



American Finance Association

Internal Capital Markets and the Competition for Corporate Resources

Author(s): Jeremy C. Stein

Source: *The Journal of Finance*, Vol. 52, No. 1 (Mar., 1997), pp. 111-133

Published by: Blackwell Publishing for the American Finance Association

Stable URL: <http://www.jstor.org/stable/2329558>

Accessed: 25/11/2009 10:04

Your use of the JSTOR archive indicates your acceptance of JSTOR's Terms and Conditions of Use, available at <http://www.jstor.org/page/info/about/policies/terms.jsp>. JSTOR's Terms and Conditions of Use provides, in part, that unless you have obtained prior permission, you may not download an entire issue of a journal or multiple copies of articles, and you may use content in the JSTOR archive only for your personal, non-commercial use.

Please contact the publisher regarding any further use of this work. Publisher contact information may be obtained at <http://www.jstor.org/action/showPublisher?publisherCode=black>.

Each copy of any part of a JSTOR transmission must contain the same copyright notice that appears on the screen or printed page of such transmission.

JSTOR is a not-for-profit service that helps scholars, researchers, and students discover, use, and build upon a wide range of content in a trusted digital archive. We use information technology and tools to increase productivity and facilitate new forms of scholarship. For more information about JSTOR, please contact support@jstor.org.



Blackwell Publishing and American Finance Association are collaborating with JSTOR to digitize, preserve and extend access to *The Journal of Finance*.

<http://www.jstor.org>

Internal Capital Markets and the Competition for Corporate Resources

JEREMY C. STEIN*

ABSTRACT

This article examines the role of corporate headquarters in allocating scarce resources to competing projects in an internal capital market. Unlike a bank, headquarters has control rights that enable it to engage in “winner-picking”—the practice of actively shifting funds from one project to another. By doing a good job in the winner-picking dimension, headquarters can create value even when it cannot help at all to relax overall firm-wide credit constraints. The model implies that internal capital markets may sometimes function more efficiently when headquarters oversees a small and focused set of projects.

THIS ARTICLE ANALYZES THE process by which internal capital markets channel limited resources to different uses inside a company. In so doing, it seeks to address two related sets of questions. First, what is the fundamental economic rationale for creating an internal capital market? That is, under what circumstances can it make sense to combine several technologically distinct projects under one roof, and have them seek funding from corporate headquarters, as opposed to setting them up as stand-alone companies that each raise external financing on their own? Second, given this rationale, what is the optimal size and scope of an internal capital market? Should headquarters be involved in funding a large number of projects, or just a few? And should these projects be unrelated to one another, or in similar lines of business?

The answers to both sets of questions flow from the insight that in a credit-constrained setting—where not all positive NPV projects can be financed—headquarters can create value by actively reallocating scarce funds across projects. For example, the cash flow generated by one division’s activities may be taken and spent on investment in another division, where the returns are higher. Or alternatively, one division’s assets may be used as collateral to raise financing that is then diverted to the other division. Simply put, individual projects must compete for the scarce funds, and headquarters’ job is to pick the winners and losers in this competition.

* MIT Sloan School of Management and NBER. This research is supported by MIT’s International Financial Services Research Center and the National Science Foundation. Thanks to Maureen O’Donnell for help in preparing the manuscript. I am grateful for the comments and suggestions of Rob Gertner, Oliver Hart, Bengt Holmstrom, Louis Kaplow, Rick Ruback, David Scharfstein, René Stulz, Luigi Zingales, two anonymous referees, and seminar participants at Harvard and the NBER.

This insight is at the heart of Williamson's (1975) description of an internal capital market: "... cash flows in the M-form firm are not automatically returned to their sources, but instead are exposed to an internal competition . . . In many respects, this assignment of cash flows to high yield uses is the most fundamental attribute of the M-form enterprise . . ." (pp. 147–148). Donaldson (1984) sounds a similar theme: "The most critical choices top management makes are those that allocate resources among competing strategic investment opportunities." (p. 95).

One very interesting aspect of the Williamson (1975) and Donaldson (1984) view is that it suggests an interdependence among otherwise completely unrelated investment projects that just happen to be located under the roof of the same company. Specifically, the extent to which any given project gets funded in an internal capital market will depend not only on that project's own *absolute* merits, but also on its merits *relative* to other projects in the company's overall portfolio. Thus for example, if a company owns two unrelated divisions A and B, and the appeal of investing in B suddenly increases, the argument would seem to imply that investment in A would decline—even if it is positive NPV at the margin—as corporate headquarters channels relatively more of its scarce resources toward B.

Several recent articles provide evidence of investment interdependence. Lamont (1996) documents what might be termed "intracompany liquidity spillovers" in the U.S. oil industry: when major oil companies' cash flows were hard-hit by the oil price decline of 1986, they cut investment across the board, including investment in their *non-oil-related divisions*. In a similar vein but with a broader sample, Shin and Stulz (1996) find that the investment of small divisions of diversified firms is strongly related to the cash flows of other divisions.¹

To capture the notions of winner-picking and investment interdependence in a formal model, two types of assumptions are required. These assumptions are developed in detail below, but here is an overview: First, there must be binding credit constraints, both when individual projects seek financing on their own from external markets, and when headquarters attempts to raise funds for multiple projects simultaneously. These credit constraints are modelled in a conventional fashion, by assuming that both project managers and headquarters derive private benefits that increase with the resources under their control. Hence there is a tendency on the part of these agents to overstate their investment prospects. This in turn implies that less-informed outside providers of finance will be skeptical, and may ration funds to them.

The second key set of ingredients is needed to give headquarters both the incentive and the authority to engage in winner-picking. I follow Gertner, Scharfstein, and Stein (1994), henceforth GSS, in stressing the role that

¹ See also Lang, Ofek, and Stulz (1996) who demonstrate that the growth of small "noncore" divisions is negatively related to the leverage of their parent firms—i.e., there appear to be leverage spillovers across projects.

control rights play in making headquarters an effective intermediary.² In the current set-up, headquarters' control rights have two distinct consequences. First, they enable headquarters to capture a share of the private benefits generated by any projects it oversees. To the extent that private benefits are correlated with overall project profitability, this gives headquarters some incentive to favor better projects. Second, headquarters' control rights also allow it to play Robin Hood, in the sense of being empowered to *take resources away* from some projects in order to give them to other, more deserving ones. An important point is that with fixed resources, the flip side of winner-picking is "loser-sticking"—forcing some projects to accept a lower level of funding than they could obtain as stand-alones. As I argue below, its authority to engage in such loser-sticking makes headquarters a unique form of intermediary, unlike a bank or a venture capital partnership.

Given the assumptions, the first step in the analysis is to show how headquarters creates value in an internal capital market. This is done most simply by considering an example where the credit constraint facing headquarters takes an extreme form: headquarters is unable to bring any more total funding to bear than the individual projects it oversees could raise operating as stand-alones. Yet even in this extreme case, headquarters can still add value by allocating the same amount of funds more efficiently. This is winner-picking in its purest form, and it makes the contrast with a bank lender especially stark: since a bank is less able to engage in winner-picking, it can never create value unless it eases credit constraints for the projects it finances.

Having developed the basic rationale for an internal capital market, I next turn to the twin questions of optimal size and scope. The optimal size of an internal capital market turns on a straightforward tradeoff: on the one hand, by putting more projects under the same roof, headquarters may be better able to relax credit constraints, thereby raising *more* in total, than the individual projects could as stand-alones. On the other hand, this comes at the expense of headquarters' monitoring effectiveness—if it is required to oversee too many projects, it may no longer be able to discern the winners and losers as accurately.

