

NBER WORKING PAPER SERIES

INFORMATION PRODUCTION AND CAPITAL ALLOCATION:
DECENTRALIZED VS. HIERARCHICAL FIRMS

Jeremy C. Stein

Working Paper 7705
<http://www.nber.org/papers/w7705>

NATIONAL BUREAU OF ECONOMIC RESEARCH
1050 Massachusetts Avenue
Cambridge, MA 02138
May 2000

I am grateful to the National Science Foundation for research support, and to Ulf Axelson, Rob Gertner, Bengt Holmstrom, Anil Kashyap, Julio Rotemberg and seminar participants at the NBER for their input. The views expressed herein are those of the author and not necessarily those of the National Bureau of Economic Research.

© 2000 by Jeremy C. Stein. All rights reserved. Short sections of text not to exceed two paragraphs, may be quoted without explicit permission provided that full credit, including © notice, is given to the source.

Information Production and Capital Allocation: Decentralized vs. Hierarchical Firms
Jeremy C. Stein
NBER Working Paper No. 7705
May 2000
JEL No. G21, G31, L22

ABSTRACT

This paper assesses different organizational forms in terms of their ability to generate information about investment projects and allocate capital to these projects efficiently. A decentralized approach—with small, single-manager firms—is most likely to be attractive when information about individual projects is “soft” and cannot be credibly transmitted. Moreover, holding fixed firm size, soft information also favors flatter organizations with fewer layers of management. In contrast, large hierarchical firms with multiple layers of management are at a comparative advantage when information can be costlessly “hardened” and passed along within the hierarchy. As a concrete application of the theory, the paper discusses the consequences of consolidation in the banking industry. It has been documented that when large banks acquire small banks, there is a pronounced decline in lending to small businesses. To the extent that small-business lending relies heavily on soft information, this is exactly what the theory would lead one to expect.

Jeremy C. Stein
Harvard Business School
Morgan 361
Boston, MA 02163
MIT Sloan School of Management and NBER
jstein@hbs.edu

I. Introduction

In this paper, I take up the question of how an organization should be designed—in terms of size, number of managers, and the division of responsibility amongst these managers—so as to most efficiently allocate capital to competing investment projects. I contrast two basic organizational forms. The first is “decentralization”, in which small, single-manager firms choose between relatively few projects. The second is “hierarchy”, in which large firms with multiple layers of management evaluate many projects. The ultimate goal is to understand what project-level characteristics lead either decentralization or hierarchy to be the preferred design.

This question can be given a very concrete motivation. Over the last several years, there has been enormous consolidation in the banking industry, both in the U.S. and worldwide. This consolidation has been accompanied by widely-voiced concerns that the resulting larger banks will lend less to small businesses, who are particularly dependent on intermediaries for financing. And indeed, a number of researchers have documented that, when two banks merge, the resulting larger entity tends to cut back significantly on its small-business lending. Moreover, the evidence suggests that the loans that are cut are, at least on average, positive-NPV.¹

Why would a newly-enlarged bank ever willingly turn its back on a profitable existing line of business in this manner? The most common informal argument is nicely summarized by Berger, Demsetz and Strahan (1999):

“The larger institutions created by consolidation may also choose to provide fewer retail services to small customers because of Williamson (1967, 1988) type organizational

¹The literature on this topic is surveyed in Berger, Demsetz and Strahan (1999). I discuss this work in detail in Section IV below.

diseconomies....That is, it may be scope inefficient for one institution to produce outputs which may require implementation of quite different policies and procedures. These diseconomies may be most likely to arise in providing services to informationally opaque small businesses for whom intimate knowledge of the small business, its owner and its local market gained over time through a relationship with the financial institution is important...these arguments do not suggest that large complex financial institutions created by consolidation would reduce services to all small customers, rather just to those customers who rely on relationships.” (pp. 165-166.)

On the one hand, this informal argument is quite clear in asserting that there exist “organizational diseconomies” which somehow prevent big banks from being the most efficient providers of certain information-intensive services, such as relationship-based small business lending. On the other hand, it is vaguer as to what the root cause of these diseconomies might be. For example, the suggestion that big banks simply have trouble engaging in multiple activities that require different technologies (“different policies and procedures”) seems less than compelling. After all, most big banks are involved in a wide range of technologically distinct activities, from check processing to credit cards to foreign-exchange trading.

So what is it about small-business lending—as opposed to various other banking activities—that might lead it to be an especially poor fit for a large banking firm? In what follows, I argue that the key distinguishing characteristic of small-business lending in this regard is that it relies heavily on information that is “soft”—i.e., information that cannot be directly verified by anyone other than the agent who produces it. For example, a loan officer who has worked with a small-company president may come to believe that the president is honest, prudent, and hardworking—in other words, the classic candidate for an unsecured “character loan”. Unfortunately, none of these attributes can be unambiguously documented in a report that the loan officer can pass on to his superiors. This situation contrasts sharply with, e.g., an application for a home mortgage loan. Here the decision of whether or not to extend credit is

likely to be made primarily based on “hard”, verifiable information, such as the income shown on the borrower’s last several tax returns.²

The model that is developed below has the feature that although decentralization necessarily has some disadvantages, it may on balance be a good method for allocating capital when the projects in question must—like small-business loans—be evaluated based on soft information. The flip side is that large, hierarchical firms have a comparative advantage in capital allocation when information about projects can be hardened, as in the case of mortgage lending.

With soft information, the fundamental advantage of decentralization is that it strengthens the research incentives of line managers. Imagine a loan officer in Little Rock who is responsible for deciding which small-business loans are worth making. The quality of the loan officer’s judgement will depend on how good a job he has done in terms of information production, which in turn will be a function of his incentives. Under full decentralization, the Little Rock loan officer is also the CEO of the local bank, and as such has the authority to allocate the bank’s funds as he sees fit. Given that he can count on having some capital to work with, he knows that his research efforts will not be wasted, and hence his incentives to do research are relatively strong. Said differently, decentralization rewards an agent who develops expertise by ensuring that he will also have access to some capital which he can use to lever that expertise.

In contrast, if the Little Rock loan officer is part of a large multi-bank hierarchy, the

²It is worth noting, however, that this divergence between mortgages and small-business loans may be shrinking over time. Improvements in technology, the accompanying growth of “infomediaries” (such as Dun and Bradstreet), and more widespread adoption of credit-scoring models would appear to make it possible for an increasing amount of information about small businesses to be hardened. See, e.g., Petersen and Rajan (2000) for a recent discussion.

following problem arises. Suppose that he spends a lot of effort learning about prospects in his area. But then somebody higher up in the organization decides that overall lending opportunities are better in Tulsa, and sharply cuts the capital allocation for Little Rock. In this case, because he doesn't get a chance to act on the information that he has produced (and because he is unable to credibly pass it on) the Little Rock loan officer's research effort goes to waste. Ex ante, this implies that the loan officer does less research in a hierarchical setting. Here the authority to allocate capital is separated from expertise—i.e., the Little Rock loan officer may be left with no capital to work with—which tends to dilute the incentives to become an expert.

So with soft information, the advantage of decentralization relative to hierarchy is higher-powered research incentives and hence better capital allocation within operating units (i.e., within the Little Rock office). Of course, there is also a countervailing cost. By committing to a fixed capital allocation for the Little Rock office, decentralization leads to inefficiencies when lending opportunities are unexpectedly weak in Little Rock and strong in Tulsa. In other words, decentralization does not allow for efficient reallocations across operating units, while a hierarchical design does. Nevertheless, if line-manager research is sufficiently valuable, decentralization may on net be the better design when information is soft.

However, things work very differently when the information produced by line managers can be hardened and passed on to their superiors. Now, not only does a hierarchy do better in terms of moving money across operating units, it can also generate more research on the part of line managers than under decentralization. This is because with hard information, these managers effectively become advocates for their units—if they can produce verifiable positive information and pass it on to their superiors, they can increase the amount of capital that they are

allocated. Here, paradoxically, separating authority from expertise actually improves research incentives, as line managers struggle to produce enough information to convince their bosses that they should get a larger share of the firm's overall capital budget.

Beyond just saying that soft information favors small firms, the model also produces several other conclusions. First, suppose that for some other exogenous reason, it becomes optimal to have a relatively large integrated firm operating in a setting where information is soft—say because there are significant synergies across the firm's different projects. In such a case, holding the firm's size and scope fixed, the softness of information will tend to imply that a flatter organizational structure, with fewer layers of management, is more attractive.

Another implication is that hierarchies tend to be characterized by inefficient levels of bureaucracy. This implication follows if one extends the model so that the hardness of information is not exogenously fixed. For example, one might assume that by devoting effort to documentation and report-writing, a line manager can harden information that would otherwise be soft. Because hard information is so (privately) valuable to line managers in a hierarchy, they will devote excessive efforts to such documentation activities. Thus in this modified setting, the costs of a hierarchy do not necessarily take the form of line managers simply being discouraged and slacking off, but instead may manifest themselves as line managers working very hard to generate the wrong kind of information. In particular, there will be too much report-writing and not enough soft-information production.

