

Information Production and Capital Allocation: Decentralized versus Hierarchical Firms

JEREMY C. STEIN*

ABSTRACT

This paper asks how well different organizational structures perform in terms of generating information about investment projects and allocating capital to these projects. A decentralized approach—with small, single-manager firms—is most likely to be attractive when information about projects is “soft” and cannot be credibly transmitted. In contrast, large hierarchies perform better when information can be costlessly “hardened” and passed along inside the firm. The model can be used to think about the consequences of consolidation in the banking industry, particularly the documented tendency for mergers to lead to declines in small-business lending.

IN THIS PAPER, I TAKE UP the following question: How does organization design influence the process by which capital is allocated 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 goal is to understand what project-level characteristics lead either decentralization or hierarchy to be the preferred design.

This question can be given a concrete motivation. Over the last several years, there has been enormous consolidation in the banking industry, both in the United States and worldwide. This consolidation has been accompanied by widely voiced concerns that the resulting larger banks will lend less to small businesses, which are especially 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 turn its back on a profitable line of business in this way? An informal argument is given by Berger, Demsetz, and Strahan (1999, pp. 165–166):

*Stein is with the Department of Economics, Harvard University and NBER. I am grateful to the National Science Foundation for research support, as well as to Ulf Axelson; Antoine Faure-Grimaud; Luis Garicano; Rob Gertner; Rick Green; Bengt Holmstrom; Anil Kashyap; Raghu Rajan; Julio Rotemberg; an anonymous referee; and seminar participants at Harvard Business School, the NBER, and the University of Chicago for their input.

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

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

On the one hand, this informal argument is quite clear in asserting that there exist “organizational diseconomies” that 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 vague 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 other banking activities—that might lead it to be a particularly poor fit for a large banking firm? In what follows, I argue that the key distinguishing characteristic of small-business lending is that it relies heavily on information that is “soft”—that is, 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 and hardworking—in other words, the classic candidate for an unsecured “character loan.” Unfortunately, these attributes cannot be unambiguously documented in a report that the loan officer can pass on to his superiors. This situation contrasts sharply with, for example, 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

With soft information, the advantage of decentralization is that it strengthens the research incentives of line managers. Under full decentralization, a line manager is also the CEO of his firm and, as such, has the authority to allocate the firm’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

² The accounting literature has long drawn a similar distinction between hard and soft information. See, for example, Ijiri (1975) and Demski (1994) for discussions and further references.

differently, decentralization rewards an agent who develops expertise by ensuring that he will also have access to some capital that he can use to lever that expertise.

In contrast, if a line manager works inside a large hierarchy, he faces the risk that somebody higher up in the organization will decide that investment opportunities are better elsewhere in the firm and will sharply cut his capital allocation. In this case, because he does not get a chance to act on the information that he has produced (and because he is unable to credibly pass it on), the line manager's research effort goes to waste. Ex ante, this implies that he does less research in a hierarchical setting. Here the authority to allocate capital is separated from expertise, which tends to dilute the incentives to become an expert.

Thus, with soft information, the advantage of decentralization relative to hierarchy is higher-powered research incentives and better capital allocation *within* operating units. However, there is also a countervailing cost: Under decentralization, reallocations *across* operating units have to be mediated by the external capital market, while under a hierarchical design, they are implemented by the integrated firm's CEO. And there can be circumstances in which a CEO is able to bring higher-quality information to bear on such across-unit reallocation decisions than the external market does. Indeed, a central assumption throughout the analysis is that the CEO of an integrated firm has an advantage over the capital market in learning about the prospects of the units that she oversees. As I argue in detail below, this assumption can be rationalized in a number of ways, perhaps most convincingly by appealing to a complementarity between the CEO's strong control rights and her incentives to produce information.

In sum, when information is soft, decentralization has both advantages and disadvantages. Balancing the two, decentralization will on net tend to be a good design under soft information to the extent that line-manager research—and the accompanying potential for efficient within-unit reallocations—is relatively valuable.

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 than the external capital market in terms of moving money across operating units, it can also generate more research on the part of line managers, and, hence, better within-unit allocations. This is because with hard information, line managers become advocates for their units; if they can produce verifiable positive information and pass it on to their superiors, they can increase their capital budgets. 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 more of the firm's resources.

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 follows if one extends the model so that the hardness of information is made endogenous. Suppose that by devoting effort to documentation, 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. Thus, in this modified setting, the costs of a hierarchy do not take the form of line managers simply slacking off; rather, they work 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 IV 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 do not have ultimate authority, I am following Aghion and Tirole (1997). However, a key 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, it is a more general proposition that line managers are discouraged when they do not have full authority. Thus, the models have different empirical implications, as the example of the banking industry suggests: The Aghion–Tirole model does not explain why large banks might be particularly unsuited to small-business lending, as opposed to credit-card or mortgage lending.

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

I. The Model

A. Basic Structure

The model considers two operating units, i and j , which may be organized either as two separate stand-alone firms or as subsidiaries under the same roof of a single integrated firm. Abusing terminology slightly, I will refer to

these operating units as “divisions” in both the decentralized and hierarchical cases. Within each division, there are two potential investment projects. Each of the projects can be allocated either zero, one, or two units of capital, and each can be in either a G (good) or B (bad) state. The probability of either state is $1/2$, and the outcomes are independent across all projects, both within and across divisions.

A project that is in the G state yields a verifiable net output of $g(1) > 0$ if it gets one unit of capital and $g(2) > 0$ if it gets two units. 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. I assume throughout that $g(2) < 2g(1)$; that is, there are decreasing returns in the good 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.

