposit insurance or coordination problems among small creditors (see the discussion in the literature review).

[9] See Maskin, Laffont, and Hildenbrand (1982) for a discussion of the concepts of verifiability and observability.

impose an ex-post penalty P on the bank even if it remains solvent. In economic terms, the combination (ρ, P) provides incentives through a disciplining effect.[10]

In contrast to regulation, we assume that supervision has discretion and only requires *observability*. We assume that while none of the bank's possible actions $a \in \mathcal{A}$ are verifiable, some are observable by supervisors while some are neither verifiable nor observable. For example, supervisors may be able to observe if a bank's risk managers participate in important meetings but not whether they are diligent in raising concerns. The supervisors can therefore prevent risk managers from not doing their work at all but still have to rely on incentives to ensure the work is done properly. Formally, this means that supervision can prevent some observable actions, that is restrict the action set to some subset $\widetilde{\mathcal{A}} \subset \mathcal{A}$. In economic terms, this makes for easier incentivizing by eliminating some "tempting" actions.[11] For simplicity, we assume that based on observability, the supervisor effectively limits the bank's choice set to two actions, $\widetilde{\mathcal{A}} = \{\overline{a}, \underline{a}\}$, where action $\underline{a}$ is the "good" action, low risk-taking with proper risk management, while action $\overline{a}$ is the "bad" action, high risk-taking with lax risk management.

In areas where even observability fails, we assume that supervisors can collect information through a monitoring effort m, for example an on-site examination that involves gathering information and distilling it into a supervisory rating. We can think of this rating as an interim signal r that is informative about the bank's action a and therefore its riskiness, the distribution of the final return x conditional on a. In economic terms, to the extent that the signal is a sufficient statistic for a with respect to x, supervision is more efficient in incentivizing the bank than regulation is contingent on final cashflow x. Given the binary choice set $\{\overline{a}, \underline{a}\}$ for the bank, we assume that the signal (the supervisory rating) is also binary, $r \in \{\overline{r}, \underline{r}\}$. Importantly, the signal is not perfect but its precision can be affected by the supervisory monitoring effort m. For simplicity, we assume that the supervisor can choose the precision directly:

$$\Pr[\underline{r} \mid \underline{a}] = \Pr[\overline{r} \mid \overline{a}] = m$$

We then assume that after observing the signal r about the bank's action a but before the realization of the return x, supervision can react to the interim information with

[10] In principle, regulation could condition on x more generally, that is with an ex-post cost/transfer $P(x)$. We abstract away from such issues since they have already been explored in the literature on regulation.

[11] In principle, it is not clear where to draw the line between an activity F and an action a. For the purposes of our model, we implicitly draw the line between verifiability and observability. Everything verifiable corresponds to F while everything observable (and unobservable) corresponds to a.

supervisory corrective actions c. As discussed in Section 2, these actions are requests to BHCs to address unsafe or unsound practices and often involve restrictions on a bank's asset growth and set of activities. In terms of the model, these corrective actions affect the distribution of the final payoff, $F(x \mid a, c)$, in a way opposite to a, that is by reducing the variance of x.

Denoting the probability of default by $\mathrm{PD}(a, c) \equiv F(D/A \mid a, c)$ and using subscripts for partial derivatives, this means $\mathrm{PD}_a(a, c) > 0$ and $\mathrm{PD}_c(a, c) < 0$.[12] In addition, since PD is bounded between 0 and 1, we assume that both a and c affect the tails of the distribution at a decreasing rate, that is $\mathrm{PD}_{aa}(a, c) < 0$ and $\mathrm{PD}_{cc}(a, c) > 0$.[13]

Costs of supervision: For tractability, we assume that both the bank and the supervisor are risk neutral and that their utilities are additively separable between payoffs and effort costs.[14] For the bank, we assume that supervisory correction c affects its utility directly, for example because the supervisory remedies involve extra work for the bank. The bank's utility is then given by

$$U(z, a, c) = z - Q(a, A) - Q(c, A).$$

For the supervisor, both monitoring m and correction c are costly; the supervisor's utility from payoff is then given by

$$V(z, m, c) = z - C(m, A) - C(c, A).$$

Note that the actions a and c affect the distribution of the bank's gross return x, which does not depend on bank size. Similarly, m serves to generate a signal about a, also independent of bank size. We therefore account for the effect of bank size on costs by including A as an explicit argument of the cost functions for the bank and the supervisor. For example, diligent risk management, measured by its effect on the return distribution, is more costly for assets of \$10 billion than for assets of \$1 billion. Since our data is on supervisory hours, it would be natural to limit the supervisor's costs to the

[12] More generally, this means that for $x > \mathrm{E}[x]$, $F_a(x \mid a, c) > 0$ and $F_c(x \mid a, c) < 0$; and for $x < \mathrm{E}[x]$, $F_a(x \mid a, c) < 0$ and $F_c(x \mid a, c) > 0$; and therefore $F_{xa} < 0$ and $F_{xc} > 0$.

[13] More generally, for $x > \mathrm{E}[x]$, $F_{aa}(x \mid a, c) > 0$ and $F_{cc}(x \mid a, c) < 0$; and for $x < \mathrm{E}[x]$, $F_{aa}(x \mid a, c) < 0$ and $F_{cc}(x \mid a, c) > 0$.

[14] Since our focus is on supervision, we will refer to the banking authority as the supervisor. As we discuss more below, this abstracts somewhat but not fully from incentive problems on the supervisor's part—as in "who monitors the monitors" (Hurwicz, 2007).

hours necessary to achieve intensity s on assets of size A multiplied by an hourly wage w, or $C(s, A) = w \times h(s, A)$. However, we prefer to maintain a general cost function C that can also include non-monetary costs, for example discomfort experienced by the supervisor when having to confront a bank.

Next, we make some intuitive assumptions about the shape of the cost function C. First, we assume that both are increasing and convex in the supervisory effort, that is $C_s > 0$ and $C_{ss} > 0$. For monitoring m, a more precise signal is costly and increasingly so since precision m is bounded above by 1. For correction c, reducing risk is costly and increasingly so since the variance of x is bounded below by 0. Second, we assume that costs are increasing and weakly concave in bank size, or $C_A > 0$ and $C_{AA} \leq 0$. Achieving the same signal precision or risk reduction for a larger bank incurs more costs. However, costs don't grow more than proportionally with size; in practice, this is reflected, e.g. by the fact that any bank has at most one chief risk officer or that larger banks tend to have larger and not just more loans.[15] Third, we assume that the cross-partial derivative is positive, that is $C_{As} > 0$. Here it is important to remember that m and c are scale-free: The total cost of achieving a given "intensity" of monitoring or risk reduction is higher for banks with more assets, $C_A > 0$; therefore, the total cost of increasing the intensity is naturally higher for banks with more assets, $C_{As} > 0$.

We could consider a number of other arguments for the cost function C, for example a bank's complexity. For simplicity, we limit the analysis to heterogeneity in bank size which we consider most important empirically. Finally, in an effort to limit the number of moving parts in the analysis, we don't explicitly consider banks' choice of initial balance sheet A, D, E or its choice of activity F. For simplicity we assume that banks' initial balance sheet composition is scale invariant by setting leverage $D/A \equiv \ell$ for all A.

Figure 3 illustrates the timeline of the model, taking into account supervision and regulation. Of course, many more elements could be added to the model. For example, we could assume that regulation is able to impose a penalty on the bank even in case of default. Alternatively, we could consider prevention of default through a bailout that benefits the bank's existing management implying a positive payoff in case of insolvency. For the purposes of this paper, we try to strike a balance between addressing what we consider key elements of supervision and maintaining a tractable model that allows for clear intuition about the main mechanisms.

[15]See Kovner, Vickery, and Zhou (2014) for evidence of similar patterns in banks' own operating costs.

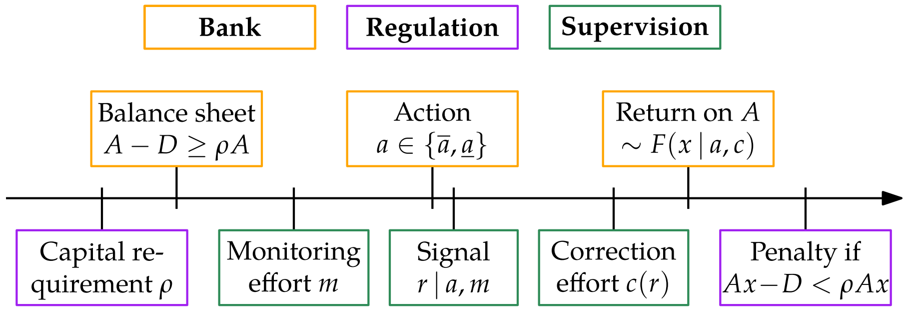

Figure 3: Model timeline with supervision and regulation

4 Data description

We use four data sources in the empirical analysis of this paper: hours data for supervisory staff at the Federal Reserve; examination information, supervisory rating, and action information from the National Examination Database (NED); financial data for domestic BHCs; and fee-schedule information of the OCC. We discuss each data source and then present summary statistics for the variables included in the regressions.

We are first to obtain information on hours spent by Federal Reserve supervisory staff from an internal database tracking information on staff activity and time allocation. The Federal Reserve supervises state member banks (SMBs) and all bank holding companies on a consolidated basis. Because we do not observe hours at non-SMB banks, which are supervised by non-Fed examiners, we focus on BHCs, which are exclusively under the purview of Fed supervisors. The hours data starts in 1998 and ends in 2014. The information is self-reported by supervisory employees on a weekly basis, and the database aims at recording each employee's time allocation across activities including monitoring banking institutions, e.g. through examinations, and other administrative and support responsibilities. Reported monitoring activity includes information on the monitored institution through its regulatory entity number (RSSD ID).

In practice, however, hours information is not always complete. First, the Federal Reserve System consists of 12 distinct districts and the Board of Governors, and reporting practices across districts may differ. We account for these differences with district and/or bank fixed effects. Second, we access information on hours pre-2000 for only a handful of districts, and have information on all districts only starting in 2006. Finally, supervisory work at the smallest institutions is often recorded using a generic bank portfolio assignment, as opposed to an institution RSSD number. By cross-checking hours information with independent information on the timing of supervisory inspection from NED, we find that consistent monitoring information with valid supervised-entity information is only available for institutions with assets of about $750 million or more; we therefore exclude institutions with less than $1 billion in assets. For each institution, we aggregate data by quarter, so that the resulting supervisory hours data is a dataset uniquely identified by a quarter and the supervised institution's RSSD ID.

We match hours information to two other data sources. First, we obtain information on bank characteristics, such as size, share of loans to total assets and employee counts, from public FR Y-9C reports, which are used to assess and monitor the financial condition of holding company organizations on a consolidated basis. The minimum asset requirement for Y-9C reporting increased from $150 million earlier in the sample to $500 million in 2006:Q1, but as discussed above we only include banks with assets greater than $1 billion. In addition, we match supervisory hours to confidential rating information and actions.

Finally, we use data on supervisory fees assessed on federally chartered commercial banks by the OCC. As the hours data discussed above can be noisy, we use this information to assess validity of the main findings. The OCC supervises nationally chartered commercial banks as well as federal savings associations (FSAs) since 2011 following the integration of the Office of Thrift Supervision (OTS) into the OCC. Similar to the Federal Reserve, the OCC is a self-funded agency that does not depend on the Congressional appropriation process. But while the Federal Reserve funds its supervisory activities through net interest margins earned on its securities portfolio (or seigniorage), the OCC levies assessments, fees, and other charges on federally chartered banks to meet the expenses of carrying out its supervisory activities. The OCC assesses semi-annual fees on its supervised entities under 12 U.S.C. 13 and 12 CFR 8. The fee schedule is adjusted by the OCC each year and determines fees as a function of bank size and bank risk, as measured by confidential supervisory ratings. We obtain data on the fee schedule from the OCC's public website.[16]

Table 1 provides summary statistics for the variables included in the regression specifications by whether banks' assets are above or below a $10 billion threshold (trend-adjusted for the growth rate of industry assets).[17] As shown in the first column of the table, on average, Fed supervisors allotted about 490 hours per calendar quarter to supervising a holding company in our sample. In terms of an 8-hour work day, and 63

[16] www.occ.treas.gov/topics/examinations/assessments-and-fees/index-assessments-fees.html. See also Kisin and Manela (2014) for another work using this same information.