The ideas in this paper build on several earlier works. I defer a full discussion of the related literature until Section V, and only note here the most direct linkages. The capital-allocation model itself is a direct extension of that in Stein (1997), with the new twist being the

explicit consideration of line-manager research incentives. In Stein (1997), line managers are treated as passive robots who do not need to be motivated. Rather, the focus is more narrowly on the CEO's incentives to shift resources across the firm's different operating units—an approach that naturally leads to a more favorable view of large organizations.

In arguing that line managers' incentives may be blunted when they are in a hierarchy and thus do not have ultimate authority, I am following Aghion and Tirole (1997). However, a critical distinction is that in the model of this paper, a hierarchical structure need not weaken line-manager incentives—indeed, it only does so when information is soft. In contrast, in Aghion and Tirole (1997), it is a more general proposition that line managers are discouraged when they do not have full authority. Thus the models have quite different empirical implications, as the example of the banking industry suggests: the Aghion-Tirole model does not say anything about why large hierarchies might be at more of a disadvantage with small-business loans than with credit cards or mortgages.

The remainder of the paper proceeds as follows. The basic model is developed in Section II. Section III considers several extensions and variations. In Section IV, I return to the banking industry, and review the relevant empirical evidence more fully in light of the theory. Section V discusses the related theoretical literature on organizational design. Section VI concludes.

II. The Model

A. Basic Structure

The model considers a firm with two divisions, i and j . Within each division, there are two potential investment projects. The projects in division i are denoted by $i1$ and $i2$

respectively, while those in division j are denoted by $j1$ and $j2$. Each of the four projects can be operated at a variable scale; in particular, each project can be allocated either zero, one or two units of capital. The firm faces an overall capital constraint, in that it only has a total of four units of capital that can be distributed.

The capital constraint can be interpreted in one of two ways. First, one can think of “capital” as simply cash that can be spent on the projects. In this interpretation, the firm faces an overall liquidity constraint. Although the roots of this liquidity constraint are not modeled in detail here, they can be easily endogenized, as is done by Stein (1997) in a very similar setting; I will elaborate on this point briefly below. Alternatively, the “capital” might represent, e.g., access to specialized production facilities. That is, the firm might have a fixed amount of lab space—which, due to adjustment costs, cannot be increased in the short run—and it must decide how to allocate this lab space across projects.

Each of the four projects can be in either a G (good) or B (bad) state. The probability of each state is $\frac{1}{2}$, and the outcomes are independent across projects. A project that is in the G state yields a net output of $g(1)$ if it gets one unit of capital, and $g(2)$ if it gets two units. Similarly, a project that is in the B state yields a net output of $b(1)$ with one unit of capital, and $b(2)$ with two units. It is assumed throughout that $g(2) < 2g(1)$, and $b(2) < 2b(1)$; that is, there are decreasing returns in either state. In addition, $g(2) > g(1) + b(1)$, which means that if one has two units of capital along with a G project and a B project, it is better to give the G project both units, as opposed to dividing the capital up equally across the projects. In other words, it makes sense to move capital from bad to good projects.

Each of the two divisions has its own division manager. The division managers are, by

virtue of research effort, able to learn about the signals of the projects that they oversee.

Specifically, if manager i makes an effort e_i , he has a probability $p(e_i)$ of observing the signals on both of his projects, $i1$ and $i2$, where the function $p(\cdot)$ is increasing, concave, and takes on values on the interval $[0, 1)$. I assume that the division managers have reservation utilities of zero, so that there is never any issue of satisfying their participation constraints.

The firm also has a CEO, who may in some circumstances undertake her own research. The CEO's research technology is described in detail below. For the time being it suffices to say that, because the CEO is overseeing a total of four projects, she is unable to learn as much about each one individually as are the division managers. Instead, the CEO is at best able to get noisy information about the aggregate prospects of each of the two divisions.

To model the incentives of the divisions managers and the CEO, I follow Stein (1997), and assume that each agent seeks to maximize the gross output from the assets under his or her control. This assumption can in turn be motivated based on: i) private benefits of control that are proportional to gross output; and ii) non-responsiveness of agents to monetary incentives.³ It has the following behavioral implication: each agent always prefers more capital to less, but conditional on being granted a certain amount of capital, each agent tries to allocate it efficiently. Or said differently, the agents in the model are empire builders, but holding the size of their empires fixed, they prefer them to be profitable.

Note that division managers' empire-building preferences create an agency problem

³Non-responsiveness to monetary incentives is a common—albeit extreme—modeling device in papers in this genre. One way to generate it is to assume that agents are infinitely averse to risk in their monetary income (though not necessarily averse to variations in non-monetary private benefits). See Aghion and Tirole (1997) for a discussion.

between them and the CEO. For example, if manager i has better information about projects $i1$ and $i2$ than the CEO, he will, if asked, always say that both projects are in the G state. As a result, the CEO cannot trust the division managers to give accurate reports about the prospects of their projects.

A similar agency problem exists between the CEO and the firm's shareholders. However, the effects of this problem are suppressed in the equilibria that I consider below. This is because I assume throughout that the amount of capital that the CEO has to work with is fixed at four units, so her own empire-building preferences do not come into play. Rather, given the capital constraint, her efforts to maximize her private benefits coincide perfectly with shareholders' objective of maximizing the firm's net profits.

Of course, the fact that the CEO is capital-constrained in the first place may be a result of shareholders trying to check her empire-building tendencies. This is precisely how the capital constraint is endogenously derived in Stein (1997), and a similar argument can be made here. Shareholders can never rely on the CEO to honestly reveal anything she might know about project prospects, since given her taste for more capital, she will always make an optimistic report. Thus the best that shareholders can do is simply to give the firm a fixed, uncontingent allocation of capital. By choosing parameter values appropriately, it is easy to set things up so that the optimal ex ante allocation from shareholders' perspective is four units.⁴

⁴If one makes $b(2)$ small enough (note that $b(2)$ can be negative, since it is a net return), shareholders will always prefer to give the firm four units of capital rather than five or more. This is because if the firm has five or more units, there will necessarily be some states of the world when two units are invested in a B project, yielding $b(2)$.

B. The First-Best Benchmark

As a benchmark for comparison, it is useful to compute the value of expected net output that is obtained in a idealized situation where a single planner has access to perfect information about all the projects. This computation is straightforward: simply enumerate each of the sixteen possible outcomes (four projects, each of which can be in one of two states) and, any time there is a G project that can be paired with a B project, have the planner invest two units in the former and none in the latter. The resulting net output, denoted by Y^{fb} , is given by:

$$Y^{fb} = (5g(2) + 3b(1) + 3g(1))/4 \quad (1)$$

C. Decentralization

The first—and simplest—organizational form to be considered is “decentralization”. Under decentralization, the CEO does no research of her own and does not attempt to get involved at all in the capital allocation process. Rather, there is simply an unconditional, ex ante commitment to give each division manager two of the four units of capital to work with. Manager i can then allocate his two units however he sees fit across projects $i1$ and $i2$, and similarly for manager j .

Given how spare the model is on other dimensions—with no operating synergies across any of the projects—decentralization is most naturally interpreted as dis-integration. That is, decentralization corresponds to the divisions being split up into two separate firms run by managers i and j , with the CEO out of the picture. Indeed, Baker, Gibbons and Murphy (1999) argue that a breakup of this sort is the only way to credibly commit that the CEO will not get

involved in the capital allocation process.⁵

The costs of decentralization are easy to see. Imagine that the signals for projects $i1$ and $i2$ are $\{G, G\}$ while those for projects $j1$ and $j2$ are $\{B, B\}$. In this case, the optimal thing to do would be to give all four units of capital to division i , but this cannot happen under decentralization. In other words, decentralization precludes making value-enhancing transfers across divisions. This can be thought of as a failure to effectively coordinate activities across divisions.

On the other hand, decentralization may allow for relatively efficient allocations within divisions. For if a division manager's research is successful and he observes $\{G, B\}$ for the two projects under his purview, he will do the right thing ex post and give both units of funding to the G project. The only question is how much effort he will put into this research ex ante.

To answer this question, note that if a division manager knows the states and can make allocations conditional on this knowledge, expected net output is $(g(2) + b(1) + g(1))/2$. If, on the other hand, the manager is uninformed, each project always gets one unit of funding, and expected net output is simply $g(1) + b(1)$. Thus the expected output gain to being informed under decentralization, denoted by Δ^d , is given by:

$$\Delta^d = (g(2) - g(1) - b(1))/2 \quad (2)$$

Assuming that the manager puts a weight of γ on expected net output relative to effort

⁵Aghion and Tirole (1997) adopt a different point of view, assuming that formal authority—including the right to make capital allocation decisions—can be irrevocably delegated to managers other than the CEO. This is an important distinction, and I will return to it later.