An innocent normalization that reduces notational clutter, which I adopt from this point on, is to set $b(1) = 0$. Finally, I assume that $-1 < b(2) < 0$, as well as that $b(2) + g(2) < 0$. That is, the net return to investing two units in a bad project is so negative that it offsets the gains from investing two units in a good project. As will become clear below, this assumption ensures that there will be credit constraints in equilibrium, since an uninformed provider of external finance will never want to give each division four units of capital to work with. And financing frictions are necessary in this model if one is to pose questions of internal capital allocation; without such frictions, all projects would get fully funded in all states of the world.³

Each of the two divisions has its own division manager. The division managers are, by virtue of research effort, able to obtain signals about the projects that they oversee. If manager i makes an effort e_i , he has a probability $p(e_i)$ of observing signals on *both* of his projects, 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. In the hierarchical setting, when the two divisions are integrated into a single firm, the firm also has a CEO, who may also 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.

To model the incentives of the divisions managers, as well as the CEO (when there is one), I assume that any agent k has a utility function of the form:

$$U_k = (y_k + I_k) - \gamma e_k, \quad (1)$$

³ In my model, the firm invests everything it raises from the outside market, due to the CEO’s empire-building tendencies. By contrast, in papers such as Antle and Eppen (1985) and Harris and Raviv (1996), there can be internal capital rationing even absent external constraints, as the CEO, acting as a principal, seeks to curb the information rents of her subordinates.

where y_k is the expected net output of agent k 's division (or of the whole firm, if k is the CEO), I_k is the amount initially invested in the division (or in the whole firm, if k is the CEO), and γ is a measure of the degree of effort aversion. Thus, each agent seeks to maximize the expected gross output from the assets under his or her control, as given by $(y_k + I_k)$, less the cost of effort. This assumption can be motivated based on: (a) private benefits of control that are proportional to gross output, and (b) nonresponsiveness of agents to monetary incentives.⁴ It has the following behavioral implication: Each agent prefers more capital to less, but conditional on being granted a certain budget, 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.

The goal of the subsections that follow is to establish two principal results. First, when information is soft, decentralization may for some (though not all) parameter values be the more efficient mode of organization. Second, when information is hard, hierarchy always dominates decentralization. To proceed, I begin by analyzing the case of decentralization and evaluating the net surplus created by investment. Next, I turn to the case of a hierarchy with soft information and again do a similar evaluation, which then allows me to compare the two regimes. Finally, I take on the case of a hierarchy with hard information and again compare the net surplus created by investment in this setting to that under decentralization.

B. Decentralization

B.1. External Financing Constraints Under Decentralization

The analysis of the decentralized case can be broken into two parts. First, one has to determine how much a decentralized firm is able to raise from the external capital market. And second, one has to figure out how efficiently this capital will be allocated internally. I assume that investors in the external capital market cannot directly observe the signals (if any) received by division managers.⁵ And given that division managers always prefer more capital to less, and are not responsive to monetary incentives, there is no incentive-compatible way for outside investors to get them to reveal their private information—division managers will always want to claim that both of their projects are in the G state, so as to get more financing. Thus, the best that outside investors can do is to give each division a fixed, un-contingent allocation of capital. Depending on parameter values, this allo-

⁴ Nonresponsiveness to monetary incentives is a common—albeit extreme—modeling shortcut. It can be generated by assuming that agents are infinitely averse to risk in their monetary income, though not to variations in private benefits. See Aghion and Tirole (1997).

⁵ This assumption is a natural one when division managers' signals are soft information. It requires more careful elaboration when division managers' signals are hard information. This issue is taken up in detail in Section I.D below.

cation will be set at either two or three units per division; an allocation of four units per division can never be optimal given the assumption that $b(2) + g(2) < 0$.

I start by simply assuming that each division gets an allocation of two units from the external market. I then solve for expected output, given division managers' endogenously determined incentives in this case. Next, I repeat the analysis, assuming instead that each division is able to raise three units of external capital. Then it is a simple matter of comparing net output across the two cases and checking which leads to higher ex ante returns for investors.

B.2. Incentives and Output with Two Units of Capital per Division

If a division manager with two units of capital is successful in his research efforts and observes $\{G, B\}$ for his two projects, he will behave efficiently ex post and give both units to the G project. The only question is how much effort he will put into research ex ante. To answer this question, note that if a division manager knows the states and allocates the two units based on this knowledge, expected net output is $(g(2) + g(1))/2$. (Again, recall that $b(1)$ is normalized to zero.) If, on the other hand, the manager is uninformed, each project always gets one unit of funding, and expected net output is simply $g(1)$. Thus, with a budget of two units, the manager's utility gain from being informed under decentralization, denoted by Δ_2^d , is

$$\Delta_2^d = (g(2) - g(1))/2. \quad (2)$$

Given equations (1) and (2), the manager's first-order conditions imply that, with two units of financing, the level of research effort under decentralization, e_2^d , satisfies

$$p'(e_2^d) = \gamma/\Delta_2^d. \quad (3)$$

With two independent divisions each behaving this way, the total expected net return on four units of capital under decentralization, denoted $Y^d(4)$, is given by

$$Y^d(4) = p(e_2^d)(g(2) + g(1)) + (1 - p(e_2^d))(2g(1)). \quad (4)$$

B.3. Incentives and Output with Three Units of Capital per Division

The analysis of the case where each division has three units of capital is similar. If a division manager knows the states and can allocate the three units conditional on this knowledge, expected net output is $(3g(2) + g(1) + b(2))/4$. Note that in this case, even if the manager is informed, there are some states—that is, those with two B projects—where it is impossible for outside investors to avoid the negative $b(2)$ outcome.

If the manager is uninformed, he must randomly allocate two units to one of the projects, and one to the other, which leads to expected net output of $(g(2) + g(1) + b(2))/2$. Thus, with three units of financing per division, the utility gain to being informed under decentralization, denoted by Δ_3^d , is given by

$$\Delta_3^d = (g(2) - g(1) - b(2))/4. \quad (5)$$

The level of research effort with three units of financing, e_3^d , is given by

$$p'(e_3^d) = \gamma/\Delta_3^d. \quad (6)$$

With two independent divisions each behaving this way, the total net return on six units of capital under invested under decentralization, $Y^d(6)$, is given by

$$Y^d(6) = p(e_3^d)(3g(2) + g(1) + b(2))/2 + (1 - p(e_3^d))(g(2) + g(1) + b(2)). \quad (7)$$

A comparison of equations (2) and (5) implies that research incentives will be stronger when a division is allocated three, rather than two units of capital, if $b(2) + g(2) - g(1) < 0$. This condition is always satisfied given the initial assumptions that $b(2) + g(2) < 0$ and $g(1) > 0$. Intuitively, being informed is more valuable when there are three units of capital to allocate, because now it helps avoid the very adverse outcome where two units are put in a bad project.