[17] Appendix B provides exact variable definitions and constructions.

work days per quarter, these hours can be converted into about a full-time examiner being assigned to the holding company. But large institutions have on average three times as many hours assigned (column 3), while smaller ones about a fifth.[18]

5 Supervision at the individual bank level

We first study the determinants of supervisory efforts m and c at an individual bank, using our model to derive the key trade-offs and then studying the relationships in the data on Fed supervisory hours. In Section 6, we show that the supervision of multiple banks adds an additional trade-off across banks.

5.1 Theoretical trade-offs

Analyzing our model by backward induction, we start with the supervisory correction c conditional on the observed signal r. Then we proceed to the bank's choice of a given the supervisory monitoring m and correction strategy $c(r)$. Finally, we study the supervisory monitoring m that determines the precision of the signal, taking into account the bank's response. Our subgame perfect approach imposes time consistency on the supervisory actions, but supervisors would be better off ex-ante if they were able to commit to a strategy that is suboptimal ex-post and therefore time-inconsistent.

Supervisory correction: The supervisory correction c has to be conditional on the signal r received about the bank's action a.[19] For a given signal r, the supervisor's expected payoff when choosing c is given by the expected asset return net of the expected default spillover:

$$\mathrm{E}[Ax] - \mathrm{E}\big[\mathrm{PD}(a,c) \,\big|\, r\big]\, N(A) - C(c, A)$$

[18] This calculation excludes hours (that may be substantial) that have not been booked by the examiner to a specific institution. In addition, the day-count translation would underestimate an actual headcount because it abstracts from time-off and other administrative or training activities that an examiner may be involved when not assigned to a bank.

[19] The probability that the bank took an action a, for a signal realization r, is given by Bayes' rule, e.g.

$$\Pr[\bar{a} \,|\, \bar{r}] = \frac{m \Pr[\bar{a}]}{m \Pr[\bar{a}] + (1-m)\,\big(1 - \Pr[\bar{a}]\big)}.$$

The supervisor's first-order condition for the optimal c^* is then

$$-\,\mathrm{E}\big[\mathrm{PD}_c(a, c^*)\,\big|\,r\big]\,N(A) = C_c(c^*, A). \tag{3}$$

The left-hand side of equation (3) is the marginal benefit of correction c: the reduction in expected default probability multiplied by the spillover losses of default, $N(A)$. The right-hand side is the marginal cost of correction c at a bank of size A. As discussed in Section 3, the costs $C(c, A)$ certainly include the hours $h(c, A)$ necessary to achieve a level c of correction at a bank of size A. In addition, the costs (and therefore marginal costs) could include, for example the non-monetary discomfort experienced by the supervisor when imposing correction measures that the bank disagrees with. If these non-monetary marginal costs are high, the trade-off in the first-order condition (3) leads to low levels of correction—just as if the monetary costs are high or supervisory resources scarce.

We are interested in the effect of perceived riskiness and of bank size on supervisory hours spent on correction. Since the signal r corresponds to the supervisor's assessment of the bank's riskiness, we consider the effect of r on hours H which implement the optimal correction, $H = h(c^*, A)$:

$$\frac{dH}{dr} = h_c(c^*, A)\,\frac{dc^*}{dr} \tag{4}$$

Since $h_c > 0$, the sign is determined by that of dc^*/dr, that is how the intensity of correction depends on the bank's rating. Implicit differentiation of the first-order condition (3) yields

$$\frac{dc^*}{dr} = \frac{\big(\mathrm{PD}_c(\underline{a}, c^*) - \mathrm{PD}_c(\overline{a}, c^*)\big)\,\frac{d}{dr}\mathrm{Pr}[\overline{a}\,|\,r]\,N(A)}{\mathrm{E}\big[\mathrm{PD}_{cc}(a, c^*)\,\big|\,r\big]\,N(A) + C_{cc}(c^*, A)}.$$

The denominator is positive since $\mathrm{PD}_{cc} > 0$ and $C_{cc} \geq 0$. The sign of dc^*/dr therefore depends on the sign of the numerator, specifically on whether the sensitivity of the default probability to supervisory correction is greater at $\overline{a}$ than at $\underline{a}$, that is whether $\mathrm{PD}_{ac} \lesseqgtr 0$.

Proposition 1. *Riskier banks receive more supervisory hours if and only if corrective action has greater impact on the probability of default for high bank risk,* $\mathrm{PD}_{ac}(a, c) < 0$.

If the condition in Proposition 1 is satisfied, which we can confirm in the data, we have $dc^*/dr > 0$ and therefore more supervisory correction after a bad signal than after a good signal. Since there are two possible signal realizations $r \in \{\overline{r}, \underline{r}\}$, we have a supervisory correction strategy $c^*(\overline{r}) = \overline{c}$ and $c^*(\underline{r}) = \underline{c}$ with $\overline{c} > \underline{c}$.

Supervisory monitoring: Given the correction strategy $c(r)$ derived above, we now consider the bank's choice of a and then the supervisor's choice of monitoring intensity m. We denote the bank's expected utility from taking action a and then being subject to supervisory correction c by $\mathcal{U}(a, c)$:

$$\begin{aligned} \mathcal{U}(a,c) = \big(1 - \mathrm{PD}(a,c)\big)\, \mathrm{E}[Ax - D \mid Ax > D] \\ - Q(a, A) - Q(c, A) \end{aligned}$$

As with $\mathcal{U}$, we denote the supervisor's expected payoff conditional on the bank's action a and the supervisory correction c by $\mathcal{V}(a, c)$:[20]

$$\mathcal{V}(a,c) = \mathrm{E}[Ax] - \mathrm{PD}(a,c)\, N(A) - C(c, A)$$

Given the signal $r \in \{\overline{r}, \underline{r}\}$ and corresponding supervisory corrections $c \in \{\overline{c}, \underline{c}\}$, the bank has an incentive to take the good action $\underline{a}$ if

$$m\,\mathcal{U}(\underline{a}, \underline{c}) + (1 - m)\,\mathcal{U}(\underline{a}, \overline{c}) \geq m\,\mathcal{U}(\overline{a}, \overline{c}) + (1 - m)\,\mathcal{U}(\overline{a}, \underline{c}). \tag{5}$$

The left-hand side of condition (5) is the bank's expected payoff from choosing $\underline{a}$: With probability m, the supervisor will observe the right signal $\underline{r}$ (low risk) and choose the appropriate low level of correction $\underline{c}$; with probability $1 - m$, however, the supervisor will observe the wrong signal $\overline{r}$ (high risk) and mistakenly choose the high level of correction $\overline{c} > \underline{c}$. The right-hand side of condition (5) is the bank's expected payoff from choosing the bad action $\overline{a}$, where $\overline{c}$ is the correct supervisory response and $\underline{c}$ means the supervisor doesn't catch the bank and therefore doesn't intervene appropriately.

Solving condition (5) for m, we see that the bank chooses the good action as long as the signal is sufficiently precise, that is the supervisor is sufficiently likely to correctly

[20] Note that $\mathcal{V}(a, c)$ contains the cost of c but not the cost of m.

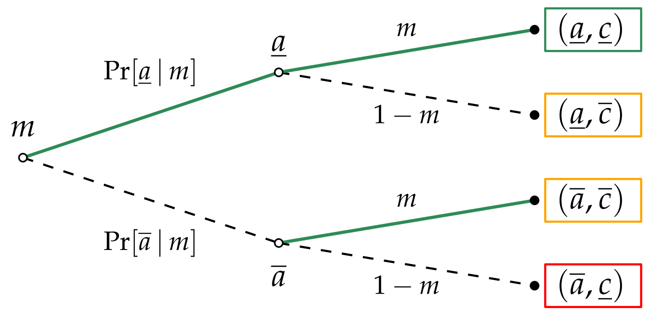

Figure 4: Uncertainty facing the supervisor when choosing monitoring

observe the bank's riskiness:[21]

$$m \geq \hat{m} := \frac{\mathcal{U}(\overline{a}, \underline{c}) - \mathcal{U}(\underline{a}, \overline{c})}{\mathcal{U}(\underline{a}, \underline{c}) - \mathcal{U}(\underline{a}, \overline{c}) + \mathcal{U}(\overline{a}, \underline{c}) - \mathcal{U}(\overline{a}, \overline{c})}$$

We assume that the supervisor has some uncertainty about the difference in the bank's effort cost Q between the actions $\underline{a}$ and $\overline{a}$. This can also be uncertain from the bank's perspective—the only important feature is that the bank knows its payoff difference at the time of choosing a while the supervisor doesn't, so that there remains some uncertainty about which action the bank chooses. This implies that the supervisor doesn't know the cutoff $\hat{m}$ precisely but views it as distributed according to a c.d.f. G.[22]

Figure 4 illustrates the uncertainty facing the supervisor when choosing the signal precision through monitoring effort m. For a given precision m, the the bank will choose $\underline{a}$ if $\hat{m} \leq m$, that is with probability $G(m)$, and $\overline{a}$ if $\hat{m} > m$, that is with probability $1 - G(m)$. In either case, the supervisor will sometimes make mistakes due to observing the wrong signal: If the bank chooses $\underline{a}$, the supervisor observes the wrong signal and mistakenly intervenes with probability $1 - m$ while if the bank shirks, the supervisor mistakenly does not intervene with probability $1 - m$. The supervisor's expected payoff

[21]For this cutoff to be well-behaved, we need (i) $\mathcal{U}(\overline{a}, \underline{c}) > \mathcal{U}(\underline{a}, \overline{c})$, that is getting away with shirking is better than being disciplined when behaving, (ii) $\mathcal{U}(\underline{a}, \underline{c}) > \mathcal{U}(\underline{a}, \overline{c})$, that is when behaving it is better if the supervisor knows, and (iii) $\mathcal{U}(\overline{a}, \underline{c}) > \mathcal{U}(\overline{a}, \overline{c})$, that is when shirking it is better if the supervisor doesn't know.

[22]Suppose from the supervisor's perspective $Q(\underline{a}, A) - Q(\overline{a}, A) = \Delta Q + \varepsilon$ where ε is uncertain with some c.d.f. $\widehat{G}(\varepsilon)$. Then ε drops out of the denominator of $\hat{m}$ but remains in the numerator so that $\hat{m}$ is uncertain with c.d.f. $G(\hat{m}) = 1 - \widehat{G}\big(\hat{m}\,(\mathcal{U}(\underline{a}, \underline{c}) - \mathcal{U}(\underline{a}, \overline{c}) + \mathcal{U}(\overline{a}, \underline{c}) - \mathcal{U}(\overline{a}, \overline{c})) - \mathcal{U}(\overline{a}, \underline{c}) + \mathcal{U}(\underline{a}, \overline{c})\big)$.

when choosing the monitoring effort m therefore is

$$\begin{aligned} &G(m)\left(m\,\mathcal{V}(\underline{a},\underline{c})+(1-m)\,\mathcal{V}(\underline{a},\overline{c})\right)\\ &\quad+\left(1-G(m)\right)\left(m\,\mathcal{V}(\overline{a},\overline{c})+(1-m)\,\mathcal{V}(\overline{a},\underline{c})\right)\\ &\quad-C(m,A)\end{aligned}$$

The supervisor's first-order condition for the optimal m^* is then

$$\begin{aligned} &\overbrace{G_m(m^*)\left(m^*\mathcal{V}(\underline{a},\underline{c})+(1-m^*)\,\mathcal{V}(\underline{a},\overline{c})-m^*\mathcal{V}(\overline{a},\overline{c})-(1-m^*)\,\mathcal{V}(\overline{a},\underline{c})\right)}^{\text{Benefit of higher likelihood of bank choosing } \underline{a}}\\ &\qquad+\underbrace{G(m^*)\left(\mathcal{V}(\underline{a},\underline{c})-\mathcal{V}(\underline{a},\overline{c})\right)+\left(1-G(m^*)\right)\left(\mathcal{V}(\overline{a},\overline{c})-\mathcal{V}(\overline{a},\underline{c})\right)}_{\text{Benefit of higher likelihood of correct signal}}=C_m(m^*,A). \end{aligned} \tag{6}$$

Similar to the first-order condition (3) for correction c, this first-order condition for m equalizes the marginal benefit of increased monitoring (left-hand side) with its marginal cost (right-hand side). While there is only one benefit of correction—reducing the probability of default—the marginal benefit of monitoring in equation (6) has two parts, illustrated by the solid green branches in Figure 4: First, a more precise signal makes it more likely that the bank chooses the right action $\underline{a}$. Second, a more precise signal makes it more likely that the supervisor receives the right signal and implements the correct response, $\underline{c}$ for $\underline{a}$ and $\overline{c}$ for $\overline{a}$.[23]

The effect of bank size: We now turn to the effect of bank size on supervisory hours. This is an important question to address since our analysis so far has been in terms of the scale-free "intensities" of supervision—monitoring and corrective action. However, there is significant heterogeneity in the size of banks supervised by the Federal Reserve, and considerable policy debate has been focused on the importance of size due to potentially non-linear effects.