(i.e., his private benefits are a fraction γ of gross output and investment is exogenously fixed), his first-order conditions imply that the level of research effort under decentralization, e^d , satisfies:

$$p'(e^d) = 1/\gamma\Delta^d \quad (3)$$

With two independent divisions each behaving this way, the total net return (on the four units of capital) under decentralization, Y^d , is given by:

$$Y^d = p(e^d)(g(2) + b(1) + g(1)) + (1 - p(e^d))(2g(1) + 2b(1)) \quad (4)$$

From equation (4) it can be seen that decentralization involves two inefficiencies relative to the first-best outcome in equation (1). First, the fact that research is costly implies that there will not always be full information about project prospects: $p(e^d) < 1$. Second, even if we approach the limit where $p(e^d) = 1$, Y^d still only reaches $(g(2) + b(1) + g(1))$, which is less than Y^{fb} . This is because under decentralization, resources cannot be moved across divisions, only within divisions.

D. Hierarchy: the Case of Soft Information

The next case to be considered is one in which there is soft information and the firm is organized as an integrated hierarchy, with the two division managers ceding formal authority to the CEO. To make this case interesting, one needs to assume that the CEO can gather some information on her own. If not, she can do no better than to always grant each division two units

of funding, thereby reproducing the decentralized outcome. In other words, with soft information, a CEO who does no research of her own may have formal authority, but (in the words of Aghion and Tirole (1997)) she has no real authority.

While it is plausible that the CEO can learn something about investment prospects, it would be unreasonable to posit that she can learn as much in total as the two division managers. Instead, I assume that the CEO can only get coarse information about the aggregated prospects of each division. Specifically, there is a probability q that the CEO's research efforts will be successful. Successful research means that, if one or both divisions are "stars"—in the sense of having both of their projects in the G state simultaneously—this star status will be revealed to the CEO. The coarseness of the CEO's research technology is captured in the fact that, even if her research is successful, she can never differentiate between a division that is "average" (has one G and one B project) and a division that is a "dog" (has two B projects).

A couple of points about this formulation deserve comment. First, the CEO's research-success probability q is for the time being an exogenous parameter. In Section III, I will discuss what happens when q is made an endogenous function of the CEO's effort. Second, the exact way that I have modeled the coarseness of the CEO's information is not critical. I could equivalently assume that successful research allows the CEO only to identify dog divisions, and that she can never distinguish between average divisions and stars; this leads to the same results.

Capital allocation in a hierarchy works as follows. When the CEO's research is unsuccessful (which happens with probability $(1-q)$), the best she can do is to just give each

division manager two units of funding, and we are back to the decentralized outcome.⁶ When the CEO's research is successful (which happens with probability q), she may choose to deviate from equal funding, and give one division more than the other. This will only happen if one division is identified as a star, and the other is not. In such a "lone-star" scenario, the CEO has three options: 1) continue to give each division two units; 2) give the star division three units and the other division one unit; or 3) give the star division all four units. It is easy to show that the CEO will choose the most extreme tilting of the capital budget—giving all four units of funding to the star division—if the following sufficient condition is met:

$$g(2)/2 > (5g(1) + b(1))/6 \quad (5)$$

This condition simply requires that decreasing returns to scale are not too pronounced in the good state. The analysis that follows is most transparent when this star-gets-everything condition is satisfied, so I will begin by assuming that it is. Later, I will come back to the case where returns decrease more sharply with scale, so that the CEO gives just three, rather than four units of funding to a division that is identified as a lone star.

If division managers' ex ante research incentives in a hierarchy were the same as under decentralization, it would follow immediately that hierarchy is the strictly dominant

⁶I do not consider the possibility that the CEO might use her information as part of a mechanism to induce division managers to reveal their own signals. One could imagine that if a manager lies about his division being a star, and the CEO subsequently catches him, he would be punished with a reduced capital allocation, as in Harris and Raviv (1996). In my set-up, such a scheme suffers from two distinct commitment problems. First, punishment involves ex-post inefficient capital allocations. And second, if in equilibrium the division managers do truthfully reveal their signals, the CEO will no longer have any incentive to do her own research.

organizational form. This is because ex post, hierarchy allows for a form of selective intervention. When the CEO knows nothing about divisional prospects, she does not interfere, and the outcome is the same as with decentralization. When the CEO does know something, her ability to shift funds towards a star division leads to an improved across-division allocation.

The problem, however, is that division managers' ex ante research incentives are weaker in a hierarchy when information is soft. To see why, suppose that the CEO's research has been successful, and that she has identified division j as a star. Division i , meanwhile, has one G and one B project. In a hierarchy, division j gets all four units of funding, and division i gets nothing. Hence any information that manager i has acquired is not put to use. In contrast, if the divisions were decentralized, and manager i had two units of funding to work with, he would find his information valuable—it would lead him to shift both units to his single G project. Thus the downside to a hierarchy is that because the CEO sometimes takes away manager i 's capital budget, the marginal return to his research efforts is reduced, and he produces less information.⁷

It is important to recognize that the negative incentive effects of hierarchy arise not simply because the CEO sometimes has her own independent information about divisional investment opportunities. It is also crucial to the argument that the CEO have the authority to take away all funding from division i —even though i might be able to raise two units if it were a stand-alone entity—when her research indicates that division j is a star.⁸ As emphasized by Stein

⁷Note that the negative incentive effect that comes from sometimes losing two units of funding is not offset by the fact that the division manager also sometimes gets two extra units (four units in total) of funding in a hierarchy. When he receives four units, the only thing he can do is invest two in each of his projects, and his information is again not of any value.

⁸Indeed, the discouragement effect occurs precisely because an “average” division (i.e., one with one G and one B project) can lose all its funding based not on its own prospects, but

(1997), it is this authority that distinguishes a CEO from, e.g., a well-informed banker. Thus what I am calling a “hierarchy” cannot be equivalently thought of as two independent firms facing a single well-informed external supplier of capital.

To formalize things, denote by Δ^{hs} the gain in expected net output that arises when, in a hierarchy with soft information, a division manager’s research efforts are successful. It is straightforward to show that:

$$\Delta^{hs} = (1 - q)(g(2) - g(1) - b(1))/2 + 3q(g(2) - g(1) - b(1))/8 = (1 - q)\Delta^d + 3q\Delta^d/4 \quad (6)$$

The level of research effort in a hierarchy with soft information, e^{hs} , satisfies:

$$p'(e^{hs}) = 1/\gamma\Delta^{hs} \quad (7)$$

Since $\Delta^{hs} < \Delta^d$, it follows that $e^{hs} < e^d$. By working through all the possible outcomes, it can then be established that expected net output in a hierarchy with soft information is given by:

$$Y^{hs} = (1 - q)\{p(e^{hs})(g(2) + b(1) + g(1)) + (1 - p(e^{hs}))(2g(1) + 2b(1))\} + q\{p(e^{hs})(6g(2) + 3b(1) + g(1))/4 + (1 - p(e^{hs}))(3g(2) + 6b(1) + 4g(1))/4\} \quad (8)$$

By comparing equations (8) and (4), one can evaluate the relative efficiency of decentralization versus hierarchy. The results to this point can be summarized as:

rather because the prospects of its counterpart division within the firm are so strong.

Proposition 1: Assume that the condition in (5) holds, $g(2)/2 > (5g(1) + b(1))/6$, so that in a hierarchy, a lone-star division gets four units of funding. Then decentralization always leads to more research effort than hierarchy: $e^d > e^{hs}$. In addition, it is possible (though not necessary) that decentralization leads to higher expected output than hierarchy; i.e., that $Y^d > Y^{hs}$.

To see why decentralization can generate higher expected output, consider a simple limiting case where $q = 1$, $p(e^d) = 1$, and $p(e^{hs}) = 0$. (The latter two conditions can always be generated by picking the proper form for the $p(\cdot)$ function.) In this case, equation (4) simplifies to $Y^d = g(2) + b(1) + g(1)$, while equation (8) simplifies to $Y^{hs} = (3g(2) + 6b(1) + 4g(1))/4$, implying that $Y^d > Y^{hs}$. What is going on here is that when the firm is organized as a hierarchy, only the CEO does any research, and the division managers are totally discouraged. Conversely, under decentralization, the division managers are highly motivated and become perfectly informed. Given that the two division managers taken together are able to gather more accurate information than the CEO, this latter effect is more than enough to outweigh any improved across-division allocation that can be obtained in a hierarchy. As a result, decentralization is the better mode of organization.

Of course, this example relies on $p(e^d)$ and $p(e^{hs})$ being relatively far apart—i.e., on division-manager effort being both important, and responsive to incentives. If $p(e^d)$ and $p(e^{hs})$ are sufficiently close to one another, it is easy to see that hierarchy becomes more efficient than decentralization: $Y^{hs} > Y^d$.

Although I have, for the sake of expositional simplicity, derived the results in Proposition 1 under the assumption that the star-gets-everything condition in (5) holds, they are in fact more

general. Even when returns to scale are more sharply decreasing in the good state—so that a lone-star division gets three, rather than four units of funding—the basic intuition can carry over. That is, when manager i generates information and gets allocated only one unit of financing, (because manager j has been deemed a star by the CEO) manager i 's information is not worthless, but it may still be less valuable than if he had been allocated two units of financing. Thus the threat of losing some—if not all—of their funding can continue to exert a chilling effect on division managers' research incentives.