Of course, it does not follow that allocating three units to each division is the ex ante optimal policy for outside investors. Indeed, as can be seen by comparing equations (4) and (7), if $b(2)$ is sufficiently negative, investors will prefer to only give two units to each division. Although the exact cutoff value of $b(2)$ for which this is the case depends on the shape of the $p(\cdot)$ function, it is easy to establish the following sufficient condition (see the Appendix for details)

LEMMA 1: *If $3(g(2) - g(1)) + b(2) < 0$, outside investors will always choose to allocate two units of capital to each division, regardless of the shape of the $p(\cdot)$ function.*

In much of what follows, I assume that the sufficient condition in Lemma 1 holds, thereby focusing on regions of the parameter space for which equation (4) summarizes the output effects of decentralization. This is done solely for expositional purposes—it cuts down on the number of different cases to be considered, without changing the important conclusions.

C. Hierarchy: The Case of Soft Information

C.1. The CEO's Research Technology

The next regime 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.

But allowing the CEO to gather information raises an important question: Why is it that the CEO is able to become better informed than external providers of capital, who are assumed to be incapable of learning anything about the projects? There are two possible answers to this question. First, if there are multiple outside investors (e.g., dispersed shareholders), free-riding problems will naturally reduce their incentives to acquire information. Second, and more subtly, even a single outside investor with the same research capabilities as the CEO may endogenously choose to do less information gathering, to the extent that this outside investor has weaker control rights than the CEO. That is, there is a complementarity between authority (either formal or informal) and research incentives.⁶

As an example, consider the differing incentives of a bank lender versus a CEO when both contemplate investing effort in learning about a specific division. The bank can use any information it acquires to guide its subsequent decision of how much to lend to the division. But suppose that in the course of its investigation, the bank also learns that the division would be more valuable if its manager were replaced, or if some of its other existing physical assets were reconfigured. Outside of default, the bank does not have the authority to impose such outcomes on a reluctant division manager. In contrast, a CEO overseeing the division can. As a result, the CEO has more to gain from acquiring information about the division in the first place, and hence does more research, even if her research technology is no different from that of the banker.

Nevertheless, even if it is plausible that the CEO learns more about investment prospects than outside investors, 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 only gets coarse information about the aggregate 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 is 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

⁶ This point is modeled by Gertner, Scharfstein, and Stein (1994).

division that is “average” (has one G and one B project) and a division that is a “dog” (has two B projects).

A few points about this formulation deserve comment. First, the CEO’s research-success probability q is, for the time being, an exogenous parameter. In Section II.B, I discuss what happens when q is made an endogenous function of the CEO’s effort. In either case, I do not allow the CEO to precommit to not doing research (i.e., to setting $q = 0$), even if this precommitment might raise ex ante expected net output.⁷ 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. Finally, I am assuming that the outcome of the CEO’s research effort is independent of the outcome of the division managers’ efforts. This simplifies the exposition slightly, but is not necessary. Similar conclusions emerge if, for example, the CEO is more likely to succeed in learning something conditional on the division managers having also been successful in their research.

C.2. External Financing Constraints in the Hierarchical Case

As before, a complete analysis of the hierarchical case involves solving for both the optimal amount of funding given to the hierarchy by the external capital market as well as the incentives which govern how these funds are allocated internally.⁸ And things are more complex now, because without further restrictions on parameters, there are four possible funding levels to be considered: four, five, six, or seven units. Even if the condition in Lemma 1 is satisfied, so that the two divisions can each raise only two units in the decentralized case, it is possible that when they are integrated, the hierarchy will be able to raise five or more units. This is due to the “ease-the-credit-constraint” effect identified by Stein (1997). Intuitively, because of the diversification inherent in a larger internal capital market, there is a lower probability that an extra unit of financing will be invested in a bad project, leading to the negative $b(2)$ outcome.

Tackling all the scenarios corresponding to the various possible funding levels is, as before, straightforward. However, such a detailed treatment of external credit constraints does not serve to advance the main goal here, which is to show that—in spite of hierarchy’s potential for beneficial across-division reallocations—with soft information, it is possible for *some parameter values* to have decentralization be the preferred organizational form. To illustrate this point in the most transparent way, I impose another restric-

⁷ This no-precommitment assumption is discussed further in Section II.A. Cremer (1995) is another paper in which a supervisor’s inability to commit to staying uninformed can weaken the incentives of the agent she supervises.

⁸ The feature that no information can be credibly revealed to the external capital market still remains. Since the CEO also derives private benefits that increase with total investment, she will have the same incentive to misrepresent investment prospects as the division managers.

tion on the parameters that is sufficient to ensure that a hierarchy will never be able to raise more than four units of funding. This restriction puts aside the ease-the-credit-constraint effect, thereby isolating the trade-off between giving a CEO the ability to make transfers across divisions, versus diluting the research incentives of the division managers. In the Appendix, I prove the following lemma.

LEMMA 2: *If $25(g(2) - g(1)) + b(2) < 0$, outside investors will only give four units of capital to a hierarchy, regardless of how efficiently the hierarchy allocates this capital internally.*⁹

Again, the point is not that the condition required in the lemma—effectively, that the negative $b(2)$ outcome is very bad—should be thought of as representing the most likely case. Rather, it enables one to focus on the simplest example of the costs and benefits of hierarchy, without worrying that this example is inconsistent with optimization on the part of investors.