The issue of optimal scope is a bit more subtle. Again, the question being asked is whether an internal capital market works better when the company adopts a strategy of unrelated diversification, or a focused strategy involving investment in similar lines of business. Traditional models of financial intermediation suggest a presumption in favor of unrelated diversification. But in this model, there is an effect pushing toward a focused strategy. This effect arises because, in an internal capital market, headquarters is primarily involved in allocating a fixed pool of resources, and hence in ranking projects on a relative basis. When doing such relative ranking, it does not matter if headquarters makes absolute errors in evaluating projects, as long as these errors are correlated across the projects.

² The emphasis on control rights distinguishes both this article and GSS from Harris and Raviv's (1994) recent work on the internal capital allocation process.

The remainder of the article is organized as follows. Section I lays out a standard project-level agency problem, and shows how it can lead to credit rationing and underinvestment when individual projects attempt to raise financing as stand-alones in the arm's length external capital market. Section II uses a simple example with two independent projects to illustrate the basic winner-picking role of headquarters in an internal capital market. In Section III the logic of the model is extended to endogenize both the number of projects and their correlation structure—i.e., to pin down the optimal size and scope of an internal capital market. Section IV discusses some implications of the theory, including what it has to say about the contrast between headquarters and other intermediaries such as bank lenders. Section V concludes.

I. The Basic Project-Level Agency Problem with External Financing

A. Assumptions

I begin by describing the technology and information structure for a single investment project that is financed in the arm's length external market.³ The original idea for the project comes from a founder. To give this idea a concrete embodiment as an asset that is valuable and for which there may be meaningful control rights, one can think of it as represented by a patent. In order for the project to move forward, the idea must be combined with managerial labor and capital investment. Unfortunately, aside from the patent, the founder has no wealth, nor does he have any managerial ability. This implies that the project will require the services of both a professional manager—who also has no wealth of his own—as well as funding from a group of financiers. Throughout the article, I assume that the objective function is to maximize the expected proceeds that the founder derives from his idea, net of any payments to the financiers, or wages to the manager. For simplicity, the rate of return required by the financiers is set to zero, as is the reservation wage of the manager.

The amount of investment in the project can be either one (1) or two (2). The return on this investment depends on the state of the world. When the state is B (for “bad”) the project returns a gross amount of cash flow y_1 if an amount 1 was invested, and returns y_2 if an amount 2 was invested. The cash flows y_1 and y_2 are assumed to be costlessly verifiable. This implies that the founder and the financiers have no difficulty getting their hands on these cash flows. Furthermore, it is assumed that $1 < y_1 < y_2 < 2$. In other words, from the perspective of the founder, the optimal level of investment in state B is just 1.

When the state of the world is G (for “good”), investment is more productive. Specifically, output is raised by a factor $\theta > 1$, so that the proceeds from investing 1 are now θy_1 and the proceeds from investing 2 are now θy_2 . It is assumed that $\theta(y_2 - y_1) > 1$, so that from the perspective of the founder, the optimal level of investment in state G is 2.

³ By “arm's length,” I mean to suggest that there are a large number of diffuse investors, so it is difficult for them to engage in any sort of monitoring activity. Thus a natural interpretation is to think of the arm's length market in this section as representing the bond market.

The ex ante probability of state G occurring is p ; the probability of state B is therefore $(1 - p)$. Project managers observe the state of the world directly, but the founder and the financiers do not. In light of this information asymmetry, I will sometimes refer to the uninformed founder and financiers collectively as “outside investors.” Clearly, if outside investors wish to make investment contingent on the state, they must rely on managers to report the state to them. However, managers have a tendency to make overoptimistic reports, because they derive private benefits that increase with the amount of resources under their control.

To model management’s inclination toward overinvestment in a simple way, I assume that these private benefits are directly proportional to the cash flow proceeds from investment. Thus if the state is B, the private benefits are sy_1 and sy_2 for investment levels of 1 and 2 respectively, where s is a fixed constant. Similarly, if the state is G, the private benefits are $s\theta y_1$ and $s\theta y_2$ for investment levels of 1 and 2 respectively. These private benefits are assumed to be unverifiable; this is what makes them “private,” in the sense that they cannot be contractually assigned to outside investors.

These assumptions capture a very basic and plausible type of financing problem. When approaching outside investors for financing, managers will have a tendency to overstate their investment prospects, in order to raise more money. Knowing this, the outside investors may be forced to ration funds to managers, never giving them more than 1 unit of financing for the project. This has the benefit of preventing overinvestment when the true state of the world is B, but it also has the cost of causing underinvestment when the true state is G.

B. Credit Rationing When There is No Revelation of Information

To see how rationing might occur, let us first consider the case where no revelation schemes are used to elicit management’s private information. (I will discuss when such revelation schemes are endogenously ruled out momentarily.) In this case, investment cannot be made contingent on the state of the world. That is, the financiers can either always provide 1 unit of financing, or always provide 2 units of financing.

In the former scenario, the expected net cash flows are $(p\theta + (1 - p))y_1 - 1$. Since the first unit of investment is always positive NPV, regardless of the state, this quantity is guaranteed to be positive. This means that it will always be possible to raise 1 unit of financing, by contractually assigning a large enough portion of the cash flow benefits to the financiers. For example, one simple contract—though by no means the only one—that does the trick is a debt contract with a promised repayment of 1. This contract can be serviced regardless of the state of the world, and thus allows the financiers to break even, with the founder keeping the rest.

In the latter scenario, where investment is 2 in all states, the expected net cash flows are $(p\theta + (1 - p))y_2 - 2$. Given the assumptions made above, investing 2 units will yield less on net than investing 1 unit if the probability

p of state G occurring is sufficiently small. From this point on, I will assume that this is the case. It follows that from the founder's perspective, the optimal uncontingent financing plan involves always allocating 1 unit to the project.⁴ The bottom line is that without revelation schemes, the information/agency problem can lead to a situation where credit is rationed, and where investment is inefficiently low in state G .

Fundamentally, credit rationing arises in this model because managers' private information is not exploited, and hence, investment is not conditioned on the state. This raises the question of whether one might be able to use an incentive scheme to elicit managers' information, and thereby make better investment decisions. In general, this is possible, but not costless. Depending on how costly it is to elicit the information, it may or may not be desirable to do so. The costs of eliciting information are a function of the private benefit parameter s . Intuitively, to get a manager to truthfully reveal that the state is B , and that investment of only 1 unit is warranted, the manager has to be given a bonus that offsets the incremental private benefits that he passes up by not being able to invest 2 units.⁵ In order to simplify the exposition, I will for the remainder of the article focus my attention on those cases where private benefits are sufficiently large to rule out the usefulness of revelation schemes.

II. The Role of Corporate Headquarters in an Internal Capital Market

A. Assumptions about Headquarters' Authority, Incentives, and Monitoring Capabilities

The analysis above illustrates how information and agency problems in external finance can lead to credit rationing and inefficient investment. However, the main focus of the article is to ask whether these problems can be partially alleviated by financing projects via an internal capital market. In an internal capital market, individual project managers no longer raise funds directly from outside investors. Rather, there is an additional agent designated to act as a go-between; namely "corporate headquarters." It is now headquarters that deals with the outside market, and that in turn passes along the funds it raises to individual projects.

I make four important assumptions about the agent that is designated to serve as headquarters:

⁴ The founder can implement this plan in a number of ways. First, one can imagine that in the "charter" for the project, he inserts a clause restricting total outside financing to 1 unit. To the extent that the amount of outside financing is verifiable, this is a feasible approach. Alternatively, without such a contractual clause, one would obtain the same limitation if the founder can negotiate frictionlessly with the financiers, since it is in the joint interest of this group of outside investors to maximize verifiable net cash flows.