To make this idea precise, in the appendix I prove the following:

Proposition 2: Assume that $g(2)/2 > (3g(1) + b(1))/4$. In this case, the CEO in a hierarchy will still tilt the capital budget toward a lone-star division, but the tilt may be less extreme, with the lone star receiving three, rather than four units of funding. Nevertheless, decentralization continues to lead to more research effort than hierarchy: $e^d > e^{hs}$. In addition, it is possible (though not necessary) that decentralization leads to higher expected output than hierarchy; i.e., that $Y^d > Y^{hs}$.

E. Hierarchy: the Case of Hard Information

The relative merits of a hierarchy increase when the information generated by division managers can be hardened and passed along to the CEO. To introduce hard information, I begin by assuming that the division managers have the same research technology as before—that is, if manager i makes an effort e_i , he has a probability $p(e_i)$ of observing the signals on both of his projects, $i1$ and $i2$. Now, however, the CEO does no separate research of her own. Instead, if a

division manager learns something, there is a chance that he may be able to credibly communicate it to the CEO. Specifically, conditional on a division manager's research being successful and yielding information about his two projects, there is a probability z that this information is verifiable, and can be shown directly to the CEO. With probability $(1 - z)$, the information is non-verifiable, and can be used by the division manager, but not credibly transmitted to the CEO.

None of the results that follow depend on the particular value of z —anything in the interval $(0, 1]$ works fine. Thus one could in principle economize on notation by setting $z = 1$. However, I allow for $z < 1$ because it leads to a richer and more natural interpretation of the role of the division managers. When $z < 1$, there continue to be times when the information remains soft and a division manager knows more than he can ever communicate to the CEO. Thus the role of the division manager goes beyond simply making reports to the CEO and then passively awaiting his capital allocation—there are times when he can add further value by choosing how to best spend his allocation.

Note that I am implicitly assuming that while hard information can be credibly transmitted from division managers to the CEO, it cannot be transmitted to outside shareholders. For if such information could be shown directly to outside shareholders, it would no longer make sense to specify that the firm faces a fixed, uncontingent liquidity constraint. Rather, in cases where a lot of positive information could be shown to shareholders, they might be willing to invest more than four units in the firm.⁹

⁹This issue does not arise if one takes the alternative interpretation that the “capital constraint” represents a shortage of some specialized physical asset (e.g., production facilities), rather than an inability to raise more cash from outside investors.

What does it mean to assume that there is a type of information that can be transmitted to the CEO but not to shareholders? Think of division managers as providing the CEO with a variety of raw, albeit well-documented data about a project. For example, if the project involves drilling for oil in a new location, the raw data might be a set of geological studies. This raw data does not literally say what the project's dollar payoff will be. Instead, the raw data must be combined with the CEO's effort and expertise (e.g., her knowledge of geology, her assessment of the costs of doing the drilling and extraction, etc.) to generate a final judgement about dollar value. The CEO's final judgement—what she concludes after studying the raw data—is itself soft information that cannot be credibly communicated to non-expert outside investors.¹⁰

Another crucial assumption is that the CEO does not observe the level of division-manager research effort, nor can she tell whether or not this research has been successful.¹¹ Consequently, when a division manager has hard information, he can choose whether or not to report it to the CEO. In particular, a division manager who gets hard information that his projects are $\{B, B\}$ may opt to simply keep quiet. Importantly, this will not lead to an “unraveling” situation where the CEO can infer that the state must be $\{B, B\}$ simply because the division manager is quiet. This is because the division manager may also be quiet as a result of not having obtained any hard information in the first place.

¹⁰Under this interpretation, the parameter z can be thought of as reflecting the ability or effort of the CEO—a higher value of z means that the CEO can better evaluate the raw data produced by the division managers. As will be made clear below, when information is hard, CEO and division-manager effort become strategic complements, as opposed to substitutes.

¹¹Of course, as in any agency model, I have been assuming all along that division-manager research effort is non-contractible—i.e., cannot be verified in court. However, it is stronger to assume that it cannot even be observed by the CEO.

The option to keep quiet when hard information is negative makes research very attractive to the division manager. To see why, assume that the CEO has a fixed conjecture e_i^c about the level of effort that manager i exerts, and a corresponding conjecture $p^c = p(e_i^c)$ about the probability that his research is successful. Now suppose that the division manager's reporting strategy is to reveal his hard information to the CEO when it is either $\{G, G\}$ or $\{G, B\}$, but to keep quiet when it is $\{B, B\}$. (As will be seen below, this reporting strategy constitutes equilibrium behavior.) Bayes' rule implies that, given his conjecture of p^c , the CEO will interpret quiet as follows:

$$\text{prob}(\{B, B\}/\text{quiet}) = 1/(4 - 3zp^c) \quad (9)$$

$$\text{prob}(\{G, G\}/\text{quiet}) = (1 - zp^c)/(4 - 3zp^c) \quad (10)$$

$$\text{prob}(\{G, B\}/\text{quiet}) = 2(1 - zp^c)/(4 - 3zp^c) \quad (11)$$

Taking the CEO's conjecture of p^c as fixed, a division manager frames the problem as follows. If he does no research, he will certainly be quiet, and the CEO will update on him according to equations (9)-(11). However, if he does devote some effort to research, he gains pure option value. If the research produces hard information that his division is either $\{G, G\}$ or $\{G, B\}$, he can speak up and thereby impress the CEO. If the research produces hard information that his division is $\{B, B\}$, he just keeps quiet and is no worse off than if he had done no research.

The bottom line is that, from the perspective of the division manager, there is now an added benefit to doing research: if the information he generates is hard and positive, it can help increase his capital budget and hence his private benefits. This contrasts with both of the

previous scenarios—decentralization and hierarchy with soft information—where division managers' research efforts had no impact on the capital they were allocated.

In the appendix, I provide a detailed characterization of the case when information can be hardened. The key results can be summarized as follows:

Proposition 3: Assume that $g(2)/2 > (3g(1) + b(1))/4$. Then in a hierarchy with hard information, there is an equilibrium with the following properties:

i) Division manager reporting strategies: Division managers reveal their hard information if it is either $\{G, G\}$ or $\{G, B\}$, and keep quiet if it is $\{B, B\}$.

ii) CEO capital-allocation policy: If the CEO is facing one division that reveals itself to be $\{G, G\}$ and one that is quiet, the $\{G, G\}$ division gets at least three (and possibly four) units of capital. In all other cases, the CEO allocates each division two units of capital.

iii) Division manager benefit from being informed: Denote by Δ^{hh} the expected gross output gain to a division manager from being informed in a hierarchy with hard information.

This output gain significantly exceeds that under decentralization: $\Delta^{hh} > \Delta^d + z/4$.

Consequently, division-manager research effort is greater than under decentralization: $e^{hh} > e^d$.

iv) Output: Expected net output is greater than under decentralization: $Y^{hh} > Y^d$.

The intuition behind parts (i) and (ii) of the proposition has already been discussed.¹² Part

¹²There are other equilibria that differ insignificantly from that described in Proposition 3. When one division reveals $\{G, G\}$ and the other reveals $\{G, B\}$, the CEO is actually indifferent between allocating two units to each, or three to the former and one to the latter. I focus on the two-two allocation in the proposition (which implies that a manager with a $\{G, B\}$ signal strictly prefers to reveal it). However, parts (iii) and (iv) of the proposition apply in either case.

(iii) gives a quantitative sense for just how much stronger division managers' research incentives are when information can be hardened. The expected output gain to being informed in a hierarchy, Δ^{hh} , is now not just larger than that under decentralization Δ^d ; it exceeds it by at least $z/4$. This is a very substantial difference, since Δ^d is effectively denominated in units of net return differentials, while $z/4$ is in units of gross capital. To take a concrete example, reasonable values of $b(1)$, $g(1)$ and $g(2)$ might be 0.05 , 0.10 and 0.19 respectively. In other words, one dollar invested in a B project yields a net present value (after all discounting) of five cents, etc. For these values, equation (2) says that $\Delta^d = 0.02$. But then if z is anywhere near one, then $z/4$ is on the order of ten times the size of Δ^d . Even for much smaller values of z , it will still be the case that $z/4$ is much larger than Δ^d .

This all just reflects the fact that in a hierarchy with hard information, research effort can actually influence a division manager's gross capital budget. In contrast, under decentralization (or in a hierarchy with soft information) research effort only enables a division manager to get a higher return from a given capital budget. To the extent that division managers' private benefits are proportional to gross output, the former effect is naturally much stronger than the latter.

Part (iv) of the proposition—regarding expected net output—follows immediately from part (iii). Now a hierarchy does better than decentralization on both dimensions of importance. Not only does it generate more information at the division-manager level, and thereby lead to better within-division allocations, but it also allows for reallocations across divisions when the CEO learns that such reallocations are value-increasing.