C.3. Incentives and Output When the Hierarchy Has Four Units of Capital

Given that it is able to raise four units of external finance, 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 corresponding decentralized outcome.¹⁰ 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: (a) still give each division two units, (b) give the star division three units and the other division one unit, or (c) 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:

$$g(2)/2 > 5g(1)/6. \tag{8}$$

⁹ In proving the lemma, I implicitly rule out schemes where the CEO raises four units of financing, irrevocably sinks it all in division i , and then later returns to the capital market again, asking for two more units for the as-yet-unfunded division j . One way to rationalize this restriction is to assume that outsiders cannot verify whether investment expenditures "belong" to one division of an integrated firm versus another; such an assumption is a common one in the literature on internal capital allocation. See Scharfstein and Stein (2000) for a discussion.

¹⁰ I do not consider mechanisms that the CEO might use to induce division managers to reveal their signals. One could imagine that if the CEO catches a manager lying about his signal, he would be punished with a reduced capital allocation, as in Harris and Raviv (1996). In my setup, such a scheme suffers from two distinct commitment problems. First, punishment involves ex post inefficient 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.

This condition simply requires that decreasing returns to scale are not too pronounced in the good state. The analysis that follows is most intuitive 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 that hierarchy is the strictly dominant 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 effort is reduced, and he produces less information.¹¹

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 would be able to raise two units if it were a stand-alone entity—when her research indicates that division j is a star.¹² As emphasized by Stein (1997), this authority distinguishes the CEO from an equally well-informed outside provider of finance, such as a banker.

¹¹ This discouragement effect is not offset by the fact that the manager also sometimes gets two extra units in a hierarchy. When he receives four units, the only thing he can do is invest two in each project, and information is again not of any value. Moreover, the feature that a stable capital budget creates stronger research incentives than a variable one is actually more general than the simple setup here suggests. Suppose instead that investment is continuous, that an amount I yields $f(I)$ in state G , 0 in state B , and that a manager has one project of each type. If he is uninformed and has a budget of K , he will invest $K/2$ in each project, yielding $f(K/2)$. If he is informed, he will get $f(K)$. Hence, the value of being informed is $V(K) = f(K) - f(K/2)$. A stable capital budget will lead to more research if $V(K)$ is concave, or if $4f''(K) - f''(K/2) < 0$. This mild condition is satisfied, for example, for any $f(I)$ of the power form $f(I) = I^\alpha/\alpha$, for $0 < \alpha < 1$.

¹² Indeed, 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 rather because the prospects of its counterpart division within the firm are so strong.

Assume that the timing of the game is such that a division manager must choose his level of research effort before he knows whether the CEO's research has been successful. Denote by Δ_4^{hs} the gain in expected utility that an individual division manager obtains when, in a hierarchy with soft information and four units of capital, his own research efforts succeed. It is straightforward to show that

$$\Delta_4^{hs} = (1 - q)(g(2) - g(1))/2 + 3q(g(2) - g(1))/8 = (1 - q)\Delta_2^d + 3q\Delta_2^d/4. \quad (9)$$

The corresponding level of research effort, e_4^{hs} , satisfies

$$p'(e_4^{hs}) = \gamma/\Delta_4^{hs}. \quad (10)$$

Since $\Delta_4^{hs} < \Delta_2^d$, it follows that $e_4^{hs} < e_2^d$. By working through all the possible outcomes, it can then be established that expected net output in a hierarchy with soft information and four units of capital is given by

$$Y^{hs}(4) = (1 - q)\{p(e_4^{hs})(g(2) + g(1)) + (1 - p(e_4^{hs}))(2g(1))\} \\ + q\{p(e_4^{hs})(6g(2) + g(1))/4 + (1 - p(e_4^{hs}))(3g(2) + 4g(1))/4\}. \quad (11)$$

By comparing equations (11) and (4), one can evaluate the relative efficiency of decentralization versus hierarchy, conditional on there being a total external capital constraint of four units in either case. The results to this point can be summarized in the following proposition.

PROPOSITION 1: *Assume that the condition in Lemma 2 holds, so that the external credit constraint is four units. Assume further that inequality (8) holds, so that in a hierarchy, a lone-star division gets all four units of funding. Then decentralization always leads to more research effort than hierarchy: $e_4^d > e_4^{hs}$. In addition, it is possible (though not necessary) that decentralization leads to higher expected output; that is, that $Y^d(4) > Y^{hs}(4)$.*

To see why decentralization can generate higher expected output, consider a simple limiting case where $q = 1$, $p(e_4^d) = 1$, and $p(e_4^{hs}) = 0$. (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 $Y^d(4) = g(2) + g(1)$, while equation (11) simplifies to $Y^{hs}(4) = (3g(2) + 4g(1))/4$, implying that $Y^d(4) > Y^{hs}(4)$. In this example, when the firm is organized as a hierarchy, only the CEO does any research, and 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 $p(e_4^d)$ and $p(e_4^{hs})$ being relatively far apart—that is, on division-manager effort being both important and responsive to incentives. If $p(e_4^d)$ and $p(e_4^{hs})$ are sufficiently close to one another, it is easy to see that hierarchy becomes more efficient than decentralization: $Y^{hs}(4) > Y^d(4)$.

Although I have derived the results in Proposition 1 under the assumption that the star-gets-everything condition in (8) 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 negative effect on division managers' research incentives.

To make this idea precise, in the Appendix I prove the following proposition.

PROPOSITION 2: *Assume that the condition in Lemma 2 holds, so that the external credit constraint is four units. Assume further that $g(2)/2 > 3g(1)/4$. 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: $e_4^d > e_4^{hs}$. In addition, it is possible (though not necessary) that decentralization leads to higher expected output than hierarchy, that is, that $Y^d(4) > Y^{hs}(4)$.*

D. Hierarchy: The Case of Hard Information

The next task is to show that the relative merits of a hierarchy increase when the information generated by division managers can be hardened and passed along to the CEO. I begin by assuming that the sufficient condition in Lemma 2 continues to hold. This allows for a simple comparison of hierarchy relative to decentralization in a situation where the total amount of available external finance is four units in either case. But once I have established that a hierarchy with hard information dominates decentralization for these parameter values, it is easy to extend the argument to show that it must dominate for *all* parameter values—that is, regardless of how many units of financing can be raised under either decentralization or hierarchy.