⁵ More precisely, I demonstrate in the NBER working paper version (Stein (1995)) that the ex ante costs of eliciting information from the manager are $(1 - p)s(y_2 - y_1)$. The benefits are that one can now invest 2 in state G instead of just 1, which is worth $p(\theta(y_2 - y_1) - 1)$. This implies that it will not be worth using a revelation scheme if p is relatively small or s is relatively large.

ASSUMPTION 1: *Headquarters has monitoring skills that enable it to acquire (possibly noisy) information about projects' ex ante prospects.*

ASSUMPTION 2: *It has no financial resources of its own.*

ASSUMPTION 3: *It is endowed by the founder with control rights which allow it to capture a fraction ϕ of the private benefits associated with any project it oversees; however, this surplus extraction by headquarters comes at the cost of diluting project managers' incentives.*

ASSUMPTION 4: *Headquarters' control rights also give it the authority to redistribute resources across projects.*

Assumption 1 is self-explanatory, and is obviously essential to any model in which headquarters performs a meaningful allocative role. Assumption 2 implies that headquarters may potentially face the same sort of financial constraints as the project managers. It has no wealth of its own, and it may also—as will become clear shortly—face credit rationing by outside investors. In other words, to the extent that headquarters ultimately creates any value in this model, it is not simply because a benevolent deep pocket has been inserted as a *deus ex machina* to directly alleviate financing constraints.

Assumptions 3 and 4 both reflect the idea that headquarters is endowed with residual rights of control (in the sense of Grossman and Hart (1986) and Hart and Moore (1990)) over the initial patent and any subsequent project assets. This contrasts with the case of external financing, where project managers are endowed with control rights.⁶ In the current context, this allocation of control rights implies two distinct sorts of consequences.

First, as Assumption 3 states, giving control rights to headquarters allows it to effectively appropriate from project management a fraction $\phi < 1$ of the private benefit.⁷ Thus, for example, if publicly verifiable cash flows are y , and total private benefits are sy , headquarters now gets to keep ϕsy , leaving only $(1 - \phi)sy$ for project management. As in Grossman and Hart (1986), this *ex post* opportunism by the controlling party has *ex ante* costs. In particular, by taking away a share of project managers' private benefits, headquarters reduces their *ex ante* incentives to put forth effort. This in turn reduces invest-

⁶ In both cases, one can think of the founder as making an initial decision as to who will be assigned control rights—headquarters or the project manager. The founder will make that choice which maximizes expected net (verifiable) cash flows. For simplicity, I am assuming that headquarters is completely distinct from the founder. This might be rationalized by arguing that the founder does not have the skills needed to do headquarters' job—i.e., to engage in monitoring—and thus must engage an outside professional. I am also assuming that headquarters is distinct from the individual project managers. Thus, in a multiple-project setting, I am ruling out the potentially interesting asymmetric possibility that, say, the manager of project i will be given control of both project i and project j , and therefore be empowered to deal with the outside market on behalf of the two projects. Again, this might be rationalized by arguing that managers of individual projects have their hands so full that they would be unable to simultaneously monitor other projects.

⁷ What I have in mind here is that headquarters' control rights give it the power to make certain threats—e.g., to fire the project manager irrespective of his performance—that enable it to extract some of the project manager's surplus.

ment output. To capture this effort-dilution in a simple reduced-form way, I assume that when a project is overseen by headquarters, all cash flows fall by a factor $k < 1$ relative to the case of manager control. Thus when the state is B, investment is 1, and headquarters is involved, verifiable cash flows are ky_1 , headquarters gets a private benefit of ϕsky_1 , and the project manager gets a private benefit of $(1 - \phi)sky_1$. Similarly, when the state is G, investment is 1, and headquarters is involved, verifiable cash flows are $k\theta y_1$, headquarters gets a private benefit of $\phi sk\theta y_1$, and the project manager gets a private benefit of $(1 - \phi)sk\theta y_1$. Analogous notation applies when the investment level is 2, with y_1 replaced everywhere by y_2 .

The second consequence of headquarters' being endowed with control rights is a bit more novel, and is captured by Assumption 4. This assumption is relevant only in the case where headquarters oversees multiple projects. In this case, I assume that its control rights give headquarters the authority to choose the level of funding for individual projects. Thus if headquarters initially controls n patents, and has somehow managed to use these patents as "collateral" to raise n units of financing from the outside market, it has a number of choices: it can give each of the projects exactly 1 unit of financing; or it can give m of them 2 units, another m of them no funding whatsoever, and $(n - 2m)$ of them 1 unit. In other words, headquarters can divide up any resources it has been able to raise unevenly across the projects if it so chooses.

The authority to redistribute resources across projects distinguishes headquarters from a nonowner intermediary such as a bank. By redistributing, headquarters may wind up giving some projects less financing than they could obtain as stand-alones. For example, headquarters may draw on the collateral value of patent i to obtain funds, but then pass these funds on to another project j , leaving the manager of project i with nothing to do. A bank cannot do this, because a bank does not have the right to prevent any given project manager from taking the patent that he controls and going to the outside market to seek a more favorable deal. In contrast, I am assuming that headquarters' control rights enable it to impose exclusivity of this sort on project managers.⁸ This assumption is very much in the spirit of Hart and Moore (1990), who write that: "... the sole right possessed by the owner of an asset is his ability to exclude others from the use of that asset." (p. 1121).

B. The Role of Headquarters with Multiple Projects: A Two-Project Example

Given the assumptions, it is clear that creating a headquarters is always value-reducing in a one-project setting. First of all, the dilution of project managers' incentives means that for any given level of investment, output will be reduced by a factor k . Now, if headquarters could somehow be used as a mechanism for monitoring and communicating information to the outside market, and thereby easing credit constraints, perhaps this cost could be

⁸ See Gertner (1986) for a different perspective on the significance of such exclusivity in an internal capital market.

offset. However, headquarters will be no more inclined to communicate truthfully with the outside market than will project-level managers.

To see this, suppose that headquarters' monitoring uncovers that the true state is B. It will not honestly report the state, because, like project management, its private benefits are proportional to output, and hence it always prefers more investment to less. The bottom line is that headquarters will face exactly the same sort of credit constraints that the project manager would if he dealt directly with the market—headquarters will never be able to raise more than 1 unit of outside financing, regardless of the true state.⁹

This result is reminiscent of Diamond's (1984) analysis of financial intermediation. Like I do, Diamond explicitly considers the fact that interposing an intermediary between the capital market and individual projects creates a second layer of agency. This leads him to the same conclusion—that an intermediary is strictly dominated when it only finances a single project.¹⁰

The bottom line is that in order for an additional agent—i.e., headquarters—to have any value in the current framework, it must be involved in financing multiple projects. To see the logic most clearly, it is useful to start with a very simple example. Suppose that there are just two projects, i and j , each of which is exactly as described in Section I above. The only difference between the projects is that their productivity shocks are not perfectly correlated. Indeed, the projects are completely independent, so that the probability of project i being in state G is p , regardless of which state project j is in, and vice versa. Finally, when headquarters monitors either of the projects, it can observe the state perfectly.

For the time being, I consider a case where putting the two projects together under one roof does not ease overall credit constraints. That is, based on the two patents it controls, headquarters always raises exactly 2 units of funding from the outside market, regardless of its information. As will become more clear shortly, this may be a conservative assumption. Given the equilibrium behavior that emerges in the model, outside investors will always provide at least 2 units of funding. Moreover, there will be in general some circumstances in which headquarters will be able to raise *more* than the sum of what the individual projects could raise. I tackle this more general case in Section III.A below.