The main question for our analysis is whether the hours spent on supervision scale proportionally with bank size, that is whether the size elasticity of hours with respect

[23]For a unique optimal m^*, we need the left-hand side of (6) to be decreasing in m^*. Jointly sufficient conditions for this are (i) $G''(m^*) \leq 0$ and (ii) $\mathcal{V}(\underline{a},\underline{c})-\mathcal{V}(\underline{a},\overline{c}) < \mathcal{V}(\overline{a},\overline{c})-\mathcal{V}(\overline{a},\underline{c})$. For a symmetric distribution G, the first condition is satisfied if $G(m^*) \geq 1/2$, that is if in equilibrium the bank is at least as likely to choose $\underline{a}$ than $\overline{a}$. The second condition says that choosing the correct supervisory correction is more important when the bank chooses $\overline{a}$ than when it chooses $\underline{a}$. We assume both conditions to be satisfied.

to assets is larger or smaller than one:

$$\frac{dH}{dA}\frac{A}{H} \lesseqgtr 1$$

In contrast to the effect of risk on supervisory hours in (4), the effect of assets has a direct and an indirect part on optimal hours $H = h(s^*, A)$ for both optimal intensities of supervision $s^* \in \{m^*, c^*\}$:

$$\frac{dH}{dA}\frac{A}{H} = \Big(\underbrace{h_A(s^*, A)}_{\text{direct}} + \underbrace{h_s(s^*, A)\frac{ds^*}{dA}}_{\text{indirect}}\Big)\frac{A}{H} \tag{7}$$

The first, direct part is the increase in hours when increasing the size of the bank while maintaining the same intensity of supervision—which itself is scale free. The second part is the change in hours if the first-order condition, (3) for c^* or (6) for m^*, implies different optimal supervisory intensities for different bank sizes.

Proposition 2. *Two competing effects determine whether supervisory hours scale proportionally with bank size:*

1. *Technological scale economies in supervision can lead to hours increasing less than proportionally with bank size.*

2. *Increasing supervisory intensity can lead to hours increasing more than proportionally with bank size.*

With a Cobb-Douglas hourly cost function, $h(s, A) = s^\sigma A^\alpha$, the expression (7) simplifies to clarify the intuition:

$$\frac{dH}{dA}\frac{A}{H} = \alpha + \sigma\frac{ds^*}{dA}\frac{A}{s}$$

Technological scale economies in supervision imply that $\alpha < 1$ which can lead the whole expression to be less than 1. However, since marginal costs of intensity are increasing, $\sigma > 1$, the effect of scale economies is counteracted if optimal intensity increases with bank size, $ds^*/dA > 0$. The main determinants of the relationship between supervisory intensity and bank size are the spillover losses $N(A)$, to which we next turn.

The effects of spillovers: The optimal level for both supervisory activities $s \in \{m, c\}$ is determined by a first-order condition, (3) for c^* and (6) for m^*, setting the marginal benefit of s equal to its marginal cost:

$$\text{benefit}_s(s^*, A) = C_s(s^*, A) \tag{8}$$

For both supervisory activities, the main benefit is to avoid the spillover losses $N(A)$. The size of spillovers therefore directly affects the optimal intensity of supervision.

Corollary 1. *An increase in supervisory concern about spillovers implies higher supervisory monitoring and correction.*

For supervisory correction, the marginal benefit is a reduction in the expected spillover loss due to a reduced probability of default. If the spillovers are larger, the marginal benefit of correction is greater and warrants an increase in the level of correction to rebalance marginal benefit and marginal cost. For supervisory monitoring, the marginal benefit is an increase in both the likelihood that the bank chooses the "good" action $\underline{a}$ and the likelihood that the supervisor correctly observes the bank's action and therefore chooses the appropriate correction. Both are more important if the spillovers are larger, and therefore warrant more supervisory monitoring.

We now turn to the effect of the slope of $N(A)$ on optimal supervision, i.e. how quickly spillovers increase in bank size. Implicit differentiation of the first-order condition (8) yields some insight into the properties of ds^*/dA:

$$\frac{ds^*}{dA} = \frac{\text{benefit}_{sA}(s^*, A) - C_{sA}(s^*, A)}{-\text{benefit}_{ss}(s^*, A) + C_{ss}(s^*, A)}$$

The denominator is positive since the marginal benefit is decreasing and the marginal cost is increasing. Therefore, the sign of ds^*/dA depends only on the numerator, i.e. comparing the effect of bank size on the marginal benefit and on the marginal cost of supervision:

$$\frac{ds^*}{dA} > 0 \quad \Leftrightarrow \quad \text{benefit}_{sA}(s^*, A) > C_{sA}(s^*, A)$$

How the marginal benefit of supervision changes with bank size is therefore determined by the effect of size on spillovers.

Corollary 2. *If spillovers increase sufficiently strongly with bank size, then the optimal intensity of supervision increases with bank size.*

Summary: Combining the results of Proposition 2 and Corollary 2, our model points to two competing effects that determine whether hours spent on monitoring m and correction c should increase more or less than proportionally with a bank's asset size. On the one hand, the existence of technological scale economies would directly imply that supervisory hours increase less than proportionally with size. On the other hand, the optimal intensity of supervision—if it is higher for larger banks—would imply that hours increase more than proportionally. While the effect of technological scale economies is unambiguous, the competing effect of changing intensity requires that banks' spillovers increase sufficiently with size in order to outweigh the higher costs of increasing supervisory intensity.

5.2 Empirical evidence

Fed supervisory hours: Table 2 presents our baseline regression specification of log supervisory hours on log bank assets and a range of controls. The coefficient on log assets captures the elasticity of supervisory hours with respect to bank assets, that is if hours increase proportionally with assets, the coefficient would be equal to 1 and if less than proportionally, less than 1.

The elasticity of supervisory hours to bank assets is 0.96 (column 1) and 0.68 when including bank fixed effects (column 2). In other words, the pooled and within variation in the data suggests an elasticity close to, or less, than 1. This means that a bank with \$10 billion in assets receives less than twice the amount of attention than a bank with \$5 billion in assets. In terms of our model, the fact that hours increase less than proportionally with size is evidence of technological scale economies in supervision, that is achieving a certain intensity of supervision requires less supervisory resources per dollar of assets at a large bank than at a small bank. The empirical finding of an overall size-coefficient less than 1 does not rule out the possibility that larger banks warrant and receive a higher intensity of supervision, which would on its own imply a size coefficient greater than one. However, if this effect exists, it appears quantitatively dominated by the technological scale economies.

To assess the linearity of the estimated relation between supervisory hours and bank size, Figure 7 shows binned scatter plots and the relationship implied by the linear specification without (left panel) and including bank fixed effects (right panel). As is evident from the figure, the estimated relation between hours and assets is approximately linear in log-scale.

Riskier banks, as measured by their supervisory rating, also display increased supervisory attention, as predicted by Proposition 1 (column 1 pooled, column 2 within). For example, as compared to a bank with the best possible rating of 1, which is the omitted category in the regression, a bank of the same size but with a rating of 3 receives 70% more hours (within variation), an effect roughly equivalent to doubling the size of the bank. A bank rated the worst-possible 5 receives 136% additional hours, roughly equivalent to the baseline hours of a bank three times its size.

In practice, two important components of supervision that are also reflected in the hours measure, are ensuring that a bank is in compliance with regulations as well as assessing its internal processes. We expect these components of supervision, as well as the monitoring and correction captured by our model, to demand increasing resources with increasing bank complexity. Because complexity and size are often related, the elasticity of supervisory hours to size may be measuring the joint effect of banks' size and complexity. We attempt to control for the effect of organizational complexity by controlling the log number of legal entities within each bank holding company.[24] Hours display a statistically significant elasticity with respect to bank complexity of 0.26 (column 3, pooled variation) and 0.14 (column 4, within variation), consistent with the idea that some possible economies of scale are lost as banks become more complex. Also, including the complexity measures lowers the estimated size-elasticity to 0.77 in the pooled variation regression. Because bank complexity does not vary much and is thus mostly accounted for by the fixed effect, the size-elasticity of hours is largely unaffected by the inclusion of the complexity measure.

In sum, the baseline regression specification of Table 2 implies the elasticity of supervisory hours with respect to assets is less than one, and that riskier banks as well as more complex banks require additional attention.

Fed supervisory actions: The model splits supervisory efforts in monitoring and intervention, but the data on hours does not. We can nonetheless observe intervention efforts indirectly using supervisory actions. In particular, we examine all formal actions and the most-severe informal ones (MOUs and 4Ms). This data verifies that it is indeed the case that corrective actions increase as the rating worsens. Table 3 provides regression estimates of the sensitivity of supervisory actions outstanding at each point in time on size and rating when including time and bank effects. As predicted by the model, actions respond to risk (but not to size; columns 1 and 3). The peak in terms of

[24]See Appendix B for an exact definition.

the volume of informal actions (column 3) is for a rating of three. Banks with this rating have about .5 more issues than banks rated one. Formal actions are monotonically increasing with a peak of .9. Adding the two effects, starting with a rating of three, banks are subject to an increasing volume of actions, with the composition also shifting from less-severe (confidential) actions to formal ones.

OCC assessment fees: To further validate the conclusions drawn from the hours data, we compare the estimated elasticities of supervisory hours with respect to bank size and risk to those of assessment fees collected by the Office of the Comptroller of the Currency (OCC) on its supervised entities. In contrast to the hours data, this fee data is a more direct measure of the supervisory cost function as fees are expressed in dollar terms. However, because of potential cross-subsidies across different bank-size or risk categories, the assessment schedule may not be directly informative of the supervisory production function at an institution level. We find that size and risk elasticities of assessment fees turn out to be very similar to those estimated on Federal Reserve supervisory hours.

The OCC's base assessment is calculated using a table with eleven categories, or brackets, each of which comprises a range of asset-size values. In addition to the base amount, which is the same for every bank in its asset-size bracket, the fee includes a marginal amount, which is computed by applying a marginal assessment rate to the assets in excess of the lower bound of the asset-size bracket. The marginal assessment rate declines as asset size increases, "reflecting economies of scale in bank examination and supervision" (Federal Register Vol. 79, No. 81, April 28, 2014).