To introduce hard information, I assume 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. 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 nonverifiable and can be used by the division manager, but not credibly transmitted to the CEO.

A critical assumption is that while hard information can be credibly transmitted from division managers to the CEO, it *cannot* be revealed to outside investors. Otherwise, the hierarchy would no longer face a fixed external credit constraint. What is the economic interpretation of such an assumption? Think of division managers as providing the CEO with a variety of raw, albeit well-documented data about a project.¹³ For example, if the project involves drilling for oil in a new location, the raw data might be a set of geological studies. These studies do not literally say what the project's dollar payoff will be. Instead, they must be combined with the CEO's specific expertise (e.g., her knowledge of geology, her assessment of the costs of doing the drilling and extraction) to generate a final judgment about dollar value.

As before, the complementarity of the CEO's authority and research incentives imply that she will in equilibrium be better able to assess such raw data than would an outside banker. Continuing with the example, because the CEO has the authority to decide how to deploy the firm's drilling equipment, she will have more incentive to become informed about the costs of drilling in different locales, which in turn makes her better able to assess the value implications of the geological studies. The CEO's final judgment—what she concludes after studying the raw data—is itself soft information that cannot be credibly communicated to outside investors.

Under this interpretation, the parameter z can be thought of as reflecting either 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. Moreover, the idea that the CEO needs to spend time evaluating raw data helps to motivate the assumption that, when information is hard, she no longer does her own separate research into divisional prospects: She is too busy trying to process her subordinates' output.¹⁴ As will be made clear below, the consequence of this assumption is that when information is hard, CEO and division-manager effort become strategic complements, as opposed to substitutes.

Another crucial assumption is that the CEO does not observe whether division-manager research has been successful. Consequently, when a division manager has hard information, *he can choose whether or not to report it*

¹³ Internal accounting or auditing mechanisms may have a role to play in making such raw data sufficiently well documented as to be credible.

¹⁴ Clearly, this aspect of the model could be more fully endogenized. For example, the CEO could face a trade-off between time spent gathering her own independent information and time spent interpreting division managers' reports, and one could establish the conditions under which she tended to find the latter optimal. This line of thinking expands on Aghion and Tirole's (1997) idea of keeping the CEO too busy to do her own separate project evaluation. Here, the CEO is kept busy not because she is overloaded with make-work, but rather because she chooses to engage in another more productive activity.

to the CEO.¹⁵ 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.

The option to keep quiet 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 strategy can constitute equilibrium behavior.) Bayes’ rule implies that, given his conjecture of p^c , the CEO will interpret quiet as follows:

$$\text{prob}(\{B, B\}/\text{quiet}) = 1/(4 - 3zp^c) \quad (12)$$

$$\text{prob}(\{G, G\}/\text{quiet}) = (1 - zp^c)/(4 - 3zp^c) \quad (13)$$

$$\text{prob}(\{G, B\}/\text{quiet}) = 2(1 - zp^c)/(4 - 3zp^c). \quad (14)$$

Taking the CEO’s conjecture of p^c as fixed, a division manager reasons as follows. If he does no research, he will certainly be quiet, and the CEO will update on him according to equations (12)–(14). However, if he does devote some effort to research, he gains pure option value. If the research produces hard information that his division is $\{G, G\}$, 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 the condition in Lemma 2 holds, so that the external credit constraint is four units. Assume further that $g(2)/2 > 3g(1)/4$. Then in a hierarchy with hard information, there is an equilibrium with the following properties:*

¹⁵ A similar assumption has been made in accounting research that addresses the topic of discretionary disclosure to the capital market. See, for example, Verrecchia (1983) and Dye (1985).

- (a) *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 $\{B, B\}$.*
- (b) *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.*
- (c) *Division manager benefit from being informed: Denote by Δ_4^{hh} the utility benefit to a division manager from being informed in a hierarchy with hard information. This benefit exceeds that under decentralization: $\Delta_4^{hh} > \Delta_4^d + z/4$. Consequently, division-manager research effort is greater than under decentralization: $e_4^{hh} > e_4^d$.*
- (d) *Output: Expected net output is greater than under decentralization: $\bar{Y}^{hh}(4) > Y^d(4)$.*

The intuition behind parts (a) and (b) of the proposition has already been discussed.¹⁶ Part (c) gives a sense for how much stronger division managers' research incentives are when information can be hardened. The utility gain to being informed, Δ_4^{hh} , is now not just larger than that under decentralization Δ_4^d ; it exceeds it by at least $z/4$. This can be a substantial difference, since Δ_4^{hh} is denominated in units of net returns, while $z/4$ is in units of gross capital. The difference arises because 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 the division manager's utility is tied to gross output, the former effect is naturally much stronger than the latter. Part (d) of the proposition follows from part (c). 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.

Although Proposition 3 deals with the case where the condition in Lemma 2 holds, so that the hierarchy faces an external credit constraint of four units, the result can easily be generalized. In the appendix, I prove the following proposition.

PROPOSITION 4: *Consider the more general case where the only restriction on $b(2)$ is that $g(2) + b(2) < 0$, while keeping the assumption that $g(2)/2 > 3g(1)/4$. Under decentralization, each division may now be able to raise either two or three units of capital, and, under hierarchy, the aggregate credit*

¹⁶ There are other equilibria that differ insignificantly from that 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 (c) and (d) of the proposition apply in either case.

constraint may be anywhere from four to seven units. Still, if information is hard, expected net output is always strictly greater in a hierarchy than under decentralization.