⁹ This is true so long as one continues to disregard the potential for revelation schemes in which incentive payments are made to managers based on their investment decisions and ultimate realized cash flows. If such schemes are relevant, then headquarters may have an advantage in fund-raising in a one-project setting. This is because headquarters only receives a *fraction* of the private benefits associated with overinvestment; thus it will be less costly to devise an incentive scheme that discourages headquarters from overstating its prospects. I ignore this issue in what follows.

¹⁰ It should be noted, though, that the conclusions here are very different from those in GSS. In that article, an internal capital market could add value even when financing only a single project. This is because in GSS, we assumed away credit constraints at the level of corporate headquarters. Similarly, in Harris and Raviv (1994), where headquarters faces no credit constraint, it is able to play a valuable role even in a one-project setting.

Given that an internal market does not have more resources to work with, the only way that it can improve things is if it allocates these resources more efficiently. Thus, for example, if project i is in state G, and project j is in state B, headquarters might invest 2 units in i , and nothing in j . In order for the problem to be interesting, I assume that this is a more efficient allocation than investing an equal amount in both projects. In other words, the marginal return on the second unit invested in i must be greater than the marginal return on the first unit invested in j , or $\theta(y_2 - y_1) > y_1$.

If this condition is satisfied, it turns out that headquarters will indeed have the proper incentives to allocate its fixed resources to the right projects. This is simply because headquarters' private benefits are directly proportional to project cash flows. Thus in attempting to maximize its private benefits, headquarters will also pick the mix of investment that maximizes verifiable cash flow value.

Although the model is extremely stylized, a general point emerges: while headquarters will want to overinvest on average in all of its projects—and hence will face credit constraints from the outside market—it still has an incentive to channel the funds it obtains toward “winners” and away from “losers.” (I should be a bit careful with the terminology here, since even the “loser” state B project is positive NPV when it receives 1 unit of investment.) The crucial distinction between headquarters and an individual project manager is that *headquarters' broader span of control allows it to derive private benefits from several projects simultaneously*, while a project manager can only derive private benefits from a single project. This implies that headquarters will sometimes be willing to take funds away from a weaker project in order to give more to a stronger one, whereas the manager of the weaker project would never do so.

It is straightforward to calculate the ex ante gain (from a founder's perspective) associated with moving from individual project financing in an external market to an internal capital market presided over by headquarters. In an external market, the total expected net output from the two projects, denoted by EM, is given by:

$$EM = 2(y_1(p\theta + (1 - p)) - 1) \quad (1)$$

Equation (1) reflects the fact that in an external market, each of the projects always invests exactly 1. In contrast, in an internal market, the total expected net output, denoted by IM, is given by:

$$IM = 2(1 - p)^2ky_1 + 2p^2k\theta y_1 + 2p(1 - p)k\theta y_2 - 2 \quad (2)$$

Equation (2) differs from (1) in two ways: First, there is the effort-dilution factor k , which tends to make IM less than EM. Second, when the projects are in different states—which happens with probability $2p(1 - p)$ —the project which is in state G receives 2 units of investment, while the project in state B receives no investment at all. This improved allocation of resources tends to

make IM greater than EM. When the latter effect is sufficiently strong, an internal capital market administered by headquarters can be the preferred mode of financing. Again, this is because headquarters has the right incentives to act on the information generated by monitoring. It is in headquarters' private interest to channel funds toward projects that are winners and away from those that are losers.

Of course, I have rigged things somewhat by making headquarters' private benefits directly proportional to project cash flows. This implies that, for a fixed amount of resources, headquarters' self-serving interests will coincide *exactly* with value maximization. The basic purpose of this assumption is to capture a notion that is implicit in much of the agency literature: managers have a general tendency to overinvest, but nonetheless, when ranking projects, they tend to prefer higher NPV ones to lower NPV ones. This seems to be what Donaldson (1984) has in mind when he argues that managers seek to maximize "corporate wealth."¹¹ And it is exactly the sort of premise that has led Jensen (1993) and others to argue that plentiful internal cash flow is more likely to create a divergence between managerial and shareholder interests when companies have relatively few attractive investment opportunities.¹²

Clearly, the assumption that private benefits and verifiable cash flows are exactly proportional is more extreme than it needs to be, and is only used to make things more transparent. One can easily generalize the model to consider the case where there is only an imperfect correspondence between private benefits and verifiable cash flows. It is straightforward to show that an internal capital market can sometimes add value in this case as well, although it typically adds less. Now there will be times when a particular project is selected for favorable funding status even when its cash flows do not merit this, simply because it offers disproportionately high private benefits to headquarters.¹³

¹¹ Donaldson's definition is as follows: "Corporate wealth is that wealth over which management has effective control. . . ." For an elaboration of this point and its significance for optimal capital structure, see also Myers (1993).

¹² Louis Kaplow has suggested an alternative formulation that leads to similar results. Suppose that private benefits are proportional to the amount invested rather than the cash flow produced, so that based on private benefits alone, managers care purely about the scale of investment, and are completely indifferent to its NPV. Then if one introduces a small amount of stock-based incentive compensation, this will tend to make managers favor higher NPV projects over lower NPV ones, even if the incentive compensation is not enough to eliminate the general bias toward overinvestment.

¹³ Such a generalization could be interesting in its own right, as it might generate predictions about the optimal portfolio of projects for headquarters to oversee. Specifically, if one allows for the possibility that the "private benefit appropriation technology" (as represented by the parameter s) varies systematically across projects or industries, this would suggest that all else equal, resources will be allocated more efficiently when headquarters restricts itself to a set of projects with similar values of s . This sort of reasoning could ultimately be developed into an argument for focused corporate strategies that is distinct from, but complementary to, that presented in Section III.B below.

III. The Nature of an Internal Capital Market: Optimal Size and Scope

In the simple example above, both the number of projects (two) and their degree of correlation (zero) were taken as exogenous. In this section, I investigate whether the basic logic of the model can be used to endogenize both of these variables, thereby delivering implications for the optimal size and scope of an internal capital market. To keep the analysis manageable, I treat the two questions separately. That is, I tackle the size issue while keeping the projects statistically independent, and then address the issue of optimal scope (i.e., the correlation structure of the projects) while holding the number of projects fixed.¹⁴

A. Size: The Tradeoff between Tightness of Credit Constraints versus Monitoring Quality

Given the minimal ingredients that are in the model so far, it would seem that the larger an internal capital market, the better. This is because of an “ease the credit constraint” effect, which I have been ignoring to this point. There are two aspects to this phenomenon. First, for any fixed number of projects, headquarters may, by virtue of its greater allocative efficiency, be able to raise more total resources from the outside market than the individual project managers. For example, even if headquarters is restricted to overseeing only two projects, there are some parameter values for which it can raise 3 units of funding, even when the individual projects can just raise 1 unit each.

Second, as the number of projects is increased, efficiency along this dimension is further enhanced. To see how, take an extreme example where headquarters oversees 1000 independent projects. In this case, the law of large numbers tells us that we can count on having roughly $1000p$ projects in state G, and $1000(1 - p)$ projects in state B. Thus if headquarters is given $1000(1 + p)$ units of financing to work with, it will allocate 2 units each to the projects that are in state G, and 1 unit each to the projects that are in state B. Investment will therefore converge to the first-best level, and credit constraints will no longer be binding.¹⁵

This line of argument suggests that the model as developed thus far—with just two projects—is not fully specified. Another ingredient is required to pin down the optimal scale of an internal capital market, if one wants to avoid the counterfactual implication that the firm should ideally be made arbitrarily large. Perhaps the most natural way to close the model is to acknowledge that headquarters’ monitoring efforts will be compromised when it oversees a large

¹⁴ A couple of recent articles—also set in an incomplete contracting framework—investigate other reasons why firms might wish to stay either small or focused. See Meyer, Milgrom, and Roberts (1992) and Rotemberg and Saloner (1994).