One might argue that the results in Proposition 3 in favor of hierarchy relative to decentralization are now “too strong”—once information can be hardened, hierarchy

unambiguously dominates decentralization for all parameter values, which seems unrealistic. Certainly the model omits a number of other factors that might tip the balance back towards decentralization. For example, other than research effort, the model assumes that division managers do not need to take any other actions. If one were to introduce another dimension of non-contractible firm-specific investment, this might (following the logic of Grossman and Hart (1986), Hart and Moore (1990), and Hart (1995)) be expected to make hierarchy less attractive.

But it is important to emphasize that I am much less interested in making absolute unconditional statements about the virtues of hierarchy compared to decentralization, and more interested in making comparative-statics statements about the circumstances under which hierarchy is more likely to be attractive. And the key result in this regard—that hierarchy looks better when information is hard, as opposed to soft—seems like it should be robust to the inclusion of various other factors into the model.

III. Further Issues

A. Does Decentralization Require Dis-Integration?

In the preceding discussion, I have been treating the concept of decentralization as loosely synonymous with dis-integration, i.e., with the two divisions being spun off as independent stand-alone entities. Given how minimalist the model is on other dimensions—with no operating synergies across the two divisions—this interpretation seems uncontroversial. However, a harder question arises if there are significant synergies, which would be lost in a breakup. Then one must ask whether it is possible to capture the benefits of decentralization in capital allocation—i.e., to commit to division managers that they will always receive two units of

funding—without a costly breakup.

Different authors have adopted different perspectives on this question. Aghion and Tirole (1997) assume that formal authority over certain decision rights can be contractually delegated to managers in the organization other than the CEO. In the context of the model here, this assumption would correspond to the idea that the CEO can indeed be contractually prohibited from getting involved in the capital allocation process. Clearly, if this is the case, and if such delegation does not interfere with the CEO's ability to carry out other aspects of her job, the benefits of decentralization can be obtained even in the context of a large integrated firm.

Baker, Gibbons and Murphy (1999) take the opposite position, arguing that formal authority over all decisions necessarily resides at the top of an organization. In other words, a CEO cannot irrevocably contract away her right to get involved in any decision. If this is a better description of reality, then there are limits to how far one can go towards the decentralized outcome without a breakup. Certainly, there are other devices short of a breakup that one can employ to reduce the CEO's ex post incentive to get involved; in this vein, Aghion and Tirole (1997) suggest putting the CEO in a situation of "overload", so that she is too busy to do much research on her own. And Baker, Gibbons and Murphy (1999) emphasize the role of reputation—in a repeated-game setting, the CEO may try to develop a reputation for being non-interventionist. But while these devices may be able to improve division managers' incentives to a degree, they will not in general have as strong an effect as a binding commitment that prevents the CEO from altering the capital budget under any circumstances.

A middle-ground view might be that even if Aghion and Tirole (1997) are right, and it is possible for the CEO to formally divest some types of decision rights, (e.g., the right to fire

certain workers without establishing just cause) it is unlikely that she can cede the specific right that is of interest here, the right to make high-level capital allocation decisions.¹³ Indeed, others (Williamson (1975), Donaldson (1984), Stein (1997) and Scharfstein and Stein (2000)) have argued that the CEO's authority to move capital across divisions is the single most defining characteristic of an integrated firm.

If this view is correct, it leads to the following qualitative conclusions. Suppose that based purely on division-manager research incentives, an analysis such as that in Section II.D cuts in favor of decentralization. Suppose further that the other, non-capital-allocation-related synergies from integration are given by X . When X is relatively small, it will make sense to decentralize "to the max" by breaking up the firm. However, when X is relatively large, it may be better to preserve X by keeping the firm integrated, but at the same time to do as much as possible with various internal devices (reputation, overload, etc.) to mimic—even if one cannot fully replicate—the decentralized outcome in terms of division-manager incentives.

B. CEO Research Incentives: Why Soft Information Favors Flatter Organizations

In modeling a hierarchy with soft information, I have thus far taken as exogenous q , the probability that the CEO's research will be successful. Now I ask what happens in the soft-information case when the CEO's incentives are also taken into account—i.e., when $q = q(e^{CEO})$, where e^{CEO} is the research effort exerted by the CEO, and where $q(\cdot)$, like $p(\cdot)$, is an increasing

¹³Once one takes account of other factors, it is hard to imagine that the CEO would ever want to divest this particular right, even if she legally could. For example, given incomplete contracting, it might be difficult for the CEO to surgically excise just her right to move capital around without interfering with her ability to conduct other aspects of her job.

concave function. For simplicity, I stick to the case where the star-gets-everything condition in (5) is satisfied.

Denote by Δ^{CEO} the expected gain in net output that arises when, in a hierarchy with soft information, the CEO's research effort is successful. This quantity is easily calculated by evaluating Y^{hs} in equation (8) for both $q = 1$ and $q = 0$, and taking the difference:

$$\Delta^{CEO} = p(e^{hs})(2g(2) - b(1) - 3g(1))/4 + (1 - p(e^{hs}))(3g(2) - 2b(1) - 4g(1))/4 \quad (12)$$

An immediate consequence of equation (12) is that the CEO's research is more valuable when division managers do less research—i.e., when $p(e^{hs})$ is low. Specifically, differentiation of (12) yields:

$$d\Delta^{CEO}/dp(e^{hs}) = b(1) + g(1) - g(2) < 0 \quad (13)$$

It then follows from the CEO's first-order conditions that $de^{CEO}/de^{hs} < 0$ also: the CEO's research effort is reduced when the division managers are working hard.¹⁴ The intuition behind this result is straightforward. Recall that the CEO, even when her research is successful, has at best coarse information. Thus while her tilting of the capital budget adds value on average, it does have a cost in some states of the world. In particular, when faced with one division that is a star (i.e., that is $\{G, G\}$) and one that is not, the CEO gives the star all four units of funding. This full tilting of the capital budget toward the star is optimal if it turns out that the non-star division

¹⁴The model of Aghion and Tirole (1997) produces a similar result.

is $\{B, B\}$. However, the tilt is too extreme if the non-star division turns out to be $\{G, B\}$ —ideally, it would be better to leave a $\{G, B\}$ division with at least one unit of funding .

In other words, because of her coarse information, an activist CEO sometimes takes away too much from a $\{G, B\}$ division. And what is the cost of taking capital away from a $\{G, B\}$ division? It depends on how profitable this division can be expected to be. If its division manager is informed, the $\{G, B\}$ division generates a greater expected return from a given allocation of capital, as this capital is always steered to the better project within the division. Thus the opportunity cost of the CEO’s activism is greater when she runs the risk of taking resources away from informed division managers.

Note the symmetry that is at work here: the CEO’s research incentives are blunted by the possibility that the division managers will become informed, much as the division managers’ incentives are blunted by the possibility that the CEO will become informed. That is, the more senior agent in the organization can be discouraged by the hard work of her subordinates, as well as vice-versa.

This symmetry has a couple of consequences. First, since the CEO’s effort e^{CEO} is a decreasing function of division managers’ efforts e^{hs} , and conversely, the model—without further restrictions on the $p(\cdot)$ and $q(\cdot)$ functions—admits the possibility of multiple equilibria.¹⁵ For example, one might have either a “control-freak” equilibrium where the CEO is highly informed and intervenes often, and where the division managers are very discouraged, or a “laissez-faire” outcome where the reverse occurs. Depending on parameter values, one equilibrium will be ex

¹⁵In other words, equilibrium involves the intersection of two downwards-sloping curves in (e^{CEO}, e^{hs}) space. Without any further restrictions on functional forms, it is possible that these two curves may cross more than once.

ante more efficient than the other, and there is no guarantee that a firm will not get stuck in the wrong equilibrium. Thus, for example, the negative consequences of the hierarchical form of organization may be more pronounced for a firm that has gotten stuck in a control-freak culture, as opposed to one that has somehow managed to maintain a laissez-faire environment.

A further implication of the model with endogenous CEO effort is that it suggests another organizational form that may be optimal in some circumstances. In the same way that decentralization can be valuable as a precommitment to get the CEO out of the picture—and thereby increase division-manager incentives—it might sometimes make sense to remove the division managers, so as to increase CEO incentives. It is easy enough to construct numerical examples that have this feature. The key is to make CEO effort both valuable and highly elastic, while making division-manager effort less so. In such cases, we are left with just the CEO overseeing all four investment projects, which can be interpreted as an integrated (i.e., large) firm with a flat management structure.

Moreover, the same logic—that removing some of her subordinates can improve the performance of the CEO—leads to a richer and more interesting rendition of organizational flatness if applied to a scaled-up version of the model. Imagine that there are eight projects within the firm, and that the firm is initially organized in a pyramid structure with three layers of management: four low-level managers, who each oversee two projects; two mid-level managers, who each oversee two low-level-managers; and the CEO. Now consider the effects of stripping out the two mid-level managers, leaving only the four low-level managers to report directly to the CEO. There will be two distinct positive incentive effects. First, the low-level managers will work harder, for the reasons outlined in Section II.D—they face less intervention from their

immediate supervisors. Second, the CEO will also work harder, for the reasons developed just above—she no longer anticipates that her research will be devalued by the competing efforts of her immediate subordinates.