One might argue that the results in Propositions 3 and 4 in favor of hierarchy relative to decentralization are now “too strong”—once information can be hardened, hierarchy dominates decentralization for all parameter values, which seems patently 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 noncontractible 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 less interested in making absolute statements about the virtues of hierarchy compared to decentralization, and more interested in making comparative-static statements about the circumstances under which hierarchy is 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.

II. Further Issues

A. *Does Decentralization Require Dis-Integration?*

I have been treating the concept of decentralization as synonymous with formal dis-integration, that is, with the two divisions being operated as independent stand-alone firms. The key assumption has been that it is impossible for a single firm to replicate the decentralized outcome, because the CEO cannot commit to either remaining uninformed or to not using her information to reallocate capital across divisions. Given that the model has no operating synergies across the two divisions, not much hangs on the distinction between decentralization and dis-integration. But what if there are significant synergies? Then one must ask if it is possible to capture the benefits of decentralization—that is, to commit to division managers that they will always receive two units of funding—without resorting to a costly breakup.

On the one hand, it is hard to argue that a CEO can definitively alienate her right to get involved in capital-allocation decisions. Indeed, several authors (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.¹⁷ On the other hand, as Aghion and Tirole (1997) point out, there are

¹⁷ Baker, Gibbons, and Murphy (1999) go a step further, claiming that formal authority over all decisions necessarily resides at the top of an organization.

other devices short of a breakup that can reduce—albeit not fully eliminate—the CEO’s ex post incentive to meddle in the capital-allocation process; they suggest keeping the CEO very busy, so she cannot do much research on her own. In a banking context, a relevant example of such a device might be the creation of a multibank holding company. By making each of the divisions a legally distinct entity, such a structure could create some impediments to moving capital freely among them.¹⁸

This logic leads to the following qualitative conclusions. Suppose that based purely on division-manager research incentives, an analysis such as that in Section I.C 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 large, it may be better to keep the firm integrated, but at the same time to do as much as possible with various internal devices (e.g., the multibank holding-company structure) 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—that is, 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 concave function. For simplicity, I stick to the case where the external credit constraint is four units—that is, where the sufficient condition in Lemma 2 is satisfied—and where the star-gets-everything inequality (8) also holds.

The outcome(s) in this setting are the Nash equilibria of the game where the CEO and the division managers choose their research efforts simultaneously. Denote by Δ^{CEO} the utility gain to the CEO when, in a hierarchy with soft information, her research effort is successful. This quantity is easily calculated by evaluating $Y^{hs}(4)$ in equation (11) for both $q = 1$ and $q = 0$, and taking the difference:

$$\Delta^{CEO} = p(e_4^{hs})(2g(2) - 3g(1))/4 + (1 - p(e_4^{hs}))(3g(2) - 4g(1))/4. \tag{15}$$

An immediate consequence of equation (15) is that the CEO’s research is more valuable when division managers do less research—that is, when $p(e_4^{hs})$ is low. Specifically, differentiation of (15) yields

$$d\Delta^{CEO}/dp(e_4^{hs}) = (g(1) - g(2))/4 < 0. \tag{16}$$

¹⁸ I am grateful to Raghu Rajan for suggesting this example.

It then follows from the CEO's first-order conditions that $de^{CEO}/de_4^{hs} < 0$ also: The CEO's research effort is reduced when the division managers are working hard.¹⁹ 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 nonstar division is $\{B, B\}$. However, the tilt is too extreme if the nonstar 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_4^{hs} , and conversely, this game—without further restrictions on the $p(\cdot)$ and $q(\cdot)$ functions—admits the possibility of multiple Nash equilibria.²⁰ 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 ante more efficient than the other, and there is no guarantee that a firm will end up in the better one. 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 precom-

¹⁹ The model of Aghion and Tirole (1997) produces a similar result.

²⁰ In other words, Nash equilibrium involves the intersection of two downwards-sloping curves in (e^{CEO}, e_4^{hs}) space. Without any further restrictions on functional forms, it is possible that these two curves may cross more than once.

mitment 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.

Thus, when we compare hard versus 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. This is because the efforts of higher-ups and subordinates are strategic substitutes when information is soft, but strategic complements when information is hard.

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, and, that by expending additional effort on documentation, a division manager can harden information that would otherwise be soft.²¹

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

²¹ An 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.

cases of absolutely “hard” and “soft” information in Section I as corresponding to the polar extremes where $\beta = 1$ and $\beta = 0$, respectively.

It is easy to see that even when β is much 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 utility gain of $p^s \Delta_4^d$ (assuming the external credit constraint is four units). In contrast, if one division manager deviates and goes after hard information, we know from part (c) of Proposition 3 that this manager’s utility gain will exceed $p^h (\Delta_4^d + z/4)$. 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 above: Hard information is more attractive to division managers, because it can help them get larger capital budgets.

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. This is because of an externality. Manager i ’s utility goes up by one 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 nonzero), 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 I.C. 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.

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

While these ideas would seem to have significant empirical content, there are 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 test of this proposition would have to control for a variety of other mechanisms that could lead to a similar outcome—for example, the well-known “replacement effect,” whereby large firms are discouraged from innovating for fear of cannibalizing their existing product lines.²² 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: (a) it is easy to identify the primary “projects” that line managers must choose among—namely, individual loan applications; and, moreover, (b) it seems quite plausible that information about these particular projects is likely to be innately soft.

The following findings emerge from the empirical literature on small-business lending. 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)).²³ Perhaps more strikingly, when large banks acquire small banks, the small-business lending of the combined organization tends to fall sharply (Peek and Rosengren (1998), Berger et al. (1998), Sapienza (2002)). Moreover, it appears that the loans that are cut as a result of consolidation are not cut simply because they are negative-NPV—that is, 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 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 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 that speak more directly to the central idea of this paper, namely that large organizations are

²² See Tirole (1988, Chapter 10) for a discussion of the replacement effect.

²³ Relatedly, 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.

not well suited to handling soft information. Three types of studies are especially worth noting.