Table 4 provides summaries for semiannual assessments (meaning that annual fees are twice as large) as a function of assets in 2007 and 2014 that we obtain from OCC bulletins. The 2014 fee structure includes a new bracket for the largest banks, with assets greater than $250 billion. This additional bracket was introduced to help the OCC recover additional costs associated with supervising large and complex banks. Starting in 2001, the OCC began imposing a surcharge of 25% on their original (size-based) assessment for national banks with a 3, 4, or 5 rating, to "reflect the increased cost of supervision" (OCC 2000-30). By 2004, the size of the surcharge had been increased to 50% for 3-rated banks and to 100% for 4- or 5-rated banks.[25]

[25] With the exception of the addition of the $250 billion asset bracket, asset brackets and base/marginal fee schedules prior to 2007 were stable over time, except for an annual inflation adjustment. Both inflation adjustments and rating surcharges were capped at $20 billion, prior to 2014, and at $40 billion thereafter.

We apply the fee structure to the universe of nationally chartered commercial banks using asset information as of 2006:Q4 and 2013:Q4 (relevant periods for fee calculations in 2007 and 2014) and compute the implied scale economies pre- and post-2008 by regressing log fees on log assets and controls in Table 5. The elasticity of OCC fees to assets is 0.7 which is nearly identical to the within-bank estimate using Fed hours data (see Table 2, column 2). The increase in OCC fees with respect to bank risk is similar although not as steep as the estimated increase in Fed hours. Relative to 1-rated institutions, fees increase by about 40% on average for 3-rated institutions and by about 70% for 4- or 5-rated institutions.

In Figure 8, we present a binned scatter plot and the linear relationship obtained when regressing OCC fees on bank assets (in log scale). Comparing these results with those obtained using Fed supervisory hours (Figure 7), we again see a very similar pattern. We discuss changes in the OCC fee schedule post-2008 in the next section.

Additional determinants of supervisory hours: To investigate further the determinants of supervision in practice, we first consider the sensitivity of supervisory hours with respect to alternative measures of bank size and complexity. We then study seasonal clustering of supervisory hours.

The model assumes that supervisory efforts are devoted to acquiring information (monitoring), and enforcing corrective actions (intervention), related to the riskiness of a bank's assets. Consistently in our main specification (Table 2), we relate hours to a bank's risk and size in terms of its assets. Alternatively, one could consider a model where supervisors are monitoring bank employees to ensure that they exert reasonable effort. We directly consider this alternative by expanding the set of regressors in the baseline regression specification of Table 2 to include the log of a bank's total employees. We see a statistically significant elasticity of about 0.2 in the pooled estimation (Table 6, column 1), which is no longer significant when considering within-bank variation (column 2). This indicates that the sensitivity to a bank's head count is either relatively small or not significant, justifying the focus on assets in our model and empirical analysis.

Besides the organizational complexity in our baseline specification (number of legal entities), the complexity of the consolidated balance sheet may also be important. To account for this possibility, we consider two measures of asset complexity: (i) the share of loans to total assets, which is higher for traditional, less complex banking organizations, and the concentration of the bank's activities as measure by the Herfindahl index of the balance sheet shares of different asset categories. We find that neither measure of balance sheet complexity enters significantly in our estimated regressions (columns 3 and 4).

Finally, as noted in Section 2, large banks are "continuously" monitored while efforts at smaller ones are clustered around full-scope examinations. Because our estimates are based on quarterly data, and ratings are changed around examinations, our point estimates may be affected by the within-year time patterns. We next use information on whether a bank is undergoing a full-scope examination and include a dummy variable for whether such an event occurs in a given quarter. We interact this indicator variable with another for large and smaller institutions. The point estimates (Table 6, columns 5 and 6) imply that supervisory efforts increase during a full scope examination, with the larger banks experiencing an increase of about 100%, and the smaller banks of about 200%. In unreported results, we also find that increases in hours around full-scope examinations are even more muted for larger banks than for smaller banks. Overall point estimates are not much affected by the inclusion of the examination dummies, but the precision of the estimates is improved and we therefore include them in the specifications below.

6 Supervisory efforts when dealing with multiple banks

We now turn to the issues in supervising multiple banks, in particular the allocation of scarce supervisory resources across several banks in need of attention. In practice each Federal Reserve district supervises the banks located within the district.

6.1 Theoretical analysis

In the theoretical analysis so far, we have considered the optimal level of a supervisory action $s \in \{m, c\}$ for a single bank with first-order conditions for optimal supervision of the form

$$\text{benefit}_s(s, A) = C_s(s, A). \tag{9}$$

In practice, supervisors have to deal with multiple banks and are allocating a given budget of supervisory resources $\overline{C}$. Note that we can invert the cost function C to infer the level of supervisory action $s_i = C^{-1}(C_i, A_i)$ that results from an allocation of C_i resources to bank i with size A_i. Then we can state the problem of allocating resources across a set of banks as

$$\max_{\{C_i\}} \left\{ \sum_i \text{benefit}\big(C^{-1}(C_i, A_i), A_i\big) \right\} \quad \text{subject to} \quad \sum_i C_i \leq \overline{C}.$$

With a Lagrange multiplier μ on the budget constraint, this yields first-order conditions of the form

$$\text{benefit}_s(s_i, A_i)\,\frac{1}{C_s(s_i, A_i)} = \mu \quad \text{for all } i. \tag{10}$$

Comparing the first-order conditions (9) and (10), we see that they are structurally equivalent and that the Lagrange multiplier μ in the multiple-banks version acts like a shadow wage that ties together all individual bank first-order conditions. This means that all the comparative statics we derive above for individual banks carry over to the analysis of multiple banks. In addition, the first-order conditions require that the ratio of marginal benefit to marginal cost be equalized across banks:

$$\frac{\text{benefit}_s(s_i, A_i)}{C_s(s_i, A_i)} = \frac{\text{benefit}_s(s_j, A_j)}{C_s(s_j, A_j)} \quad \text{for all } i, j \tag{11}$$

The fact that first-order conditions are linked through the Lagrange multiplier μ implies that supervisory attention has to be reallocated across banks in response to shocks. For example, since the two actions m and c happen sequentially—first monitoring to collect information, then corrective action conditional on the information—satisfying condition (11) for both requires reallocating supervisory hours after the signals $\{r_i\}$ are realized, away from banks with good ratings and towards banks with bad ratings. From the perspective of an individual bank, a shock such as more bad signals at other banks corresponds to an increase in the Lagrange multiplier, since the shock implies that supervisory resources are more scarce than before. Similarly, a shock that increases spillovers $N(A)$ also implies an increase in the Lagrange multiplier.

Proposition 3. *Shocks that increase the scarcity of resources and therefore the Lagrange multiplier μ lead to a decrease in attention to unaffected banks:*

$$\begin{aligned}
\frac{dH}{d\mu} &= h_s(s^*, A)\,\frac{ds^*}{d\mu} \\
&= h_s(s^*, A)\,\frac{-C_s(s_i, A_i)}{-\text{benefit}_{ss}(s_i, A_i) + \mu C_{ss}(s_i, A_i)} < 0
\end{aligned}$$

6.2 Empirical evidence

We provide evidence of reallocation when resources become scarce. First, we consider how attention to small banks may have changed post-2008 as supervisors have reassessed potential spillovers $N(A)$ from the largest banks. Second, we study how attention toward a bank changes, holding its risk fixed, when other banks in the same district become stressed.

Changes in supervision post-2008: While Figure 1 shows that Fed supervisory staff experienced a significant expansion post-financial crisis, this is by no means conclusive evidence of more intensive supervision. First, a number of financial institutions, including investment banks and credit card companies, became bank holding companies during the financial crisis, so much of the expansion may be because of the increase in total BHC assets. In our hours data, we can start to disentangle the two. Figure 9 shows that total Fed hours increased about 50% post-financial crisis, matching the increase in employment. As shown by the blue area, while newly formed BHCs accounted for a significant fraction of the increase in total hours, hours at existing BHCs also increased, with much of the increase being accounted for by BHCs with assets greater than $10 billion. In contrast, smaller institutions appeared to have experienced a reduction in supervisory hours. While this pattern suggests a reallocation of attention from small to large banks, it does not account for possible changes in bank risk.

We therefore study this issue with our regression model. We first analyze post-2008 changes in the elasticity of hours with respect to bank size by adding to the baseline regression specification an interaction between log assets and a dummy variable marking the post-2008 sample. Hours sensitivity increased by a relatively small amount of about 0.1 in absolute terms in the post-2008 period (Table 7, columns 1 and 2). To study the effect at different bank sizes, we split the post-2008 level increase in hours between large and small banks. For the large banks (assets $\geq$ $10 billion), there is a notable level increase on the order of 60% in the post-2008 period (columns 3 and 4). In contrast, hours spent at the small institutions declined by about 20–30%. These point estimates thus confirm a reallocation away from small banks amid a gradual resource expansion. In Figure 10, we illustrate this effect graphically. We first estimate residuals from our baseline specification without bank fixed effects (Table 2, column 1). We then average residuals of log hours by quarter and size-group, and plot their evolution over time. Residuals for large banks (blue) are about 50% larger in the post-2008 sample, as opposed to smaller institutions for which we observe a decline.

In Table 3, we also show that supervisory actions declined in the post-2008 sample for the smallest banks (columns 3 and 4). While it is possible to conceive that post-2008, supervisors became less concerned about the smallest institutions at the same time that attention increased at the largest banks, our model suggests that when supervisory resources are scarce and relatively inelastic, an increase in the optimal relative attention to large banks can then only be implemented by an absolute reduction of attention at small banks. We confirm this intuition with empirical estimates of the pre/post-2008 size elasticity of OCC fees.

As discussed in Section 5.2, the overall elasticities of OCC fees to bank size and risk are very similar to those of Fed hours. If the decline in Fed supervisory hours at small banks is due to lower concern, we would expect to see a similar decline in the size-elasticity of OCC fees for small banks. If, however, the decline is due to limited resources, we would not expect any change in OCC fees for small banks. In 2008, the OCC introduced an additional fee bracket for banks with assets greater than $250 billion to cover increasing costs associated with the supervision of large and complex banks. The elasticity of fees to bank assets is modestly affected by this additional bracket (Table 5, column 2), but on average, fees are about 50% higher in 2014 than in 2007 for banks with assets greater than $10 billion (column 3), while for smaller banks fees are essentially unchanged. In addition, the increase in fees for large banks matches very closely the increase in hours of Fed supervisors post-2008.

Attention reallocation within Federal Reserve districts: We now turn to reallocation of attention at the district level. We assume that, in the short run, resources are relatively fixed within each of the twelve Federal Reserve districts, allowing us to study the impact on supervisory attention at one bank when other banks in the same district warrant relatively more attention because they are under stress. At the district level, we therefore compute the share of assets in distress (assets of banks rated 4 or 5) and study its effect on hours at large and small banks.[26] As shown of Table 8, we find that an increase in stress at other banks in the district implies a sizable decline of attention at a given bank when controlling for the bank's risk and size. Point estimates imply that attention at a given bank would decline by 74 percent, should all remaining banks in a district become stressed (column 1, pooled variation) and 50 percent (column 2, bank fixed effects). The effects seem to be smaller for the largest banks (columns 3 and 4), but the difference is

[26]When calculating the share, we always exclude the bank under observation.

significant only in pooled variation (column 3).

Consistent with the decline in supervisory attention, we also find that supervisory actions at a bank of a given rating decline when other banks are stressed (columns 2 and 4 of Table 3).

In sum, both the post-2008 evidence for smaller banks and the results using cross-district variation support the prediction of Proposition 3. We use both in the next section to study the effectiveness of supervision.

7 Impact of supervision

In the model, we assume that an increase in supervisory attention can lower bank risk. In this section we study whether and to what extent this is the case empirically. Furthermore, we discuss the limits of assessing the effectiveness of supervision based on observable outcomes.

7.1 Identifying the impact of supervision

Identifying the impact of supervision is challenging because of a standard simultaneity issue. As per our regression estimates, more bank risk leads to more supervisory attention in terms of supervisory hours and actual corrective actions, and thus simply regressing hours or actions on outcomes would lead to positive coefficient—suggesting that supervisors *increase* bank risk. We exploit as an exogenous source of variation the previously documented reallocation of supervisory attention from small-to-large banks post-2008 and from any bank when others in a Fed district are under severe stress. In particular, we study whether, for a bank with given risk, outcomes worsen over the next year as attention is reallocated away. Outcomes are measured in two ways: first, as actual or near failures (rating of 4 or 5) over the next year and, second, as the year-ahead standard deviation of a BHC's return on assets (ROA).