Thus when we compare hard vs. soft information in terms of their implications for organizational form, we now have a new conclusion. Not only does soft information tend to favor decentralization (i.e., smaller firms with fewer projects), but holding fixed a firm's size and scope, soft information also favors a flatter, more streamlined management structure.¹⁶ To put it simply, extra layers of management are more costly when information is soft.

C. Hardness of Information is Endogenous: Excess Bureaucracy in Hierarchies

I have been assuming throughout that the hardness of information is exogenous—either a division manager's information can or cannot be credibly transmitted to the CEO, but there is nothing that the division manager can do to influence this. An alternative approach is to posit that the degree of hardness is endogenous. In particular, there may be cases where, by expending additional effort on documentation, a division manager can harden information that would otherwise be soft.¹⁷

To capture this idea formally, I return to the version of the model in Section II.E, where the CEO does no research of her own, and make one modification. Now each division manager

¹⁶A large flat firm is all the more likely to be attractive if, in addition to soft information, there are other operating synergies that cut in favor of integration.

¹⁷An example would be an academic department chair, who, having already decided that he wants to make a tenure offer to somebody, undertakes the process of writing for outside letters. The letters may provide no new information to the department chair, but they can help him to credibly sell the case to a less-well-informed dean.

can choose between one of two research technologies: he can either put one unit of effort into acquiring soft information, or he can put one unit of effort into acquiring hard information. If he opts for soft information, there is a probability p^s that his research is successful. If instead he chooses to go after hard information, there is a probability p^h that his research is successful, and conditional on success, a probability z that the information actually turns out to be verifiable. I assume that $p^h = \beta p^s$ and $\beta < 1$. Thus effort devoted to acquiring hard information is less productive. A high value of β means that information about the project in question is by its nature relatively easy to document, so that not too much of a price is paid to make it hard; one can think of the previous cases of absolutely “hard” and “soft” information in Section II as corresponding to the polar extremes where $\beta = 1$ and $\beta = 0$ respectively.

It is easy to see that even when β is substantially less than one, the only equilibrium may be one in which both division managers opt to go after hard information. If both managers go after soft information, the CEO never learns anything, and we are effectively in a decentralized outcome, with each division manager’s effort yielding him a private benefit of $p^s \gamma \Delta^d$. In contrast, if one division manager deviates and goes after hard information, we know from part (iii) of Proposition 3 that this manager’s private benefit will exceed $p^h \gamma (\Delta^d + z/4)$. So even when p^h is much smaller than p^s , hard information can yield higher private benefits. This is just an application of the logic developed in Section II.E: hard information is much more attractive to division managers, because, unlike soft information, it can help them get a larger capital budget.

Moreover, while hard information is, all else equal, more valuable to the firm as a whole—it enables the CEO to make value-enhancing reallocations—division managers’ preference for it is far too strong. That is, there are low values of β for which the firm would be better off

if the division managers pursued soft information, but division managers' private incentives are such that they pursue hard information instead.¹⁸ This is because of an externality. Manager i gets a first-order benefit of γ when he generates hard information that lands his division one additional unit of capital. However, manager j loses an equal amount, and the gain to the firm as a whole is only proportional to the improved net return on this one unit.

This logic implies that when project information is innately hard to document (i.e., when β is small but non-zero) the costs of a hierarchical form of organization may manifest themselves not just as division managers slacking off and doing no research, as in Section II.D. Rather, one may observe division managers working extremely hard at creating the wrong kind of information. In other words, a hierarchy may be characterized by a great deal of bureaucracy, in the sense of division managers generating lots of reports that are very well-documented, but ultimately not terribly informative. And conversely, if a hierarchical firm is broken up, the excessive bureaucracy vanishes, and division managers instead produce only soft information.

IV. Empirical Implications: A Closer Look at Banks' Small-Business Lending Practices

The most basic implication of the theory developed above is that large, hierarchical firms are at a comparative disadvantage when information about individual investment projects is innately soft. Moreover, if one takes the model seriously, the hardness or softness that is most relevant for organizational form has to do with information about those "small" projects that are overseen by line managers. In other words, what matters is the nature of information that, in a

¹⁸Note that when $\beta = 1$, the firm is actually more valuable as a result of the fact that division managers choose to go after hard, rather than soft information. This corresponds to the case where information is intrinsically hard, as in Section II.E.

hierarchical setting, would be produced far from the ultimate decisionmaker, the CEO.¹⁹

While these ideas would seem to have significant empirical content, there are many challenges in mapping the theory into a set of precise, differentiating predictions that can be readily tested. For example, one interesting implication of the theory is that small firms might be better able than large ones to engage in certain types of new product development, since the prospects of many new products must often be assessed based on soft information. However, any empirical test of this proposition would have to control for a variety of other mechanisms that could lead to a broadly similar outcome—e.g., the well-known “replacement effect”, whereby large firms are discouraged from innovating for fear of cannibalizing their existing product lines.²⁰ Implementing such a control would most likely require a careful case-by-case analysis of the new products in question.

These sorts of complications underscore why it can be particularly informative to look at small-business lending by banks. Here we have a well-defined “industry” where: i) it is easy to identify the primary “projects” that line managers must choose among—namely, individual loan applications; and moreover ii) it seems quite plausible that information about these particular projects is likely to be innately soft.

¹⁹There are of course a number of big strategic investment decisions where no information produced by line managers comes into play. (E.g.: Should the company look for a merger partner? Or expand overseas?) In such cases, one might think of the CEO—perhaps with the help of outside consultants—more or less doing all the research directly herself, so that information transmission within the firm is not an issue. Whether the information produced by the CEO in these cases is hard or soft is less relevant for organizational form.

²⁰See Tirole 1988 (Chapter 10) for a discussion of the replacement effect.

The following findings emerge from the empirical literature on small-business lending.²¹ First, small banks invest a much greater share of their assets in small-business loans than do large banks (Nakamura (1994), Berger, Kashyap and Scalise (1995), Berger and Udell (1996), Peek and Rosengren (1996) and Strahan and Weston (1996)).²² Perhaps more strikingly, when large banks acquire small banks, the small-business lending of the new combined organization tends to fall significantly (Peek and Rosengren (1998), Berger et al (1998), Sapienza (1998)). Moreover, it appears that the loans that are cut as a result of consolidation are not cut simply because they are negative-NPV—i.e., because the acquiring bank is cleaning out the bad loans of the target. Two pieces of evidence support this view. Berger et al (1998) establish that many of the loans that are cut in the process of consolidation are picked up by other banks in the same local market. And Sapienza (1998) finds that there is no relationship between a borrowing firm's credit quality and the likelihood that it will have its lending relationship severed following a merger.

Although these patterns are broadly consistent with the theory developed above, they do not really pinpoint the exact mechanism at work. That is, they do not explain why large banks might be disadvantaged at small-business lending. However, other findings are beginning to emerge which speak more directly to the central idea of this paper, namely that large organizations are not well-suited to handling soft information. Three types of studies are especially worth noting.

²¹A recent survey by Berger, Demsetz and Strahan (1999) provides a much more detailed and comprehensive discussion of this literature.

²²Relatedly, Brickley, Linck and Smith (2000) document that large banks are the dominant players in densely populated metropolitan areas (where presumably borrowers are more likely to be big firms) while small banks have a greater role in suburban and rural regions.

First, it appears that it is not just bank size per se that discourages small-business lending, but rather organizational complexity. For example, DeYoung, Goldberg and White (1997) show that controlling for a bank's size and age, its proclivity for making small-business loans is also negatively related to the number of branches it has, as well as to its being part of a multi-bank holding company. Keeton (1995) finds similar results, with a particularly negative effect on small-business lending for banks that are owned by out-of-state holding companies.

Second, there is some evidence that in making small-business loans, large banks tend to shy away from those "difficult credits" where soft information is likely to be most important in assessing whether or not the loan is positive-NPV. Berger and Udell (1996) find that large banks charge about 100 basis points less on small-business loans than do small banks, and require collateral about 25% less often. One interpretation of this result is that large banks only lend to those small customers whose financial position is so strong that detailed further investigation is not needed.

Finally, Cole, Goldberg and White (1997) use a new survey of small-business finance to look at differences in the loan approval process across large and small banks. They show that large banks (over \$1 billion in assets) tend to base loan approvals primarily on standard criteria obtained from financial statements. In contrast, "small banks deviate from these criteria more and appear to rely on their impression of the character of the borrower to a larger extent."²³ This evidence fits very nicely with the spirit of the model developed above.