¹⁵ Although the amount of investment will be equal to the first-best level in this example, output will be lower, because of the effort-dilution parameter k .

number of projects.¹⁶ As a simple device to capture this notion, I make the following assumption: when headquarters oversees n projects, there is a probability $M(n)$ that its monitoring will be successful, and that it will be able to learn the true state of each of the n projects. With probability $(1 - M(n))$, monitoring is unsuccessful, and headquarters learns nothing about any of the projects. $M(n)$ is a decreasing function.¹⁷

With this one addition, the model now delivers two new variables endogenously— n , the optimal number of projects overseen by headquarters, and F , the optimal amount of total resources allotted to these projects. Although it is difficult to obtain closed-form expressions for the optimal values of n and F as a function of primitive parameters, it is straightforward to identify them in any given case. The logic is as follows. First, pick a number of projects n . This implies a value of $M(n)$. Then pick an integer level of total funding F . Denote by $\pi^M(n, F)$ the per-project profits if monitoring is successful and headquarters learns the true state of each project. Similarly, denote by $\pi^N(n, F)$ the per-project profits if monitoring fails and headquarters learns nothing. Then the ex ante expected profit is given by

$$\pi(n, F) = M(n)\pi^M(n, F) + (1 - M(n))\pi^N(n, F) \quad (3)$$

This value can be directly (if sometimes tediously) calculated for any values of n and F . Holding fixed n , one first optimizes over F —i.e., one identifies the optimal level of funding $F^*(n)$ for an internal capital market of given size. Finally, it remains only to pick the value of n that maximizes $\pi(n, F^*(n))$.

A couple of examples illustrate the effects at work.¹⁸ One simple piece of intuition is that improvements in the monitoring technology might make it more attractive to have larger internal capital markets. Indeed, this effect arises for a wide range of parameter values.¹⁹

Example 1: Set $y_1 = 1.01$; $y_2 = 1.68$; $\theta = 2$; $p = 0.4$; and $k = 1$. For these values, it is easy to show that no matter what $M(2)$ is, $F^*(2) = 2$. That is, credit constraints will never be eased with two projects, no matter how good the monitoring technology. However, $F^*(3) = 4$ if $M(3) > 0.538$. (Otherwise, $F^*(3) = 3$.) In other words, credit constraints can be eased by increasing the number of

¹⁶ A more sophisticated approach would be to allow for the possibility that headquarters can add a larger staff, with different staff members delegated to the task of monitoring different subsets of the firm's projects. Of course, to model this satisfactorily, one would have to explicitly take account of the added layers of agency problems that arise in such an environment. See Section V for a further discussion of how such a model might work.

¹⁷ This "all or nothing" formulation has the advantage of preserving ex post symmetry across the projects—there are never some projects about which headquarters is better informed than others. This simplifies the analysis considerably, without, I suspect, changing any of the important qualitative conclusions.

¹⁸ In both of the examples below, I focus only on the choice between $n = 2$ and $n = 3$, implicitly assuming that $M(4)$ is sufficiently small as to always make $n = 4$ a suboptimal choice.

¹⁹ This effect is trivially present in the extreme case where the monitoring technology becomes perfect, so that $M(n) = 1$ for all n . In this case, the optimal n converges to infinity, and investment approaches the first-best level, as suggested above.

projects under management to three, so long as the monitoring technology still works sufficiently well with three projects.

Suppose initially that $M(2) = 0.6$ and $M(3) = 0.5$. In this case, there is no easing of credit constraints with either two or three projects. This implies that a two-project setup is strictly preferred, since it involves better monitoring, and the same average level of per-project financing. In particular, it can be shown that $\pi(2) = 0.461$, while $\pi(3) = 0.454$.

Now consider improving the monitoring technology, so that $M(2) = 0.9$, and $M(3) = 0.8$. This raises $F^*(3)$ to 4, so that in the three-project setup, the average level of per-project funding is increased. This effect is strong enough to overcome the fact that monitoring works less well with three projects, with the net result being that the optimal size of the internal capital market increases—we now have $\pi(2) = 0.485$ and $\pi(3) = 0.487$.

While this basic effect seems intuitively plausible, it is not a general feature of the model that improvements in the monitoring technology always have a monotonic effect on the optimal size of an internal capital market. This is because there is an offsetting force at work: when the monitoring technology improves, it may make it possible to generate a substantial easing of credit constraints with a smaller number of projects. Thus it becomes less important to add projects in an effort to boost the level of per-project funding.

Example 2: Set $y_1 = 1.1$; $y_2 = 1.9$; $\theta = 1.5$; $p = 0.4$; and $k = 1$. The key difference in this example is that there can be easing of credit constraints with only two projects—in particular, $F^*(2) = 3$ if $M(2) > 0.833$. Also, $F^*(3) = 4$ if $M(3) > 0.361$.

If we start out with a relatively weak monitoring technology, with $M(2) = 0.4$ and $M(3) = 0.3$, neither a two-project nor a three-project setup eases credit constraints, so as before, a two-project setup is preferred on the grounds of better monitoring, with $\pi(2) = 0.330$ and $\pi(3) = 0.327$. If the monitoring technology is improved to the point where $M(2) = 0.8$ and $M(3) = 0.6$, it becomes possible to ease credit constraints, but only with three projects. This implies that a three-project setup is preferred, with $\pi(2) = 0.339$ and $\pi(3) = 0.343$.

The nonmonotonicity can be seen by improving the monitoring technology even further, to the point where $M(2) = 1$ and $M(3) = 0.65$. Now the monitoring technology is so good that it is possible to relax credit constraints even with two projects. Given the configuration of parameters, it is thus better to have only two projects, with $\pi(2) = 0.348$ and $\pi(3) = 0.346$.

B. Scope: Focus versus Unrelated Diversification as Corporate Strategies

The next question to be addressed is whether, holding size fixed, an internal capital market works more efficiently when it adopts a strategy of unrelated diversification or a “focused” strategy involving closely related projects. To keep the formal analysis simple, I will focus on the case considered in Section II above, where there are two projects, i and j , and where headquarters always has just 2 units of total funding available. In so doing, I will be abstracting

away from any ease-the-credit constraint effects that might be associated with diversification.²⁰ The only remaining issue is whether by choosing a focused strategy, headquarters can do a better job in the pure winner-picking dimension.

B.1. A Trivial Argument in Favor of Diversification: Uncorrelated Project Outcomes

Even keeping the total resources available to headquarters fixed, there is another reason why diversification might make sense. Suppose the model is generalized to allow project outcomes to be correlated, while maintaining the assumption that headquarters can predict these outcomes without error. Then if one thinks of diversification as just making project outcomes less correlated, it follows immediately that the value of an internal capital market increases with the degree of diversification. The intuition is most easily seen in the limiting case of “perfect focus” where the projects are perfectly correlated, so that if i is in state G , j is always in state G as well, and vice-versa. In this case, winner-picking is impossible, because one project is never better than the other. Consequently, an internal market will always be dominated by the external market, so long as the effort-dilution parameter k continues to be less than one.

While this argument makes logical sense, and may capture part of what is relevant, it is clearly a bit superficial for two reasons. First, one might question the somewhat simplistic link that is drawn between project correlation and the notion of focus. One can imagine situations in which a company adopts what anybody would reasonably call a focused strategy, yet there remains a good deal of independence across projects. For example, consider a pharmaceutical company that invests in research and development (R&D) for several different new drugs. This sounds like a relatively focused strategy, yet the outcomes likely have large idiosyncratic components. In this case, there will still be ample room for headquarters to engage in winner-picking. A second weakness of this argument in favor of diversification is that it relies to some degree on the assumption that project outcomes can be predicted without error. As will be seen below, when this assumption is relaxed, the conclusion can be reversed.