BHC failures are defined by a termination of the BHC in the regulatory data where one of two conditions exist: either (i) the reason for this termination is given as the failure of the BHC, or (ii) there exists a subsidiary of the BHC that is identified as failing within one quarter of the BHC termination. A bank fails in the next year if these conditions are satisfied at some point in the next four quarters. Because of the low incidence of actual failures, we additionally identify banks that have a failing rating in the next year, defined as banks that have a rating of 4 or 5 at some point in the next four quarters. While these events are not as severe as a failure, they do represent very negative outcomes, as banks' ability to function as credit providers is severely hampered with ratings of 4 or 5. The standard deviation of ROA at time t is calculated over the subsequent four quarters.

As discussed in Section 6, for given rating and size, hours and supervisory actions declined both for small banks post-2008 and for institutions in a district where a large share of other banks' assets are under stress (ratings of 4 and 5). Corresponding to this reduced supervisory attention, Table 9 shows that holding the current rating and size fixed, small banks have been 3 percent more likely to face an actual or near failure in the following year and that the standard deviation of their ROA is about 10 basis points higher irrespective of whether time effects are excluded (columns 1 and 4) or included in the specification (columns 2 and 5). Similarly, holding the current rating and size fixed, banks in a district with lots of stressed assets experience a much higher probability of bad outcomes and higher standard deviation of ROA when time effects are excluded. When time effects are included, point estimates have the same sign but loose significance at standard confidence levels, suggesting that much of the identification of district stress result from common cross-district stress variation, such as during the 2008–09 crisis period.

In studying the direct relation between the source of exogenous variation in attention and outcomes, these results are "reduced-form" IV estimates. For completeness, columns 3 and 6 of Table 9 show the corresponding IV estimates where the two measures above instrument for hours and attention. As shown in the table, doubling hours allotted to a firm is associated with a decline in the probability of bad outcomes of 5 percentage points (column 3) and a decline in the standard deviation of ROA of 11 basis points (column 6). These point estimates are consistent with the model assumption that an increase in supervisory attention is associated with a decline in bank risk.

In unreported results, we also studied the same outcomes for commercial banks housed within holding companies compared to those that are not and are therefore not subject to bank holding company consolidated supervision. A similar argument to the one above would suggest that all else equal and abstracting from selection issues, banks housed within BHCs will be safer because of the additional resources spent, and we indeed find this to be the case in the data.[27]

[27] Results are available upon request.

7.2 Assessing the effectiveness of supervision

The results in the previous section identify that supervision is effective at reducing the probability of bad outcomes as well as reducing the standard deviation of ROA. In order to assess the effectiveness, however, we would ideally like to study the counter-factual of no supervision. Table 10 lists the empirically observed probabilities of bad outcomes and standard deviations of ROA for different BHC ratings as well as the increase in supervisory attention coinciding with the ratings (coefficients from column 2 of Table 2). Note that all measures of bad outcomes are increasing in the supervisory rating even though supervisory attention increases significantly with ratings. What can we learn from the fact that bank outcomes are not equalized across ratings?

Default probabilities under optimal supervision: We can assess this issue through the lens of our model by comparing the theoretical default probabilities $\mathrm{PD}(a, c)$ that result from the four possible combinations of bank actions $a \in \{\overline{a}, \underline{a}\}$ and supervisory corrections $c \in \{\overline{c}, \underline{c}\}$.

Proposition 4. *Under optimally chosen supervision, the default probability* $\mathrm{PD}(a, c)$ *satisfies the following ranking:*

$$\mathrm{PD}(\underline{a}, \overline{c}) < \mathrm{PD}(\underline{a}, \underline{c}) < \mathrm{PD}(\overline{a}, \overline{c}) < \mathrm{PD}(\overline{a}, \underline{c})$$

There is clear intuition behind this ranking of the four default probabilities. First, it is clear that not intervening when necessary leads to the highest possible default probability $\mathrm{PD}(\overline{a}, \underline{c})$. Second, in case the bank chooses the good action $\underline{a}$, the "right" supervisory response $\underline{c}$ results in a *higher* default probability than the "wrong" response $\overline{c}$ so that $\mathrm{PD}(\underline{a}, \underline{c}) > \mathrm{PD}(\underline{a}, \overline{c})$. While this may seem unintuitive at first, it is due to the fact that correction is costly: The optimal supervisory response is chosen trading off costs and benefits so high correction is only worthwhile after a bad signal where it has high expected benefit. Third, the fact that correction is costly also implies that $\overline{c}$ will never fully "offset" the effect of $\overline{a}$ so that $\mathrm{PD}(\overline{a}, \overline{c}) > \mathrm{PD}(\underline{a}, \underline{c})$. Even though we have assumed that supervision could fully correct a bank's mistakes, the fact that it is costly means that optimal supervision will not attempt to fully correct a bank's mistakes.

Default probabilities conditional on ratings: Note, however, that the default probability $\mathrm{PD}(a, c)$ is conditional both on the bank's action a as well as the supervisor's response $c(r)$ while the default probabilities in Table 10 are conditional only on the rating

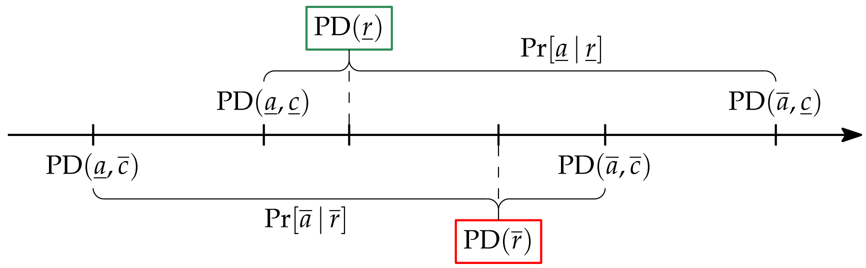

Figure 5: Relationship between empirical and theoretical default probabilities

r. The theoretical analogue to the empirically observed default probabilities therefore are

$$\begin{aligned}
\mathrm{PD}(\underline{r}) &= \Pr[\underline{a} \mid \underline{r}]\, \mathrm{PD}(\underline{a}, \underline{c}) + \big(1 - \Pr[\underline{a} \mid \underline{r}]\big)\, \mathrm{PD}(\overline{a}, \underline{c}), \\
\mathrm{PD}(\overline{r}) &= \Pr[\overline{a} \mid \overline{r}]\, \mathrm{PD}(\overline{a}, \overline{c}) + \big(1 - \Pr[\overline{a} \mid \overline{r}]\big)\, \mathrm{PD}(\underline{a}, \overline{c}).
\end{aligned}$$

Since the bank action is unobserved, the empirical default probabilities are convex combinations of the conditional default probabilities $\mathrm{PD}(a, c)$. As illustrated in Figure 5, it does not necessarily follow from Proposition 4 that the empirical default probabilities satisfy $\mathrm{PD}(\underline{r}) < \mathrm{PD}(\overline{r})$. However, we can give conditions for this to occur.

Corollary 3. *The empirical default probability* $\mathrm{PD}(r)$ *is increasing in rating r if signal precision m is sufficiently large and/or if* $\big|\mathrm{PD}(a, \underline{c}) - \mathrm{PD}(a, \overline{c})\big|$ *for* $a \in \{\overline{a}, \underline{a}\}$ *is sufficiently small.*

The fact that, in the data, the probability of default is significantly higher for banks with worse ratings is therefore indicative of supervisors investing heavily in monitoring so that their signals about bank actions are very precise and/or of limits on the effectiveness of supervisory correction. The latter possibility seems unlikely, given both the large documented response of supervisory attention to a bad rating and the considerable estimated reduction of bad bank outcomes due to supervision.

Unconditional default probabilities: The remaining question is what can be inferred by an outside observer after a bank default, without being able to condition on either bank action a or supervisor response c. Since the conditional default probability $\mathrm{PD}(a, c)$ is highest if the bank chooses $\overline{a}$ and the supervisor mistakenly chooses $\underline{c}$, one might conclude that this combination is the most likely cause when observing a bank failure. However, this reasoning would ignore the likelihood of the conditioning event $(\overline{a}, \underline{c})$. As illustrated by Figure 6, given optimal supervisory monitoring m^*, the bank

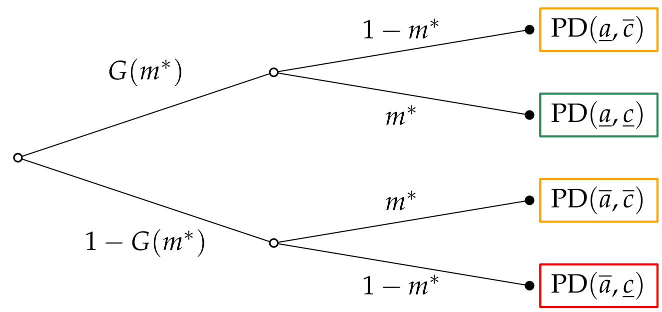

Figure 6: Event tree leading to conditional default probabilities

chooses the bad action $\overline{a}$ with probability $1 - G(m^*)$ and the supervisor misses this and responds only with $\underline{c}$ with probability m^*. The *unconditional* probability of observing a bank failure due to the bank taking high risk and the supervisor not noticing is therefore given by

$$\mathcal{P}(\overline{a}, \underline{c}) \equiv \big(1 - G(m^*)\big)\,(1 - m^*)\,\mathrm{PD}(\overline{a}, \underline{c}).$$

While the conditional default probabilities $\mathrm{PD}(a, c)$ can clearly be ranked from smallest to largest (Proposition 4), it is much harder to clearly rank the unconditional default probabilities.

Proposition 5. *Under optimal supervision, only an incomplete ranking of unconditional default probabilities is possible:*

$$\mathcal{P}(\overline{a}, \overline{c}) > \mathcal{P}(\underline{a}, \overline{c}) \tag{12}$$

$$\textit{and} \quad \mathcal{P}(\underline{a}, \underline{c}) > \mathcal{P}(\underline{a}, \overline{c}) \tag{13}$$

The first inequality, (12), states that we are more likely to observe a default where the bank chose the "bad" action $\overline{a}$ and the supervisor intervened appropriately with $\overline{c}$ than a default where the bank chose the "good" action $\underline{a}$ and the supervisor mistakenly intervened with $\overline{c}$. The second inequality, (13), states that a default where the bank chose the "good" action and the supervisor intervened appropriately with $\underline{c}$ is also more likely than one with $(\underline{a}, \overline{c})$. Both inequalities therefore involve the probability of a bank failure where the bank chose the low-risk action $\underline{a}$ but the supervisor received the wrong signal and chose to intervene with $\overline{c}$, further reducing the bank's risk. As shown in Proposition 4, this combination leads to the lowest conditional probability of default. Inequalities (12) and (13) tell us that such a bank failure is less likely than either case of correct supervisory response, $(\overline{a}, \overline{c})$ and $(\underline{a}, \underline{c})$, respectively.

The one unconditional default probability missing from the inequalities (12) and (13) is the one corresponding to the most troublesome scenario, where the bank takes high risk but it goes undetected by the supervisor, $(\overline{a}, \underline{c})$.

Corollary 4. *Without making further assumptions on parameters and functional forms, we cannot infer from a bank failure whether the scenario $(\overline{a}, \underline{c})$ was a more or less likely than any of the other three possibilities, $(\underline{a}, \overline{c})$, $(\overline{a}, \overline{c})$, and $(\underline{a}, \underline{c})$.*

8 Conclusion

This paper provides a new perspective on bank supervision by combining a theoretical framework and an empirical investigation that speak to the objectives and resource constraints of supervision. Our model uses building blocks from the theory of contracts and incomplete information to derive key implications of an optimally chosen supervision strategy. Taking into account that any supervisory strategy is subject to resource constraints, we highlight the trade-off between the costs and benefits of supervision and derive comparative statics that are confirmed in the data.