While much of the foregoing discussion has implicitly treated the softness of information

²³More precisely, standard measures of borrower credit quality do a better job of explaining (in an R^2 sense) the loan approval decisions of large banks.

in small-bank lending as an exogenously fixed parameter, there is evidence that this parameter is changing over time, with improvements in information technology, widespread adoption of credit scoring models, and the growth of “infomediaries” such as Dun and Bradstreet. Petersen and Rajan (2000) document that the physical distance between small firms and their bankers has been growing—from an average of 16 miles in the 1970's to 68 miles in the 1990's—a pattern which they interpret as evidence that an increasing amount of hard information is being brought to bear on credit decisions. If their interpretation is correct, and if this trend continues, then the model of this paper suggests that the comparative advantage of small banks in small-business lending should diminish in the future.

V. Related Theoretical Work

The ideas in this paper are related to several distinct strands of earlier theoretical work. Rather than attempting a comprehensive survey, I will just briefly discuss a few of the most direct linkages. One branch of the literature takes the perspective that firms are organized so as to be maximally efficient at the processing and communication of various types of information. (See, e.g., Sah and Stiglitz (1986), Radner (1993), Bolton and Dewatripont (1994), Harris and Raviv (1999)). Although information production and transmission are clearly central to my story as well, I differ from these other works in a couple of ways. First, I focus explicitly on managers' private incentives to create and pass on various types of information, whereas the above-mentioned papers abstract away from agency problems within the firm. Second, the notion of authority is more prominent in my model. In particular, the CEO does more than just listen to and act on reports from her subordinates; she actually controls the resources that these

subordinates work with and derive utility from.

The soft-information variant of the model—which emphasizes how the CEO’s capital-allocation authority can discourage division managers from doing research—is, as has already been noted, closely related to Aghion and Tirole (1997). On a similar note, Rotemberg and Saloner (1994) argue that firms may wish to avoid being too broad in scope. For if there are credit constraints at the firm level, such narrowness can help the CEO commit to employees that she will adopt any good ideas that they generate, thereby strengthening ex ante research incentives. More generally, the idea that agents’ incentives are weaker when they do not have control over asset-allocation decisions is familiar from the work of Grossman and Hart (1986), Hart and Moore (1990), and Hart (1995).

However, a sharp distinction between my model and these “costs-of-integration” theories arises when information is hard, rather than soft. With completely hard information, there is no downside to integration in my model. To the contrary, the fact that division managers do not have control actually serves to heighten their incentives, as they struggle to produce enough positive information to convince the CEO to give them a larger share of the capital budget. Thus the model not only paints a generally more favorable picture of the incentive effects of integration than much of the recent literature, its comparative statics with respect to the hardness/softness of information also imply more nuanced empirical implications. The empirical distinctions among the theories are underscored by the facts from the banking industry: the other cost-of-integration stories cannot easily explain why large banks might be at more of a disadvantage in small-business lending than in, say, credit-card or mortgage lending.

In some ways, the hard-information version of the model—with managers trying to

convince the CEO to give them more capital—is also reminiscent of the influence-cost literature (Milgrom (1988), Milgrom and Roberts (1988), Meyer, Milgrom and Roberts (1992)). However, when information is innately hard (i.e., when $\beta = 1$), the welfare implications are reversed. Unlike in the influence-cost models, division managers' efforts to sway the CEO are productive here, rather than wasteful.²⁴ It should be noted, though, that this positive result is quite sensitive to the details of the information structure. When the hardness of information is endogenous and β is low (as in Section III.C), division managers' efforts to attract more capital can lead to inefficient levels of bureaucracy, a result very much in the spirit of the influence-cost theories.²⁵

Finally, it is worth touching on the connection between my model and recent work by Hart and Moore (1999). In their model, some agents (specialists) have ideas about individual assets, while other agents (coordinators) have ideas about how to use multiple assets together. These ideas are mutually exclusive, so that only one agent's idea can be implemented with a given asset. Moreover, there is no ex post renegotiation. In this setting, the organization-design problem is to allocate decision rights ex ante in such a way as to make sure that the best ideas get implemented ex post.

Perhaps the most significant distinction between the models has to do with Hart and Moore's assumption that the ideas of different agents in the firm are mutually exclusive. This exclusivity in turn rests on the premise that the firm's assets can be combined in certain

²⁴The positive, information-creating effects of self-interested advocacy have also been emphasized by Rotemberg and Saloner (1995) and Dewatripont and Tirole (1999).

²⁵Other papers that stress the negative aspects of intra-firm struggles for capital include Rajan, Servaes and Zingales (2000) and Scharfstein and Stein (2000).

synergistic ways—a premise that may be more appropriate in some situations than others.²⁶ In contrast, in my model, the research of the CEO and the division managers can be ex post complementary. Specifically, the CEO’s research can help her decide to allocate a certain amount of capital to a given division manager, who then draws on his own more detailed knowledge to make the right sub-allocations to individual projects within the division. As a result, there is no issue of ex post inefficiency when both the CEO and the division managers have “ideas”. Rather, the problem is the ex ante one of creating incentives for them to generate these ideas in the first place.

VI. Conclusions

By way of conclusion, it is useful to point out a limitation of the model developed above. At a fundamental level, the question being asked is: “What organizational form—decentralization or hierarchy—does the best job of allocating capital to competing investment projects?” But in addressing the question, the entire focus has been on information production and transmission inside firms. As a result, the idea that valuable information might also be generated by outside investors—e.g., by traders in the stock market—has been ignored.²⁷ In particular, I have been implicitly assuming that outside investors have a fixed prior about investment opportunities, such that they always find it optimal to give a firm exactly one unit of funding for each project that it

²⁶For example, it is hard to imagine that Jack Welch, the CEO of General Electric, has many ideas about how to use the assets of NBC and GE’s aircraft division together in a synergistic fashion. In the context of my model, Welch would not spend his time thinking about how to generate synergies between these two divisions; instead he would think about which division should get more capital.

²⁷This idea is stressed by Holmstrom and Tirole (1993), among others.

has under its control.

Clearly, it would be nice to incorporate endogenous stock-market information production into the model. There are several further issues that this sort of extension might allow one to address. For example, the ability of outside investors to undertake their own research may alter the net benefits of decentralization relative to hierarchy. Recall that the drawback to splitting up divisions i and j in the model is that there is then no possibility of moving resources from i to j when j has the better investment prospects. But if outside investors can reproduce some of the information that an integrated-firm CEO would otherwise generate, then stock prices might be useful in guiding reallocations from i to j .²⁸

Moreover, further subtleties arise when one recognizes that the quality of stock-market information is likely to depend on firm size and scope. On the one hand, a potential advantage of decentralization is that it leads to stock prices that are specific to narrow, “pure-play” sets of assets, thereby providing more precise guidance for investment decisions. On the other hand, pushed too far, decentralization may dampen the overall amount of stock-market information that is produced: given fixed costs of information acquisition, very small firms may not attract much interest from either sophisticated investors or stock analysts. By taking such factors into consideration, one might hope to develop a more complete understanding of the link between organizational form and the efficiency of capital allocation.

²⁸As noted above, even if the stock market (or any other outside capital provider) produces the same information about divisions i and j as an integrated-firm CEO, it does not follow that decentralization is equivalent to hierarchy. This is because under decentralization there is no analog to the CEO’s authority to take all funding away from a non-star division.

Appendix

Proof of Proposition 2:

First, straightforward calculation establishes that a sufficient condition for the CEO to prefer an allocation of at least three units to a lone-star division (rather than just giving two units to each division) is: $g(2)/2 > (2g(1) + b(1))/3$. This condition is clearly satisfied if, as the proposition requires, $g(2)/2 > (3g(1) + b(1))/4$. So it must be that a lone-star division will get at least three units of funding.

Now suppose that a lone-star division does in fact get exactly three units of funding. It is easy to show that equation (6) in the text is modified so that the gain Δ^{hs} in expected net output when a division manager's research is successful is given by:

$$\Delta^{hs} = \Delta^d - q(2g(2) - 3g(1) - b(1))/16 \quad (\text{A.1})$$

But given the condition in the proposition, we know that $(2g(2) - 3g(1) - b(1))$ is positive. So $\Delta^{hs} < \Delta^d$, which in turn implies that $e^{hs} < e^d$. Thus even when the tilt in the capital budget is less extreme, a hierarchy induces less research effort than under decentralization. That a hierarchy also may also lead to lower output—i.e., that we may have $Y^{hs} < Y^d$ —follows from the same reasoning as in the example that was used to illustrate Proposition 1 in the text.

Proof of Proposition 3:

I begin by assuming that the division-manager reporting strategies are as described in part

(i) of the proposition; I will verify momentarily that these strategies are optimal for the division managers. With these reporting strategies, the CEO updates on a quiet division manager using equations (9)-(11) in the text. Given these updating rules, let us ask what the CEO does when she faces one division manager that reports $\{G, G\}$ and another one that is quiet. It is not hard to show that, for any value of $z p^c$, it will be optimal for the CEO to give at least three units of capital to the $\{G, G\}$ division if the following sufficient condition (which is the one required in the statement of the proposition) holds: $g(2)/2 > (3g(1) + b(1))/4$. Conversely, it is easy to see that there is never any advantage to deviating from the equal-funding allocation under any other circumstances.