B.2. A More Subtle Argument in Favor of Focus: Correlated Project-Evaluation Errors

As described in the introduction, the basic intuition for why a focused strategy might make more sense than unrelated diversification is straightforward. Given that headquarters’ job is to assess the *relative*, rather than absolute merits of projects, no harm will come from errors in judgment, so long

²⁰ In other words, I am setting aside the possibility that with two projects, headquarters might have a better chance of raising three—rather than two—units of funding by adopting a strategy of unrelated diversification.

as these errors are correlated across the projects. And a focused strategy is more likely to involve correlated project-evaluation errors than a strategy of unrelated diversification. To capture this intuition formally, I need to generalize the model slightly. There are three key features that must be added: 1) headquarters' signals are no longer perfectly accurate—i.e., there are prediction errors that arise when attempting to forecast project outcomes; 2) these prediction errors can be correlated across projects; and 3) the more focused the strategy, the more highly correlated are the prediction errors.

The specifics of the information structure are as follows. First, the unconditional probability of either project i or project j being in state G remains at p . Also as before, project *outcomes* continue to be independent, so that $\text{prob}(G^i/G^j) = p$ also. This independence will continue to hold regardless of the degree of focus—in the spirit of the pharmaceutical company example above, I am assuming that even with a focused strategy, there are some projects that are winners, and some that are losers. Headquarters no longer observes the project outcomes directly, but instead gets a noisy signal. For either project, the signal can be either H (for “high”) or L (for “low”). It is assumed that $\text{prob}(H^i/G^i) = \text{prob}(L^i/B^i) = q$, with $\frac{1}{2} < q < 1$, and analogously for project j . The larger is q , the more informative is the signal.

As suggested above, a crucial ingredient is that prediction errors are correlated. To capture this, I assume that when project j gets a “false” low (high) signal, project i is more likely to get a false low (high) signal as well. Thus:

$$\text{prob}(L^i/G^i, G^j, L^j) = (1 - q)(1 + \alpha) > \text{prob}(L^i/G^i) = (1 - q). \quad (4)$$

Analogously,

$$\text{prob}(H^i/B^i, B^j, H^j) = (1 - q)(1 + \alpha) > \text{prob}(H^i/B^i) = (1 - q). \quad (5)$$

Here α is a parameter between 0 and 1 that is a measure of the degree of focus—the more focused is the strategy, the greater is α .²¹

The rationale behind such an interpretation is as follows. When headquarters monitors two closely related projects, any misjudgments it makes about one project are likely to also show up in its assessment of the other project. For example, in the drug-development case sketched above, headquarters may simply be too bullish on drugs in general, because it underestimates the effects of future changes in health-care regulation on the prices it can charge. In contrast, if the projects under consideration are unrelated, it is less likely that the same misjudgment will be made across the projects. Thus if headquarters is choosing between allocating funds to drug R&D or aircraft carrier production, there is no reason to believe that overoptimism on the prospects for drugs will be matched by a similar misjudgment of aircraft carrier prospects.

²¹ This sort of information structure with correlated prediction errors is a necessary ingredient in tournament-style models—see Lazear and Rosen (1981), Green and Stokey (1983), Nalebuff and Stiglitz (1983), and Shleifer (1985). It is also a key feature of the herding model of Scharfstein and Stein (1990).

Given the information structure, there is still one more probability that needs to be specified. While I have assumed that a false low (high) signal on project j increases the probability of a false low (high) signal on project i , I *do not* assume that such a false low (high) signal on project j increases the odds of a *correct* low (high) signal on project i . That is, I assume that

$$\text{prob}(L^i/B^i, G^j, L^j) = \text{prob}(L^i/B^i) = q \quad (6)$$

and analogously that

$$\text{prob}(H^i/G^i, B^j, H^j) = \text{prob}(H^i/G^i) = q. \quad (7)$$

The purpose of this last simplification is just to ease the notational clutter; one can relax it without altering the basic results.

With these assumptions in place, all the various conditional probabilities can now be pinned down, either by Bayes' rule or by the requirements that certain probabilities sum to one. There are a total of 16 possible configurations. The Appendix enumerates these configurations, and for each lists: 1) the probability of it occurring; 2) the associated payoff with an external market that cannot reallocate funds across projects; and 3) the associated payoff with an internal market that can reallocate funds. To clarify the notation, the configuration GGHH refers to the case where both projects have G (good) outcomes and both signals were H (high). For simplicity, the effort-dilution parameter k is set to one in the Appendix.

Given the information in the Appendix, it is straightforward to demonstrate:

PROPOSITION: *The value created by an internal capital market is an increasing function of the focus parameter α . The positive effect of α is more pronounced when q is smaller—i.e., when headquarters' signals are less accurate.*

The intuition behind the proposition is apparent from the table in the Appendix. Increasing focus raises the probability weight on configurations like GGHH and BBHH. At the same time, it lowers the probability weight on configurations like GGHL and BBHL. On the one hand, this implies that when there is a high degree of focus, and the signal is mistaken for one project, it is more likely to also be mistaken for the other project. The configuration BBHH is an example of this—BBHH is a scenario where both signals are misleadingly positive, and it occurs more frequently as α is raised. The flip side of this is that with more focus, there is less likely to be an outcome like BBHL, where the signals are different but the projects are actually equally attractive.

The reason that this shifting around of probability weights adds value is because getting two misleading signals simultaneously—as in BBHH—does no harm. Both projects each receive one unit of funding, which is the correct allocation of resources. In contrast, getting different signals when the projects are equally attractive—as in BBHL—*does* lead to a misallocation of resources. Of course, all this only matters to the extent that signals tend to be mistaken in the first place. This is why focus matters more when q is lower—as q

approaches one, the signals become perfectly accurate, and the entire issue of correlated prediction errors becomes moot.²²

B.3. Pulling it Together: the Optimal Degree of Focus

Taken together, the arguments above suggest that there are two opposing forces at work. On the one hand, to the extent that i) focus increases the correlation of *project outcomes*; and ii) these outcomes can be predicted with relatively little error, then a focused strategy will tend to make an internal capital market less valuable. On the other hand, to the extent that i) there are errors in assessing projects; and ii) focus increases the correlation of these *errors*, then a focused strategy will tend to make an internal capital market work better.

It should be emphasized that in discussing both the optimal size and scope of an internal capital market, I have been posing a normative set of questions: how *should* a founder set up a firm in order to maximize expected profits? Of course, the implicit assumption that the founder can completely control the number and variety of projects undertaken by a firm is unrealistic. These decisions are at least in part made by managers, who have different interests. For example, it is clear that if headquarters wants to maximize its private benefits, it will prefer a larger size n than the founder. It is also possible that for any given size, headquarters will be more inclined than the founder to a policy of unrelated diversification.²³

This distinction between the normative and positive perspectives should be borne in mind when interpreting the empirical evidence. If founders make all the decisions, any unrelated diversification actually observed would be presumed to be value-enhancing. In contrast, if managers make the decisions, we may observe diversification strategies that are value-reducing. And indeed, the message that emerges from a large empirical literature on the subject is that the value consequences of unrelated diversification are at best highly suspect.²⁴ One interpretation of this evidence is that the normative argument made for focus in this article is correct, but that in reality, diversification decisions are not made solely in the interests of outside shareholders.