Larger banks receive more attention in the form of supervisory hours than smaller banks. However, hours increase less than proportionally with size. Through the lens of the model, this is evidence of technological scale economies in supervision that outweigh the effect of increased concerns about the larger banks. When studying resource allocation with multiple banks, we also find evidence of substitution effects of supervisory efforts indicating binding resource constraints. Using exogenous shifts in resource allocations, we also find evidence that supervision lowers the probability of failure and reduces the riskiness of returns, consistent with the model assumptions.

Much of the literature studying the role of regulation and supervision in the 2008 financial crisis has focused on possible policy distortions arising from institutional design (Agarwal et al., 2014; Carletti et al., 2015) or incentive problems (Lucca et al., 2014). This paper takes a complementary approach by focusing on the limits and trade-offs of supervision based on limited resources. However, the model can speak to some issues of supervisory incentives as well. For example, the non-verifiability of the supervisors' information opens up the possibility of incentive problems for the supervisors themselves. Our model allows supervisory costs to include non-monetary costs such as discomfort from confronting a bank, and therefore implicitly allows for a wedge between the strategy chosen by the supervisor and one chosen by a social planner. However, there are potentially much richer incentive issues in the spirit of "who monitors the monitors" (Hurwicz, 2007) that could be considered but are beyond the scope of this paper.[28]

Finally, our analysis imposes time consistency on the supervisor since we derive the ex-post optimal intervention strategy. If, instead, supervisors were able to commit to an intervention strategy ex-ante, risk-taking incentives for the bank could be improved further. However, supervisors would then have to defend imposing sub-optimally harsh corrective actions ex-post. While time-consistency issues apply to public policies in general (Kydland and Prescott, 1977), they are a particular challenge to bank supervision because of its opacity.

[28]See, for example, Masciandaro and Quintyn (2013) for an extensive survey on the governance of supervision.

References

Acharya, V. V. (2009). A theory of systemic risk and design of prudential bank regulation. *Journal of Financial Stability 5*(3), 224–255.

Agarwal, S., D. Lucca, A. Seru, and F. Trebbi (2014). Inconsistent regulators: Evidence from banking. *The Quarterly Journal of Economics*.

Berger, A. N. and S. M. Davies (1998). The information content of bank examinations. *Journal of Financial Services Research 14*(2), 117–144.

Bhattacharya, S., M. Plank, G. Strobl, and J. Zechner (2002). Bank capital regulation with random audits. *Journal of Economic Dynamics and Control 26*(7), 1301–1321.

Brunnermeier, M. K., A. Crockett, C. A. Goodhart, A. Persaud, and H. S. Shin (2009). The fundamental principles of financial regulation.

Carletti, E., G. Dell'Ariccia, and R. Marquez (2015). Supervisory incentives in a banking union. Working Paper.

Cetorelli, N. and S. Stern (2015). Same name, new businesses: Evolution in the bank holding company. *Liberty Street Economics Blog* (September 28, 2015).

Cole, R. A. and J. W. Gunther (1995). A camel rating's shelf life.

Decamps, J.-P., J.-C. Rochet, and B. Roger (2004). The three pillars of basel ii: optimizing the mix. *Journal of Financial Intermediation 13*(2), 132 – 155. Bank Capital Adequacy Regulation under the New Basel Accord.

Dewatripont, M. and J. Tirole (1994). *The prudential regulation of banks*. MIT Press.

Diamond, D. W. and P. H. Dybvig (1983). Bank runs, deposit insurance, and liquidity. *Journal of Political Economy 91*(3), 401–419.

Eisenbach, T. M., A. Haughwout, B. Hirtle, A. Kovner, D. O. Lucca, and M. C. Plosser (2015). Supervising large, complex financial institutions: What do supervisors do? *FRB of New York Staff Report* (729).

Flannery, M. J. (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.

Goldsmith-Pinkham, P., B. Hirtle, and D. O. Lucca (2016). Parsing the content of bank supervision. Working Paper.

Harris, M. and A. Raviv (2012). How to get banks to take less risk and disclose bad news. *Chicago Booth Research Paper* (12-55).

Hirtle, B., A. Kovner, and M. Plosser (2016). The impact of supervision on bank performance. Working Paper.

Hirtle, B. J. and J. A. Lopez (1999). Supervisory information and the frequency of bank examinations.

Hurwicz, L. (2007). But who will guard the guardians? *Nobel Prize Lecture*.

Jensen, M. C. and W. H. Meckling (1976). Theory of the firm: Managerial behavior, agency costs and ownership structure. *Journal of Financial Economics 3*(4), 305–360.

Kisin, R. and A. Manela (2014). Funding and incentives of regulators: Evidence from banking. Working Paper.

Kovner, A., J. Vickery, and L. Zhou (2014). Do big banks have lower operating costs? *Federal Reserve Bank of New York Economic Policy Review 20*(2), 1–27.

Kydland, F. E. and E. C. Prescott (1977). Rules rather than discretion: The inconsistency of optimal plans. *The journal of political Economy*, 473–491.

Lucca, D., A. Seru, and F. Trebbi (2014). The revolving door and worker flows in banking regulation. *Journal of Monetary Economics 65*(0), 17 – 32.

Marshall, D. A. and E. S. Prescott (2001). Bank capital regulation with and without state-contingent penalties. *Carnegie-Rochester Conference Series on Public Policy 54*(1), 139–184.

Marshall, D. A. and E. S. Prescott (2006). State-contingent bank regulation with unobserved actions and unobserved characteristics. *Journal of Economic Dynamics and Control 30*(11), 2015–2049.

Masciandaro, D. and M. Quintyn (2013). The evolution of financial supervision: the continuing search for the holy grail. *SUERF 50th Anniversary Volume Chapters*, 263–318.

Maskin, E., J. Laffont, and W. Hildenbrand (1982). The theory of incentives: An overview. In *Advances in Economic Theory (invited lectures from the 4th World Congress of the Econometric Society)*. Cambridge University Press.

Merton, R. C. (1995). A functional perspective of financial intermediation. *Financial Management*, 23–41.

Mishkin, F. S. (2001). Prudential supervision: Why is it important and what are the issues? In *Prudential Supervision: What Works and What Doesn't*, pp. 1–30. University of Chicago Press.

Prescott, E. S. (2004). Auditing and bank capital regulation. *Federal Reserve Bank of Richmond Economic Quarterly 90*(4), 47–63.

Rochet, J.-C. (2004). Macroeconomic shocks and banking supervision. *Journal of Financial Stability 1*(1), 93–110.

Rochet, J.-C. (2007). *Why Are there So Many Banking Crises? The Politics and Policy of Bank Regulation*. Princeton University Press.

Tirole, J. (2013). Financial regulation. In *Rethinking Macro Policy II: First Steps and Early Lessons*.

Table 1: Summary statistics. This table presents summary statistics for the variables included in the regression specifications. Large (small) BHCs are defined in terms of assets above (below) a $10 billion asset threshold adjusted for the growth rate of industry assets. For detailed variable definitions see Section 4 and Appendix B. Sample is 1998Q1-2014Q4.

	All		**Small BHCs**		**Large BHCs**	
	Mean	StDev	Mean	StDev	Mean	StDev
Hours	484.36	1357.14	95.53	168.89	1372.02	2202.04
Assets ($ millions)	36026.70	188138.43	2441.27	1611.61	112699.08	328264.78
Log(Hours)	4.14	2.15	3.33	1.73	6	1.87
Log(Assets)	8.43	1.50	7.64	0.54	10.26	1.38
Rating	1.95	0.74	1.98	0.79	1.90	0.63
Log(N BHC Subsidiaries)	2.48	1.31	1.92	0.78	3.78	1.37
Log(N BHC Employees)	7	1.47	6.23	0.63	8.75	1.33
Asset Conc. (HHI)	0.51	0.10	0.52	0.10	0.47	0.10
Loan Share	0.66	0.12	0.67	0.11	0.62	0.13
Ongoing Full Exam	0.28	0.45	0.24	0.42	0.38	0.48
Large BHC	0.30	0.46	0	0	1	0
Share of Problem District Assets	0.02	0.06	0.02	0.06	0.02	0.06
MOU/4M actions	0.12	0.41	0.09	0.31	0.19	0.57
Formal actions	0.07	0.29	0.05	0.24	0.09	0.38
ROA SD	0.28	0.32	0.30	0.33	0.25	0.29
Fail/Downgrade4-5	0.06	0.23	0.07	0.25	0.03	0.16
Observations	15364		10684		4680	

Table 2: Baseline specification for supervisory hours. This table presents estimates of supervisory hours on bank size (assets), complexity (number of subsidiaries) and dummy variables for supervisory ratings. Sample is 1998Q1-2014Q4. Standard errors clustered by bank reported in brackets. Significance: * $p < 0.1$, ** $p < 0.05$, *** $p < 0.01$

	Log(Hours)			
	(1)	(2)	(3)	(4)
Log(Assets)	0.96***	0.68***	0.77***	0.62***
	[0.02]	[0.11]	[0.04]	[0.12]
Log(N BHC Subsidiaries)			0.26***	0.14**
			[0.04]	[0.07]
Rating = 2	0.23***	0.15**	0.22***	0.13**
	[0.05]	[0.06]	[0.05]	[0.06]
Rating = 3	0.94***	0.70***	0.89***	0.66***
	[0.09]	[0.09]	[0.08]	[0.09]
Rating = 4	1.31***	1.08***	1.26***	1.03***
	[0.11]	[0.11]	[0.11]	[0.11]
Rating = 5	1.61***	1.36***	1.52***	1.29***
	[0.16]	[0.16]	[0.15]	[0.16]
District, Date FEs?	Yes	Yes	Yes	Yes
Bank FEs?	No	Yes	No	Yes
Adj. R^2	0.49	0.56	0.51	0.57
Obs.	17969	17969	17235	17235
Distinct BHCs	785	785	727	727

Table 3: Supervisory actions. This table presents estimates of supervisory confidential (MOUs and 4Ms) and formal actions on bank size (assets), dummy variables for supervisory ratings, each bank's district share of other assets in distress (banks rated 4 or 5), as well as an interaction between post-2008 and small-BHC dummies. Small BHCs are defined in terms of assets above (below) a \$10 billion asset threshold adjusted for the growth rate of industry assets. Sample is 1998Q1-2014Q4. Standard errors clustered by bank reported in brackets. Significance: * $p < 0.1$, ** $p < 0.05$, *** $p < 0.01$

	MOUs and 4Ms		Formal Actions	
	(1)	(2)	(3)	(4)
Log(Assets)	0.03	-0.01	-0.05*	-0.06**
	[0.03]	[0.03]	[0.03]	[0.03]
Rating = 2	-0.01	-0.00	-0.00	-0.00
	[0.02]	[0.02]	[0.01]	[0.01]
Rating = 3	0.52***	0.52***	0.13***	0.13***
	[0.05]	[0.04]	[0.02]	[0.02]
Rating = 4	0.34***	0.36***	0.61***	0.61***
	[0.06]	[0.06]	[0.06]	[0.06]
Rating = 5	0.08*	0.12***	0.88***	0.89***
	[0.04]	[0.04]	[0.04]	[0.04]
Post-2008 $\times$ (Small BHC)		-0.30***		-0.08**
		[0.07]		[0.04]
Share of Problem District Assets		-0.29***		-0.11*
		[0.11]		[0.07]
Date FEs?	Yes	Yes	Yes	Yes
Bank FEs?	Yes	Yes	Yes	Yes
Adj. R^2	0.46	0.48	0.54	0.55
Obs.	25196	24792	25196	24792
Distinct BHCs	930	866	930	866

Table 4: OCC general assessment fee schedule. This table shows the OCC assesment fee schedule on federally chartered commercial banks and savings association as a function of asset size. Source: 12 CFR 8 and OCC bulletins.