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 multibank holding company.²⁴ 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 percent less often. One interpretation 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" (p. 1).²⁵ 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 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 (2002) document that the distance between small firms and their bankers has been growing—from an average of 16 miles in the 1970s to 68 miles in the 1990s—a pattern which they interpret as evidence that an increasing amount of hard information is being brought to bear on credit decisions. If they are correct, and if this trend continues, then the model suggests that the comparative advantage of small banks in small-business lending may diminish in the future.

²⁴ It should be noted that the implications of the theory for the multibank-holding-company variable are ambiguous. Perhaps it proxies for complexity and the number of layers of management, which would tend to discourage small-business lending. But as argued in Section II.A, a holding-company structure may also make it harder to move capital across divisions, which would help soft-information research incentives and small-business lending.

²⁵ More 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.

IV. Related Theoretical Work

The ideas in this paper are related to several 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 central to my story as well, I differ from these other works in a couple of ways. First, I focus on managers' incentives to create and pass on various types of information, whereas the above-mentioned papers abstract 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.

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-a boss” 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 costs-of-a-boss 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 inately 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.²⁶ But this positive result is sensitive to the details of the information structure. When the hardness of information is endogenous and β is low (as in Section II.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.²⁷

Finally, it is worth touching on the connection between my model and 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 problem is to allocate decision rights ex ante in such a way as to make sure that the best ideas get implemented ex post. In contrast, in my model, the "ideas" 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 knowledge to make the right suballocations to individual projects within the division. Thus, there is no ex post inefficiency when both the CEO and the division managers have ideas. Rather, the issue is the ex ante one of creating incentives for them to generate these ideas in the first place.

V. 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 focus has been on information production and transmission *inside firms*. The notion that valuable information might also be generated by outside investors—for example, by traders in the stock market—has been downplayed.²⁸ In particular, I have been assuming that, because their research incentives are weaker than those of a CEO, outside investors acquire no information about investment prospects and base their decisions simply on their initial priors.

Clearly, it would be nice to incorporate a richer and more fully endogenous view of stock-market information production into the model. There are several further issues that this sort of extension might allow one to address. In particular, the quality of stock-market information may—like the quality of internally generated information—depend on firm size and scope. On the

²⁶ The positive, information-creating effects of self-interested advocacy have also been emphasized by Rotemberg and Saloner (1995) and Dewatripont and Tirole (1999).

²⁷ Other papers that stress the negative aspects of intrafirm struggles for capital include Rajan, Servaes, and Zingales (2000) and Scharfstein and Stein (2000).

²⁸ This idea is stressed by Holmstrom and Tirole (1993), among others.

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, tiny 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.

Appendix

Proof of Lemma 1: An upper bound on the surplus created by a division with a three-unit allocation can be obtained by assuming that the division manager’s research succeeds with certainty, yielding expected net output equal to $(3g(2) + g(1) + b(2))/4$. Conversely, a lower bound for a division with a two-unit allocation can be obtained by assuming that the division manager’s research never succeeds, in which case expected net output is simply $g(1)$. Comparing these two, it must be that two units is preferred if $3(g(2) - g(1)) + b(2) < 0$, which is the sufficient condition given in the lemma. Q.E.D.

Proof of Lemma 2: An upper bound on the surplus created by a hierarchy with a five-unit allocation can be obtained by assuming that information is perfect throughout the hierarchy, so that the five units are always allocated efficiently. In this case, expected net output is $(7g(1) + 25g(2) + b(2))/16$. Conversely, a lower bound for a hierarchy with a four-unit allocation can be obtained by assuming that all agents are uninformed, so that each project always gets one unit of capital, leading to expected net output of $2g(1)$. Comparing these two, it must be that four units is preferred if $25(g(2) - g(1)) + b(2) < 0$, which is the sufficient condition given in the lemma. Similar logic also implies that four units are preferred to six or seven, given the same condition. Q.E.D.

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 $g(2)/2 > 2g(1)/3$. This condition is satisfied if, as the proposition requires, $g(2)/2 > 3g(1)/4$. 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 (9) in the text is modified so that the utility gain Δ_4^{hs} when a division manager’s research is successful is given by

$$\Delta_4^{hs} = \Delta_4^d - q(2g(2) - 3g(1))/16. \quad (\text{A1})$$

But given the condition in the proposition, we know that $(2g(2) - 3g(1))$ is positive. So $\Delta_4^{hs} < \Delta_4^d$, which, in turn, implies that $e_4^{hs} < e_4^d$. 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 lead to lower output—that is, that we may have $Y^{hs}(4) < Y^d(4)$ —follows from the same reasoning as in the example that was used to illustrate Proposition 1 in the text. Q.E.D.

Proof of Proposition 3: I begin by assuming that division-manager reporting strategies are as described in part (a) 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 (12)–(14) 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 zp^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: $g(2)/2 > 3g(1)/4$. Conversely, it is easy to see that there is no 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 (a) and (b) of the proposition. To prove part (c), 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 utility gain Δ_{4i}^{hh} to manager i if his research is successful:

$$\begin{aligned} \Delta_{4i}^{hh} = & \Delta_4^d + z/4 + z(4g(2) - 4g(1))/16 + zp(e_j)(2g(1) - 2g(2))/16 \\ & + z^2p(e_j)(2g(1))/16. \end{aligned} \tag{A2}$$

This implies that

$$\Delta_{4i}^{hh} > \Delta_4^d + z/4 + zp(e_j)(2g(2) - 2g(1))/16 + z^2p(e_j)(2g(1))/16. \quad (A3)$$

It then follows immediately from (A3) that $\Delta_{4i}^{hh} > \Delta_4^d + z/4$ for any value of $p(e_j)$, which establishes part (c) of the proposition. Part (d) 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. Q.E.D.

Proof of Proposition 4 (Sketch): There are two cases to consider. In the first, the credit constraint facing a decentralized division remains at two units, but that for a hierarchy is relaxed, and exceeds four units. In this case, revealed preference on the part of outside investors implies that the hierarchy must do at least as well as it did with four units under the conditions of Proposition 3, so it continues to dominate decentralization.