²² Note however, that if q becomes *too* low, the signals are so unreliable that it will no longer be optimal for an internal capital market to make resource allocation decisions based on them. Thus the comparative statics result on q referred to in the proposition only really applies in a region where q is sufficiently high to make winner-picking viable in the first place.

²³ This might arise either because: i) headquarters is risk-averse; or ii) as noted above, unrelated diversification can be better for relaxing credit constraints, and thereby increasing the resources under the control of headquarters.

²⁴ Contributions to this literature include: Rumelt (1974), Mueller (1987), Porter (1987), Ravenscraft and Scherer (1987), Wernerfelt and Montgomery (1988), Morck, Shleifer, and Vishny (1990), Kaplan and Weisbach (1992), Lang and Stulz (1994), Berger and Ofek (1995), and Comment and Jarrell (1995).

IV. Discussion

A. *When Will Internal Capital Markets Be Most Useful?*

The model has straightforward implications about the circumstances under which internal capital markets will tend to be the most effective sort of financial arrangement. When external markets are relatively “undeveloped”—in the sense of information and agency problems being particularly pronounced—and therefore credit constraints are very binding, it becomes more important to ensure that the limited funds can be efficiently reallocated across projects.

For concreteness, think of an Eastern European country that is just making the transition to a market-based economy. If the accounting and auditing technology is still primitive, or the legal protection afforded outside investors is weak, this might make it easy for managers to siphon off private benefits—i.e., the parameter s will be generally high for most projects. This in turn will make it hard for individual projects to attract the first-best level of financing.²⁵ In such circumstances, internal capital markets could potentially play a very valuable role. Of course, a corollary to this line of reasoning is that as external capital markets develop over time—i.e., the auditing and/or legal technology improves and the parameter s is effectively reduced—the comparative advantage of internal markets will be eroded, and one should expect (or hope) to see some breaking up of large multi-division firms.²⁶

B. *The Distinction between Headquarters and a Bank Lender*

What is it about the model of this article that makes it specifically about the internal capital allocation process? Couldn't one interpret what I label “corporate headquarters” as being equivalent to an outside bank lender? After all, both are intermediaries that perform a monitoring function, and that ostensibly help to make investment more efficient than it would be if projects were financed in the arm's length external market.

To answer this question, it is useful to compare this model to Diamond's (1984) model of a bank. In Diamond's model, the sole purpose of a bank is to bring more (or at least cheaper) funding to bear than the individual projects could raise on their own. There is no winner-picking function whatsoever—projects are *ex ante* identical and all get the same level of funding. In contrast, in my model, a primary role of headquarters is to reallocate a fixed pool of funds among winners and losers. Headquarters may also be able to ease credit

²⁵ Shleifer and Vishny (1995) argue that the ability of managers to flat-out loot companies in many of these transition economies is a first-order consideration in the design of financing and governance arrangements.

²⁶ Hubbard and Palia (1995) use exactly this reasoning to argue that the conglomerate mergers of the 1960s in the United States made sense at the time they were done, even though later it became optimal to undo them. The environment changed, in that external markets became more efficient, thereby reducing the comparative advantage of internal capital markets in allocating resources.

constraints, as was seen above, but unlike a bank, this is not necessary for it to create value.²⁷

Of course, a bank cannot engage in the sort of winner-picking that headquarters does in my model, because it does not have the authority to. As noted earlier, with fixed resources, the flip side of winner-picking is loser-sticking—making losers invest less than they could as stand-alones. This in turn requires being able to enforce an exclusive financing arrangement with the losers. For if this exclusivity feature were not present, any project manager observing that he is in state B (and not yet knowing the signals of his colleagues in other divisions) would, in an effort to protect his private benefits, preemptively secede from the joint-financing scheme.²⁸

While I have given headquarters an edge over a Diamond-style bank by endowing it with stronger control rights, I have also taken something away, by recognizing explicitly the difficulty of overseeing large numbers of projects. Thus the starkest distinction between the two models emerges when the monitoring technology $M(n)$ is such that headquarters optimally oversees a small number of projects, and there is no easing of credit constraints whatsoever—i.e., headquarters performs a pure winner-picking function. In this case, the two models have the opposite implications for the type of portfolio to be held by the intermediary. In contrast to the small (and possibly focused) set of projects overseen by headquarters, Diamond's model suggests that a bank ideally ought to lend to many unrelated borrowers, because such large-scale diversification is so central to its ability to create value through credit-constraint relaxation.²⁹

²⁷ Boyd and Prescott (1986) present a model of financial intermediary coalitions in which—unlike in Diamond (1984)—projects *are* evaluated ex ante, with better evaluations leading to more funding. In their model, this reallocation is made possible by the assumption that all project managers in the coalition can enter into binding contracts before the results of the evaluations are known, and that the contracts can be conditioned on these results, since the results are completely verifiable. Because they allow for complete contracting in this way, the Boyd-Prescott results do not really speak to the question of project ownership rights, and thus it is unclear whether they are better interpreted as being about a bank, an internal capital market, or some other institutional arrangement.

²⁸ Another important distinction between headquarters and a bank lender is that headquarters' control rights give it a greater incentive to acquire information about the projects it oversees. (See GSS (1994) and Aghion and Tirole (1994)). That is, if monitoring is costly, headquarters will optimally choose to devote more resources to monitoring. This feature is suppressed in the current model because I have assumed for simplicity that monitoring is costless.

²⁹ One can also contrast headquarters with other intermediaries such as venture capital partnerships. Again, the most important difference comes down to the strong control rights that enable headquarters to enforce loser-sticking and hence winner-picking. While venture capitalists usually hold some fraction of the voting rights in the companies they manage, they do not typically have the absolute authority to prevent these companies from seeking funding elsewhere (although they often do retain the right of first refusal on future rounds of financing—see Sahlman (1990)).

V. Conclusions

One theme of this article is that self-interested, empire-building type behavior on the part of corporate managers is not always a completely bad thing. Depending on the span of control that managers are given, these self-interested tendencies can sometimes be harnessed, and made to more or less line up with some aspects of shareholder value maximization.

Headquarters' role in an internal capital market is an example of this point. On the one hand, headquarters' general disposition to overinvest leads to credit constraints when dealing with the external market. Thus seen as an intermediary, headquarters may not be very successful in terms of making *more* funding available to individual projects. On the other hand, for a *given* amount of available funding, headquarters may—in the pursuit of its own self-interests—do a reasonable job of reallocating resources across projects. Thus a partial solution to the agency problem of empire-building project managers is to take control from them and give it to an empire-building supervisor, who is no more noble but who has less parochial interests.

It seems that this basic principle is a general one, and might be usefully applied to a number of other questions in organizational design. For example, one issue that has not been raised in this article is the optimal structure of a *multilevel* capital allocation hierarchy. To illustrate the sorts of questions that might be addressed, suppose a firm has four projects that it must allocate funds to. One way to go about this would be to use a simple two-layer structure of the sort studied above—i.e., there would be a single supervisor (“headquarters”) who would allocate all of the firm’s resources across the four project managers. Alternatively, one could add another level to the hierarchy, so that headquarters first splits the firm’s resources between two “middle managers.” Each of these middle managers in turn is given two of the four projects to oversee, and can divide his own allotment among these projects as he sees fit.

On the one hand, there are clearly benefits to the former, two-layer design. This design gives the party making the allocation decision the most direct access to the private benefits generated by all the projects together, and therefore offers the greatest potential for internalizing good investment decisions. On the other hand, one might generate an interesting tradeoff by explicitly taking into consideration the idea raised in Section III.A—namely, that any single agent has limited information acquisition capabilities. That is, it seems plausible that with the latter design, the two middle managers can learn more in total about the four projects than can a single supervisor. Thus there is a tension between quality of information (which argues for more layers in the capital allocation process) and a desire to avoid parochialism (which argues for a single omnipotent supervisor who can reap private benefits directly from *all* projects). Exploring the nature of this sort of tradeoff might provide an interesting avenue for further work.