If the amount of the total balance sheet assets (consolidated domestic and foreign subsidiaries) is: ($ millions)		The Semiannual Assessment will be:		
Year 2007				
Over	But Not Over	This Amount ($)	Plus	Of Excess Over ($ millions)
0	2	5,480	0	0
2	20	5,480	0.000227454	2
20	100	9,574	0.000181963	20
100	200	24,131	0.000118274	100
200	1,000	35,958	0.000100078	200
1,000	2,000	116,020	0.000081883	1,000
2,000	6,000	197,903	0.000072785	2,000
6,000	20,000	489,043	0.000061932	6,000
20,000	40,000	1,356,091	0.000050403	20,000
40,000		2,364,151	0.000033005	40,000
Year 2014				
Over	But Not Over	This Amount ($)	Plus	Of Excess Over ($ millions)
0	2	5,997	0	0
2	20	5,997	0.000236725	2
20	100	10,258	0.000189379	20
100	200	25,408	0.000123092	100
200	1,000	37,717	0.000104156	200
1,000	2,000	121,041	0.000085218	1,000
2,000	6,000	206,259	0.000075749	2,000
6,000	20,000	509,255	0.000064454	6,000
20,000	40,000	1,411,611	0.000048553	20,000
40,000	250,000	2,382,671	0.000033132	40,000
250,000		9,340,391	0.0000328	250,000

Table 5: OCC general assessment fees. This table shows the relation between OCC general assessments as a function of commercial banks' assets and ratings. The fees are calculated for the universe of all federally chartered commercial banks that filed Call Reports in 2006:Q4 and 2013:Q4 using the fee schedule in Table 4 and rating surcharges discussed in Section 5.2. Assets are actual, while ratings are generated from a uniform distribution. The $10 billion asset threshold is expressed in nominal terms.

	Log(Fees)		
	(1)	(2)	(3)
Log(Assets)	0.70***	0.69***	0.68***
	[0.00]	[0.00]	[0.00]
Post-2008 × Log(Assets)		0.01***	
		[0.00]	
Post-2008 × (Assets ≥ $10bn)			0.47***
			[0.03]
Post-2008 × (Assets < $10bn)			0.03***
			[0.00]
Rating = 2	-0.01*	-0.01*	-0.01*
	[0.01]	[0.01]	[0.01]
Rating = 3	0.40***	0.40***	0.40***
	[0.01]	[0.01]	[0.01]
Rating = 4	0.68***	0.68***	0.68***
	[0.01]	[0.01]	[0.01]
Rating = 5	0.69***	0.69***	0.69***
	[0.01]	[0.01]	[0.01]
Constant	-0.03	-0.02	0.04*
	[0.02]	[0.02]	[0.02]
Adj. R^2	0.99	0.99	0.99
Obs.	2866	2866	2866
Distinct NAs	1772	1772	1772

Table 6: Additional determinants of supervisory hours. This table presents estimates of supervisory hours on alternative measures of bank size (employees) and complexity (loan share and asset concentration) as well as dummy variables indicating whether a banking institution is undergoing a full-scope examination. Large (small) BHCs are defined in terms of assets above (below) a $10 billion asset threshold adjusted for the growth rate of industry assets. Sample is 1998Q1-2014Q4. Standard errors clustered by bank reported in brackets. Significance: $^{*}\ p < 0.1$, $^{**}\ p < 0.05$, $^{***}\ p < 0.01$

	Log(Hours)					
	(1)	(2)	(3)	(4)	(5)	(6)
Log(Assets)	0.59***	0.53***	0.80***	0.63***	0.78***	0.65***
	[0.08]	[0.16]	[0.03]	[0.12]	[0.04]	[0.12]
Rating = 2	0.23***	0.13**	0.22***	0.13**	0.14***	0.05
	[0.05]	[0.06]	[0.05]	[0.06]	[0.05]	[0.06]
Rating = 3	0.90***	0.66***	0.90***	0.67***	0.71***	0.51***
	[0.08]	[0.09]	[0.09]	[0.09]	[0.08]	[0.08]
Rating = 4	1.27***	1.03***	1.27***	1.03***	1.05***	0.85***
	[0.11]	[0.11]	[0.11]	[0.11]	[0.10]	[0.11]
Rating = 5	1.54***	1.29***	1.54***	1.29***	1.22***	1.02***
	[0.15]	[0.16]	[0.15]	[0.17]	[0.15]	[0.16]
Log(N BHC Subsidiaries)	0.24***	0.13**	0.23***	0.14**	0.26***	0.16**
	[0.04]	[0.07]	[0.04]	[0.07]	[0.04]	[0.06]
Loan Share			-0.27	0.48		
			[0.52]	[0.77]		
Asset Conc. (HHI)			0.44	-0.49		
			[0.56]	[0.76]		
Log(N BHC Employees)	0.21**	0.12				
	[0.08]	[0.16]				
(Ongoing Exam) × (Large BHC)					1.01***	0.85***
					[0.07]	[0.07]
(Ongoing Exam) × (Small BHC)					1.99***	2.10***
					[0.04]	[0.04]
District, Date FEs?	Yes	Yes	Yes	Yes	Yes	Yes
Bank FEs?	No	Yes	No	Yes	No	Yes
Adj. R^2	0.51	0.57	0.49	0.55	0.62	0.68
Obs.	17234	17234	16846	16846	17235	17235
Distinct BHCs	727	727	716	716	727	727

Table 7: Supervisory hours by size and over time. This table extends the baseline regression specification in Table 2 to include interactions of asset size or size dummies with a post-2008 dummy. Large (small) BHCs are defined in terms of assets above (below) a \$10 billion asset threshold adjusted for the growth rate of industry assets. Sample is 1998Q1-2014Q4. Standard errors clustered by bank reported in brackets. Significance: * $p < 0.1$, ** $p < 0.05$, *** $p < 0.01$

	Log(Hours)			
	(1)	(2)	(3)	(4)
Log(Assets)	0.92***	0.65***	0.89***	0.63***
	[0.03]	[0.12]	[0.02]	[0.07]
Post-2008 × Log(Assets)	0.10***	0.13***		
	[0.03]	[0.03]		
Rating = 2	0.15***	0.08	0.14***	0.08
	[0.05]	[0.06]	[0.05]	[0.06]
Rating = 3	0.74***	0.54***	0.65***	0.49***
	[0.08]	[0.08]	[0.07]	[0.08]
Rating = 4	1.10***	0.94***	1.02***	0.88***
	[0.11]	[0.11]	[0.10]	[0.10]
Rating = 5	1.33***	1.16***	1.32***	1.15***
	[0.16]	[0.15]	[0.16]	[0.15]
Post-2008 × (Large BHC)			0.57***	0.60***
			[0.08]	[0.10]
Post-2008 × (Small BHC)			-0.33***	-0.19***
			[0.05]	[0.05]
District, Ongoing FEs?	Yes	Yes	Yes	Yes
Date FEs?	Yes	Yes	No	No
Bank FEs?	No	Yes	No	Yes
Adj. R^2	0.61	0.68	0.61	0.68
Obs.	17969	17969	17969	17969
Distinct BHCs	785	785	785	785

Table 8: Reallocation of supervisory hours. This table extends the baseline regressions specification in Table 2 to include each bank's district share of other assets in distress (banks rated 4 or 5). Large (small) BHCs are defined in terms of assets above (below) a $10 billion asset threshold adjusted for the growth rate of industry assets. Sample is 1998Q1-2014Q4. Standard errors clustered by bank reported in brackets. Significance: * $p < 0.1$, ** $p < 0.05$, *** $p < 0.01$

	Log(Hours)			
	(1)	(2)	(3)	(4)
Log(Assets)	0.89***	0.62***	0.88***	0.62***
	[0.02]	[0.10]	[0.02]	[0.10]
Rating = 2	0.17***	0.08	0.17***	0.08
	[0.05]	[0.06]	[0.05]	[0.06]
Rating = 3	0.73***	0.54***	0.73***	0.54***
	[0.08]	[0.08]	[0.08]	[0.08]
Rating = 4	1.12***	0.95***	1.12***	0.94***
	[0.11]	[0.11]	[0.11]	[0.11]
Rating = 5	1.41***	1.21***	1.42***	1.21***
	[0.16]	[0.15]	[0.16]	[0.15]
Share of Problem District Assets	-0.74***	-0.59**	-0.80***	-0.61**
	[0.29]	[0.26]	[0.28]	[0.26]
Post-2008 × (Large BHC)	0.91***	0.79***	0.78***	0.75***
	[0.09]	[0.10]	[0.10]	[0.10]
(Share Distress) × (Large BHC)			0.20**	0.07
			[0.09]	[0.08]
District, Date, Ongoing FEs?	Yes	Yes	Yes	Yes
Bank FEs?	No	Yes	No	Yes
Adj. R^2	0.62	0.68	0.62	0.68
Obs.	17943	17943	17943	17943
Distinct BHCs	780	780	780	780

Table 9: Supervisory outcomes following reallocation of supervisory hours. This table presents estimates of next-year's probability of being rated 4 or 5 or failing as well as the standard deviation of ROA in the next year on dummy variables for supervisory ratings, each bank's district share of other assets in distress (banks rated 4 or 5) as well as an interaction between post-2008 and small-BHC dummies. Small BHCs are defined in terms of assets above (below) a $10 billion asset threshold adjusted for the growth rate of industry assets. Sample is 1998Q1-2014Q4.

	(fail or rating to 4 or 5)$_{t+1,t+4}$			σ(ROA)$_{t+1,t+4}$		
	(1)	(2)	(3)	(4)	(5)	(6)
Log(Assets)	0.04***	0.03***	0.06***	0.11***	0.01	0.06*
	[0.01]	[0.01]	[0.02]	[0.02]	[0.02]	[0.03]
Rating = 2	0.01	0.01	0.02**	0.02	0.01	0.03
	[0.01]	[0.01]	[0.01]	[0.01]	[0.01]	[0.02]
Rating = 3	0.10***	0.09***	0.13***	0.08***	0.06***	0.14***
	[0.01]	[0.01]	[0.02]	[0.02]	[0.02]	[0.03]
Rating = 4	0.71***	0.70***	0.74***	0.17***	0.14***	0.25***
	[0.02]	[0.02]	[0.03]	[0.04]	[0.04]	[0.06]
Rating = 5	0.58***	0.58***	0.63***	0.30***	0.28***	0.43***
	[0.03]	[0.03]	[0.04]	[0.09]	[0.09]	[0.10]
Post-2008	-0.02**			-0.15***		
	[0.01]			[0.02]		
Post-2008 × (Small BHC)	0.03***	0.03***		0.09***	0.10***	
	[0.01]	[0.01]		[0.02]	[0.02]	
Share of Problem District Assets	0.17***	0.06		0.39***	0.04	
	[0.05]	[0.06]		[0.07]	[0.07]	
Log(Hours)			-0.05***			-0.11***
			[0.02]			[0.03]
Estimator	OLS	OLS	IV	OLS	OLS	IV
District, Ongoing FEs?	No	No	No	No	No	No
Date FEs?	No	Yes	No	No	Yes	No
Bank FEs?	Yes	Yes	No	Yes	Yes	No
Adj. R^2	0.57	0.58	0.22	0.25	0.32	-0.25
Obs.	24811	24811	17917	22512	22512	16271
Distinct BHCs	866	866	780	815	815	743

Table 10: Supervisory outcomes by rating. This table shows estimates of the sensitivity of log(Hours) to rating dummies from Table 2, next-year's probability of being rated 4 or 5 or failing on dummy variables for supervisory ratings, and standard deviation of ROA in the next year by rating category. Sample is 1998Q1-2014Q4.

rating_t	$\frac{\partial \log(\text{hours}_t)}{\partial \text{rating}_t}$	(fail or rating to 4 or 5)$_{t+1,t+4}$	$(\text{fail})_{t+1,t+4}$	$\sigma(\text{ROA})_{t+1,t+4}$
1	.	0.003	0.000	0.225
2	0.154	0.019	0.005	0.268
3	0.696	0.123	0.013	0.385
4	1.077	0.873	0.086	0.520
5	1.356	0.808	0.202	0.604