In the second case, suppose that each decentralized division can raise three units. In this case, a hierarchy must be able to raise at least six units, and maybe more. One can thus put a lower bound on the value created by a hierarchy by assuming that it raises exactly six units. In a hierarchy with six units, it is easy to verify that the following constitutes an equilibrium, for any parameter values: (a) a division manager speaks up if his information is $\{G, G\}$ and keeps quiet otherwise; (b) when just a single division reports $\{G, G\}$, the CEO gives it four units of financing; and (c) in all other cases, both divisions get three units of financing. It then follows from arguments analogous to those in the proof of Proposition 3 that research incentives are stronger in a hierarchy, and, that as a result, net output is also greater in a hierarchy. Q.E.D.

REFERENCES

- Aghion, Philippe, and Jean Tirole, 1997, Formal and real authority in organizations, *Journal of Political Economy* 105, 1–29.
- Antle, Rick, and Gary D. Eppen, 1985, Capital rationing and organizational slack in capital budgeting, *Management Science* 31, 163–174.
- Baker, George, Robert Gibbons, and Kevin J. Murphy, 1999, Informal authority in organizations, *Journal of Law, Economics, and Organizations* 15, 56–73.
- Berger, Allen N., Rebecca S. Demsetz, and Philip E. Strahan, 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., Anil K. Kashyap, and Joseph M. Scalise, 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., Anthony Saunders, Joseph M. Scalise, and Gregory F. Udell, 1998, The effects of bank mergers and acquisitions on small business lending, *Journal of Financial Economics* 50, 187–229.
- Berger, Allen N., and Gregory F. Udell, 1996, Universal banking and the future of small business lending, in A. Saunders and I. Walter, eds.: *Financial System Design: The Case for Universal Banking* (Irwin, Burr Ridge, IL).

- Bolton, Patrick, and Mathias Dewatripont, 1994, The firm as a communication network, *Quarterly Journal of Economics* 109, 809–839.
- Brickley, James A., James S. Linck, and Clifford W. Smith, 2000, Boundaries of the firm: Evidence from the banking industry, Working paper, Rochester University.
- Cole, Rebel A., Lawrence G. Goldberg, and Lawrence J. White, 1997, Cookie-cutter vs. character: The microstructure of small business lending by large and small banks, Working paper, New York University.
- Cremer, Jacques, 1995, Arm's length relationships, *Quarterly Journal of Economics* 110, 275–295.
- Demski, Joel S., 1994, *Managerial Uses of Accounting Information* (Kluwer Publishers, Boston, MA).
- Dewatripont, Mathias, and Jean Tirole, 1999, Advocates, *Journal of Political Economy* 107, 1–39.
- DeYoung, Robert, Lawrence G. Goldberg, and Lawrence J. White, 1997, Youth, adolescence and maturity of banks: Credit availability to small business in an era of banking consolidation, Working paper, New York University.
- Donaldson, Gordon, 1984, *Managing Corporate Wealth* (Praeger Publishers, New York).
- Dye, Ronald A., 1985, Disclosure of nonproprietary information, *Journal of Accounting Research* 23, 123–145.
- Gertner, Robert H., David S. Scharfstein, and Jeremy C. Stein, 1994, Internal versus external capital markets, *Quarterly Journal of Economics* 109, 1211–1230.
- Grossman, Sanford J., and Oliver D. Hart, 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 John Moore, 1990, Property rights and the nature of the firm, *Journal of Political Economy* 98, 1119–1158.
- Hart, Oliver D., and John Moore, 1999, On the design of hierarchies: Coordination versus specialization, Working paper, Harvard University.
- Holmstrom, Bengt, and Jean Tirole, 1993, Market liquidity and performance monitoring, *Journal of Political Economy* 101, 678–709.
- Ijiri, Yuji, 1975, *Theory of Accounting Measurement* (American Accounting Association, Chicago, IL).
- 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, Paul R. Milgrom, and John Roberts, 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 John Roberts, 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 Eric S. Rosengren, 1996, Small business credit availability: How important is size of lender?, in A. Saunders and I. Walter, eds.: *Financial System Design: The Case for Universal Banking* (Irwin, Burr Ridge, IL).
- Peek, Joe, and Eric S. Rosengren, 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, 2002, Does distance still matter? The information revolution in small business lending, *Journal of Finance*, forthcoming.

- Radner, Roy, 1993, The organization of decentralized information processing, *Econometrica* 61, 1109–1146.
- Rajan, Raghuram G., Henri Servaes, and Luigi Zingales, 2000, The cost of diversity: The diversification discount and inefficient investment, *Journal of Finance* 55, 35–80.
- Rotemberg, Julio J., and Garth Saloner, 1994, Benefits of narrow business strategies, *American Economic Review* 84, 1330–1349.
- Rotemberg, Julio J., and Garth Saloner, 1995, Overt interfunctional conflict (and its reduction through business strategy), *RAND Journal of Economics* 26, 630–653.
- Sah, Raaj Kumar, and Joseph E. Stiglitz, 1986, The architecture of economic systems: Hierarchies and polyarchies, *American Economic Review* 76, 716–727.
- Sapienza, Paola, 2002, The effects of banking mergers on loan contracts, *Journal of Finance* 57, 329–367.
- Scharfstein, David S., and Jeremy C. Stein, 2000, The dark side of internal capital markets: Divisional rent-seeking and inefficient investment, *Journal of Finance* 55, 2537–2564.
- Stein, Jeremy C., 1997, Internal capital markets and the competition for corporate resources, *Journal of Finance* 52, 111–133.
- Strahan, Philip E., and James Weston, 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).
- Verrecchia, Robert E., 1983, Discretionary disclosure, *Journal of Accounting and Economics* 5, 179–194.
- Williamson, Oliver E., 1967, The economics of defense contracting: Incentives and performance, in R. McKean, 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.