Appendix

The Value of an Internal Capital Market as a Function of the Degree of Focus

	Outcome/Signal Configuration	Probability	Payoff:	Payoff:
			External Market	Internal Market
1.	GGHH	$p^2(q^2 + \alpha(1 - q)^2)$	$2\theta y_1$	$2\theta y_1$
2.	GBHH	$p(1 - p)q(1 - q)$	$y_1(1 + \theta)$	$y_1(1 + \theta)$
3.	BGHH	$(1 - p)p(1 - q)q$	$y_1(1 + \theta)$	$y_1(1 + \theta)$
4.	BBHH	$(1 - p)^2(1 - q)^2(1 + \alpha)$	$2y_1$	$2y_1$
5.	GGHL	$p^2(q(1 - q) - \alpha(1 - q)^2)$	$2\theta y_1$	θy_2
6.	GBHL	$p(1 - p)q^2$	$y_1(1 + \theta)$	θy_2
7.	BGHL	$(1 - p)p(1 - q)^2$	$y_1(1 + \theta)$	y_2
8.	BBHL	$(1 - p)^2(q(1 - q) - \alpha(1 - q)^2)$	$2y_1$	y_2
9.	GGLH	$p^2(q(1 - q) - \alpha(1 - q)^2)$	$2\theta y_1$	θy_2
10.	GBLH	$p(1 - p)(1 - q)^2$	$y_1(1 + \theta)$	y_2
11.	BGLH	$(1 - p)pq^2$	$y_1(1 + \theta)$	θy_2
12.	BBLH	$(1 - p)^2(q(1 - q) - \alpha(1 - q)^2)$	$2y_1$	y_2
13.	GGLL	$p^2(1 - q)^2(1 + \alpha)$	$2\theta y_1$	$2\theta y_1$
14.	GBLL	$p(1 - p)(1 - q)q$	$y_1(1 + \theta)$	$y_1(1 + \theta)$
15.	BGLL	$(1 - p)pq(1 - q)$	$y_1(1 + \theta)$	$y_1(1 + \theta)$
16.	BBLL	$(1 - p)^2(q^2 + \alpha(1 - q)^2)$	$2y_1$	$2y_1$

REFERENCES

- Aghion, Philippe, and Jean Tirole, 1994, Formal and real authority in organizations, Working paper, MIT and Nuffield College, Oxford.
- Berger, Philip, and Eli Ofek, 1995, Diversification's effect on firm value, *Journal of Financial Economics* 37, 39–65.
- Boyd, John H., and Edward C. Prescott, 1986, Financial intermediary-coalitions, *Journal of Economic Theory* 38, 211–232.
- Comment, Robert, and Gregg A. Jarrell, 1995, Corporate focus and stock returns, *Journal of Financial Economics* 37, 1995, 67–87.
- Diamond, Douglas, 1984, Financial intermediation and delegated monitoring, *Review of Economic Studies* 51, 393–414.
- Donaldson, Gordon, 1984, *Managing Corporate Wealth* (Praeger Publishers, New York).
- Gertner, Robert H., 1986, The organization of capital market transactions: Exclusive contracts and vertical integration under asymmetric information, Working paper, University of Chicago.
- Gertner, Robert H., David S. Scharfstein, and Jeremy C. Stein, 1994, Internal versus external capital markets, *Quarterly Journal of Economics* 109, 1211–1230.
- Green, Jerry, and Nancy Stokey, 1983, A comparison of tournaments and contests, *Journal of Political Economy* 91, 349–364.
- Grossman, Sanford, and Oliver 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, 1994, Capital budgeting, incentives, and information, Working paper, University of Chicago and Northwestern University.
- Hart, Oliver, and John Moore, 1990, Property rights and the nature of the firm, *Journal of Political Economy* 98, 1119–1158.
- Hubbard, R. Glenn, and Darius Palia, 1995, A re-examination of the conglomerate merger wave in the 1960s: An internal capital markets view, Working paper, Columbia University.
- Jensen, Michael, 1993, The modern industrial revolution, exit, and the failure of internal control systems, *Journal of Finance* 48, 831–880.
- Kaplan, Steven N., and Michael S. Weisbach, 1992, The success of acquisitions: Evidence from divestitures, *Journal of Finance* 47, 107–138.

- Lamont, Owen, 1997, Cash flow and investment: Evidence from internal capital markets, *Journal of Finance* 52, 83–109.
- Lang, Larry H. P., Eli Ofek, and René Stulz, 1996, Leverage, investment, and firm growth, *Journal of Financial Economics* 40, 3–29.
- Lang, Larry H. P., and René Stulz, 1994, Tobin's q , corporate diversification, and firm performance, *Journal of Political Economy* 102, 1248–1280.
- Lazear, Edward, and Sherwin Rosen, 1981, Rank-order tournaments as optimal labor contracts, *Journal of Political Economy* 89, 841–864.
- Morck, Randall, Andrei Shleifer, and Robert W. Vishny, 1990, Do managerial objectives drive bad acquisitions?, *Journal of Finance* 45, 31–48.
- Mueller, Dennis C., 1987, *The Corporation: Growth, Diversification, and Mergers* (Harwood, Chur, Switzerland).
- Meyer, Margaret, Paul Milgrom, and John Roberts, 1992, Organizational prospects, influence costs, and ownership changes, *Journal of Economics and Management Strategy* 1, 9–35.
- Myers, Stewart C., 1993, Still searching for optimal capital structure, *Journal of Applied Corporate Finance* 6, 4–14.
- Nalebuff, Barry, and Joseph Stiglitz, 1983, Prices and incentives: Towards a general theory of compensation and competition, *Bell Journal of Economics* 14, 21–43.
- Porter, Michael, 1987, From competitive advantage to corporate strategy, *Harvard Business Review* 65, 43–59.
- Ravenscraft, David J., and Frederic M. Scherer, 1987, *Mergers, Sell-offs, and Economic Efficiency* (Brookings Institute, Washington, D.C.).
- Rotemberg, Julio, and Garth Saloner, 1994, Benefits of narrow business strategies, *American Economic Review* 84, 1330–1349.
- Rumelt, Richard P., 1974, *Strategy, Structure, and Economic Performance* (Harvard University Press, Cambridge, Mass.).
- Sahlman, William A., 1990, The structure and governance of venture-capital organizations, *Journal of Financial Economics* 27, 473–521.
- Scharfstein, David S., and Jeremy C. Stein, 1990, Herd behavior and investment, *American Economic Review* 80, 465–79.
- Shin, Hyun-Han, and René M. Stulz, 1996, An analysis of the divisional investment policies of diversified firms, NBER Working paper No. 5639.
- Shleifer, Andrei, 1985, A theory of yardstick competition, *Rand Journal of Economics* 16, 319–327.
- Shleifer, Andrei, and Robert W. Vishny, 1997, A survey of corporate governance, *Journal of Finance* 52, 35–55.
- Stein, Jeremy C., 1995, Internal capital markets and the competition for corporate resources, NBER Working paper No. 5101.
- Wernerfelt, Birger, and Cynthia A. Montgomery, 1988, Tobin's q and the importance of focus in firm performance, *American Economic Review* 78, 246–254.
- Williamson, Oliver E., 1975, *Markets and Hierarchies: Analysis and Antitrust Implications* (Collier Macmillan Publishers, Inc., New York).