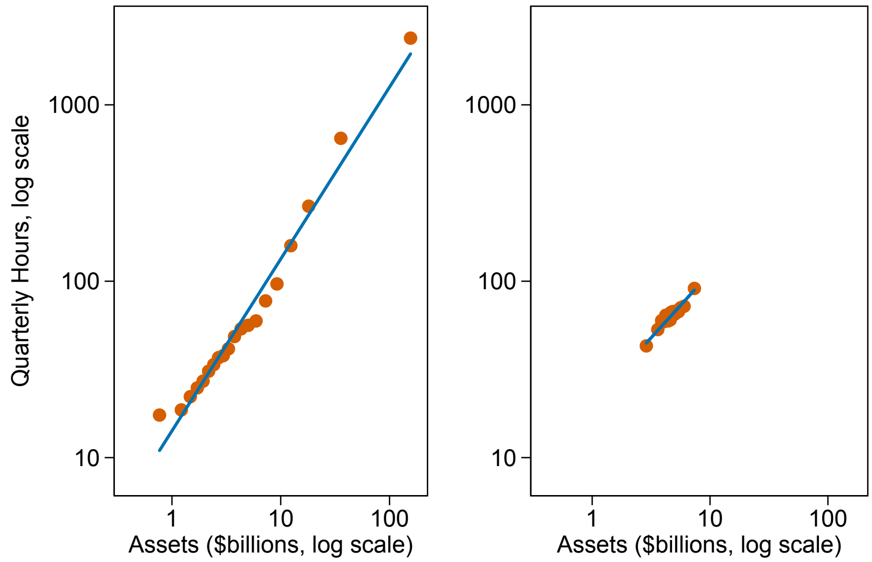

Figure 7: Relation between supervisory hours and assets. This figure presents a binned scatter plot and the fitted line of supervisory hours on BHC size obtained when controlling for rating, district and date-quarter dummies (column 1 of Table 2, left-panel) as well as bank fixed effects (column 2 of Table 2, right-panel).

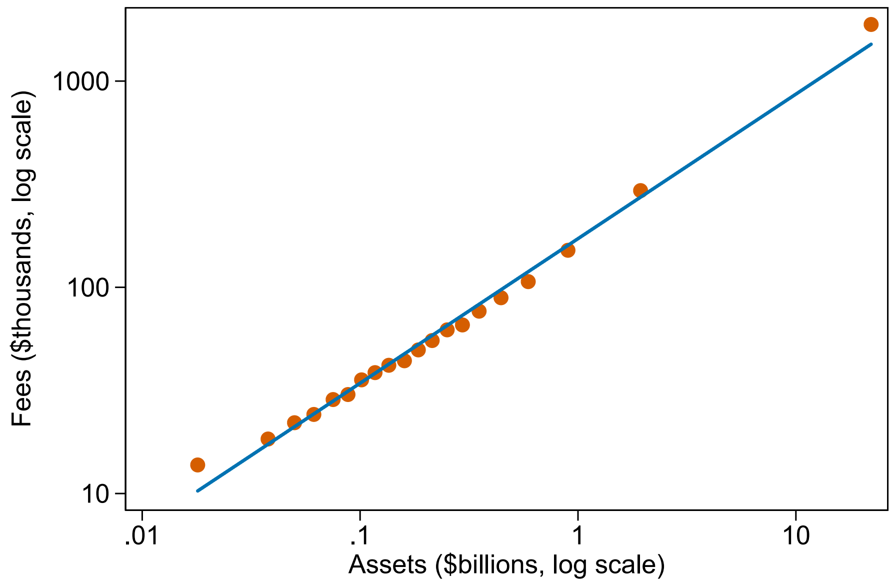

Figure 8: Relation between OCC semiannual fees and assets. This figure presents a binned scatter plot and the fitted line of OCC fees on commercial bank assets as computed in Table 5 on assets.

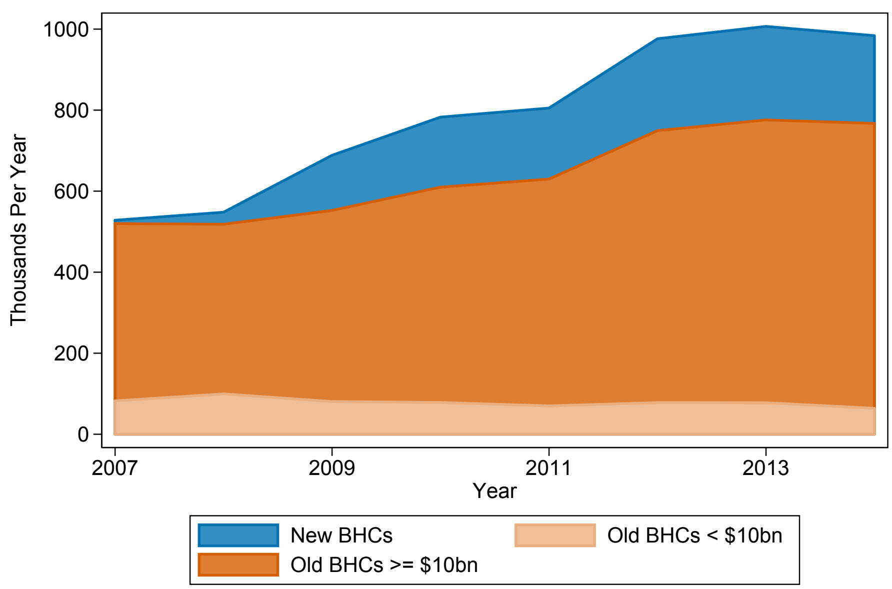

Figure 9: Total supervisory hours. This figure presents total hours of Fed supervisors matched to BHCs by whether a BHC existed as of 2007:Q1 and size.

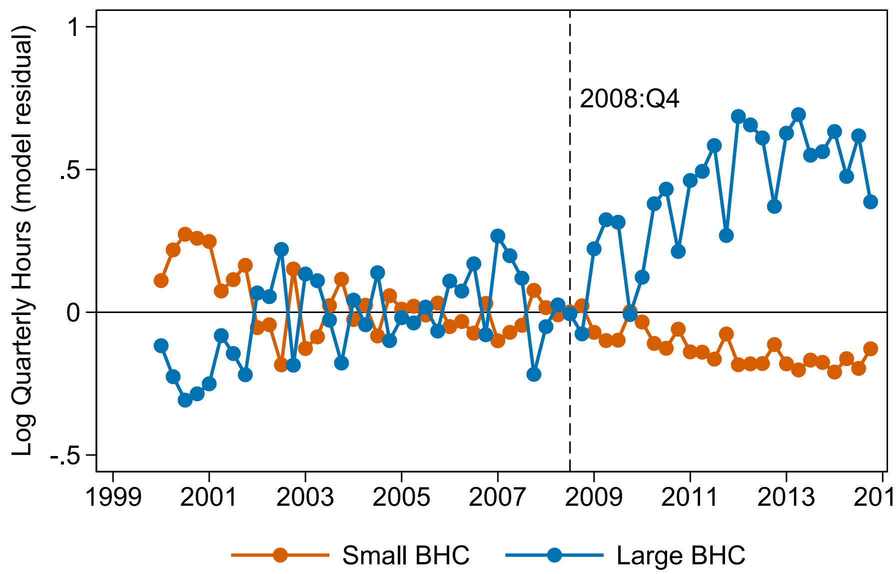

Figure 10: Average supervisory hour residuals over time. This figure shows the time series of residuals of the baseline regression (column 1, Table 2) of the log of supervisory hours on the log of assets and other controls. The residuals are averaged by date and whether the BHC is large or small on that date. Large (small) BHCs are defined in terms of assets above (below) a $10 billion asset threshold adjusted for the growth rate of industry assets.

Appendix

A Proof of Proposition 4.

Since $\overline{a} < \underline{a}$ and $\underline{c} < \overline{c}$ and F is decreasing in a, c for $x = \ell$, we have

$$\mathrm{PD}(\overline{a}, \underline{c}) = \max_{a,c} \mathrm{PD}(a, c).$$

Given $\overline{c} > \underline{c}$, we have

$$\mathrm{PD}(\underline{a}, \overline{c}) < \mathrm{PD}(\underline{a}, \underline{c}).$$

Next, consider the case of a perfectly informative signal. Then the first-order condition (3) yields

$$-\mathrm{PD}_c(\overline{a}, \overline{c}) N(A) = C_c(\overline{c}, A) \tag{14}$$

$$\text{and} \quad -\mathrm{PD}_c(\underline{a}, \underline{c}) N(A) = C_c(\underline{c}, A). \tag{15}$$

With $\overline{c} > \underline{c}$, we have $C_c(\overline{c}, A) > C_c(\underline{c}, A)$ so the left-hand sides of (14) and (15) imply

$$-\mathrm{PD}_c(\overline{a}, \overline{c}) > -\mathrm{PD}_c(\underline{a}, \underline{c}). \tag{16}$$

Since high c can compensate for high a, we have $F(x \,|\, a, c) = \hat{F}(x \,|\, a - c)$. Then (16) implies that $\overline{a} - \overline{c} > \underline{a} - \underline{c}$ and therefore

$$\mathrm{PD}(\overline{a}, \overline{c}) > \mathrm{PD}(\underline{a}, \underline{c}).$$

With a less than perfect signal, $\overline{c}$ is slightly lower and $\underline{c}$ slightly higher which strengthens the inequality. □

B Detailed variable definitions

Hours and Rating: See the discussion in Section 4.

Ongoing Full Exam: A dummy indicator for whether a full-scope exam is ongoing for a holding company in a given quarter. Source: NED confidential data.

Assets $\geq$ 10bn: An indicator for whether a bank has total assets greater than \$10 billion. The size threshold is deflated by the aggregate growth rate in BHC assets to account for growing size of BHCs. Asset data from FR-Y9C, item BHCK2170.

N Bank Employees: The number of employees. Employee data from FR-Y9C, item BHCK4150.

HHI Asset Concentration: Measure of business concentration using a Herfindahl-Hirschman Index (HHI) for asset concentration. See Kovner, Vickery and Zhou(2014). Data from FR-Y9C. Sum of squares of the following asset types (all as proportion of total assets, item BHCK2170):

- Total loans (BHCK2122)
- Total trading assets (BHCK3545)
- Fed funds and repo assets (pre-2002: BHCK1350; post 2002: BHDMB987 + BHCKB989)
- Investment securities (BHCK1754 + BHCK1773)
- The book value (not to exceed fair value), less accumulated depreciation, if any, of all real estate other than bank premises actually owned by the bank and its consolidated subsidiaries (pre-2000: BHCK2744 + BHCK2745; post-2000: BHCK2150)
- Premises and other fixed assets (BHCK2145)
- Investments in unconsolidated subsidiaries (BHCK2130)
- Direct and indirect investments in real estate ventures (BHCK3656)
- Cash (BHCK0081 + BHCK0395 + BHCK0397)
- All other assets (pre-2001: BHCK3164 + BHCKb026 + BHCK5507 + BHCK2160; post-2001 and pre-2006: BHCK0426 + BHCK2160 + BHCK2155; post-2006: BHCK0426 + BHCK2160)

Loan Share: Proportion of total BHC assets that are loans. Data from FR-Y9C, BHCK2122, BHCK2170

N Bank Subsidiaries: The count of all legal subsidiaries of a given BHC, computed using NIC data. This series was first contructed by Nicola Cetorelli and Samuel Stern at the NY Fed. More detailed is available in Cetorelli and Stern (2015).

BHC Failure: Dummy that indicates whether the BHC fails in the next year. A BHC fails when it terminates and the reason for its termination or the termination of a subsidiary within one quarter is failure. Data item: RSSD9061.

BHC Failure or Failing Rating: Dummy that indicates whether a BHC fails in the next year and/or has a rating of 4 or 5 in the next year. Ratings data from NED. See BHC failure for description of failing.

Standard deviation of ROA: Standard deviation of ROA from t+1 to t+4 (over the next year). ROA is calculated as 400 $\times$ net income / assets. Asset item is BHCK2170; net income is BHCK4340.