If the CEO follows these capital-allocation rules, then it must in fact be optimal for a division manager to speak up when his information is either $\{G, G\}$ or $\{G, B\}$. In the former case, he may get a third or fourth unit of capital by speaking up, and in the latter case, he ensures that he will not be reduced down to one or zero units (which could happen if he were quiet and the other division reported $\{G, G\}$). It is also at least weakly better for the division manager to remain silent when his information is $\{B, B\}$. Indeed, for many parameter values it is strictly better, since a division that reports $\{B, B\}$ will get allocated zero units of capital in circumstances when a quiet division would get one unit.

These arguments establish parts (i) and (ii) of the proposition. To prove part (iii), assume that when a $\{G, G\}$ division manager is paired with a quiet one, the former is allocated exactly three units of capital. This will establish a lower bound on the gains to being informed, as being informed would be strictly more attractive if a $\{G, G\}$ division were to get four units of capital in this situation. Now take the perspective of manager i , assuming that manager j exerts effort of e_j

and therefore has a probability of research success of $p(e_j)$. (This implies that there is a probability $zp(e_j)$ that manager j will uncover hard information, and a probability $zp(e_j)/4$ that manager j will be able to document to the CEO that his division is $\{G, G\}$.) A little algebra yields the following expression for the expected gross output gain Δ_i^{hh} to manager i if his research is successful:

$$\Delta_i^{hh} = \Delta^d + z/4 + z(4g(2) - 4g(1))/16 + zp(e_j)(2g(1) + 2b(1) - 2g(2))/16 + z^2 p(e_j)(2g(1))/16 \quad (\text{A.2})$$

This implies that:

$$\Delta_i^{hh} > \Delta^d + z/4 + zp(e_j)(2g(2) + 2b(1) - 2g(1))/16 + z^2 p(e_j)(2g(1))/16 \quad (\text{A.3})$$

It then follows immediately from (A.3) that $\Delta_i^{hh} > \Delta^d + z/4$ for any value of $p(e_j)$, which establishes part (iii) of the proposition. Part (iv) is then obvious, since in a hierarchy with hard information there is both more information produced than under decentralization, plus the added advantage that the CEO reallocates funds across divisions according to a value-maximizing criterion.

References

Aghion, Philippe and Tirole, Jean, 1997. "Formal and Real Authority in Organizations," Journal of Political Economy 105, 1-29.

Baker, George, Gibbons, Robert and Murphy, Kevin J., 1999. "Informal Authority in Organizations," Journal of Law, Economics and Organizations 15, 56-73.

Berger, Allen N., Demsetz, Rebecca S. and Strahan, Philip E., 1999. "The Consolidation of the Financial Services Industry: Causes, Consequences, and Implications for the Future," Journal of Banking and Finance 23, 135-194.

Berger, Allen N., Kashyap, Anil K. and Scalise, Joseph M., 1995. "The Transformation of the U.S. Banking Industry: What a Long, Strange Trip It's Been," Brookings Papers on Economic Activity, 55-218.

Berger, Allen N., Saunders, Anthony, Scalise, Joseph M. and Udell, Gregory F., 1998. "The Effects of Bank Mergers and Acquisitions on Small Business Lending," Journal of Financial Economics 50, 187-229.

Berger, Allen N. and Udell, Gregory F., 1996. "Universal Banking and the Future of Small Business Lending," in Saunders, A. and Walter I., (Ehs.) Financial System Design: The Case for Universal Banking. Irwin, Burr Ridge IL, pp. 559-627.

Bolton, Patrick and Dewatripont, Mathias, 1994. "The Firm as a Communication Network," Quarterly Journal of Economics 109, 809-839.

Brickley, James A., Linck, James S. and Smith, Clifford W. 2000. "Boundaries of the Firm: Evidence from the Banking Industry," working paper, Rochester University.

Cole, Rebel A., Goldberg, Lawrence G. and White, Lawrence J., 1997. "Cookie-Cutter vs. Character: The Microstructure of Small Business Lending by Large and Small Banks," working paper, NYU.

Dewatripont, Mathias and Jean Tirole, 1999. "Advocates," Journal of Political Economy 107, 1-39.

DeYoung, Robert, Goldberg, Lawrence G. and White, Lawrence J., 1997. "Youth, Adolescence and Maturity of Banks: Credit Availability to Small Business in an Era of Banking Consolidation," working paper, NYU.

- Donaldson, Gordon, 1984. Managing Corporate Wealth. Praeger Publishers, New York.
- Grossman, Sanford J. and Hart, Oliver D. 1986. "The Costs and Benefits of Ownership: A Theory of Vertical and Lateral Integration," Journal of Political Economy 94, 691-719.
- Harris, Milton and Artur Raviv, 1996. "The Capital Budgeting Process, Incentives and Information," Journal of Finance 51, 1139-1174.
- Harris, Milton and Artur Raviv, 1999. "Organization Design," working paper, University of Chicago.
- Hart, Oliver D., 1995. Firms, Contracts, and Financial Structure. Oxford University Press, Oxford.
- Hart, Oliver D. and Moore, John, 1990. "Property Rights and the Nature of the Firm," Journal of Political Economy 98, 1119-1158.
- Hart, Oliver D. and Moore, John, 1999. "On the Design of Hierarchies: Coordination Versus Specialization," working paper, Harvard University.
- Holmstrom, Bengt and Tirole, Jean, 1993. "Market Liquidity and Performance Monitoring," Journal of Political Economy 101, 678-709.
- Keeton, William R., 1995. "Multi-Office Bank Lending to Small Businesses: Some New Evidence," Federal Reserve Bank of Kansas City Economic Review 81 63-75.
- Meyer, Margaret, Milgrom, Paul R. and Roberts, John, 1992. "Organizational Prospects, Influence Costs, and Ownership Changes," Journal of Economics and Management Strategy 1, 9-35.
- Milgrom, Paul R., 1988. "Employment Contracts, Influence Activities, and Efficient Organizational Design," Journal of Political Economy 96, 42-60.
- Milgrom, Paul R. and Roberts, John, 1988. "An Economic Approach to Influence Activities in Organizations," American Journal of Sociology 94 Supplement, 154-179.
- Nakamura, Leonard I., 1994. "Small Borrowers and the Survival of the Small Bank: Is Mouse Bank Mighty or Mickey?," Federal Reserve Bank of Philadelphia Business Review, November/December, 3-15.
- Peek, Joe and Rosengren, Eric S., 1996. "Small Business Credit Availability: How Important is Size of Lender?," in Saunders, A. and Walter I., (eds.) Financial System Design: The Case for Universal Banking. Irwin, Burr Ridge IL, pp. 628-655.

Peek, Joe and Rosengren, Eric S., 1998. "Bank Consolidation and Small Business Lending: It's Not Just Bank Size That Matters," Journal of Banking and Finance 22, 799-819.

Petersen, Mitchell A. and Raghuram Rajan, 2000. "Does Distance Still Matter? The Information Revolution in Small Business Lending," working paper, University of Chicago.

Radner, Roy, 1993. "The Organization of Decentralized Information Processing," Econometrica 61, 1109-1146.

Rajan, Raghuram G., Servaes, Henri and Zingales, Luigi, 2000. "The Cost of Diversity: The Diversification Discount and Inefficient Investment," Journal of Finance 55.

Rotemberg, Julio J. and Saloner, Garth, 1994. "Benefits of Narrow Business Strategies," American Economic Review 84, 1330-1349.

Rotemberg, Julio J. and Saloner, Garth, 1995. "Overt Interfunctional Conflict (And Its Reduction Through Business Strategy)," RAND Journal of Economics 26, 630-653.

Sah, Raaj Kumar and Stiglitz, Joseph E., 1986. "The Architecture of Economic Systems: Hierarchies and Polyarchies," American Economic Review 76, 716-727.

Sapienza, Paola, 1998. "The Effects of Banking Mergers on Loan Contracts," working paper, Northwestern University.

Scharfstein, David S. and Stein, Jeremy C., 2000. "The Dark Side of Internal Capital Markets: Divisional Rent-Seeking and Inefficient Investment," Journal of Finance 55.

Stein, Jeremy C., 1997. "Internal Capital Markets and the Competition for Corporate Resources," Journal of Finance 52, 111-133.

Strahan, Philip E. and Weston, James, 1996. "Small Business Lending and Bank Consolidation: Is There Cause for Concern?," Federal Reserve Bank of New York Current Issues in Economics and Finance 2, 1-6.

Tirole, Jean 1988. The Theory of Industrial Organization. MIT Press, Cambridge MA.

Williamson, Oliver E., 1967. "The Economics of Defense Contracting: Incentives and Performance," in McKean, R., (ed.) Issues in Defense Economics. Columbia University Press, New York.

Williamson, Oliver E., 1975. Markets and Hierarchies: Analysis and Antitrust Implications. Collier Macmillan, New York.

Williamson, Oliver E., 1988. "Corporate Finance and Corporate Governance," Journal of Finance 43, 567-591.