

NBER WORKING PAPER SERIES

INTERNAL CAPITAL MARKETS
AND THE COMPETITION FOR
CORPORATE RESOURCES

Jeremy C. Stein

Working Paper No. 5101

NATIONAL BUREAU OF ECONOMIC RESEARCH
1050 Massachusetts Avenue
Cambridge, MA 02138
May 1995

This research is supported by MIT's International Financial 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, Luigi Zingales and seminar participants at Harvard and the NBER. This paper is part of NBER's research program in Corporate Finance. Any opinions expressed are those of the author and not those of the National Bureau of Economic Research.

© 1995 by Jeremy C. Stein. All rights reserved. Short sections of text, not to exceed two paragraphs, may be quoted without explicit permission provided that full credit, including © notice, is given to the source.

INTERNAL CAPITAL MARKETS
AND THE COMPETITION FOR
CORPORATE RESOURCES

ABSTRACT

This paper examines the role of corporate headquarters in allocating scarce resources to competing projects in an internal capital market. Unlike a bank lender, headquarters has control rights that give it both the authority and the incentive 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 its own relationship with the outside capital market is fraught with agency problems and it therefore cannot help at all to relax overall firm-wide credit constraints. One implication of the model developed here is that internal capital markets may function more efficiently when companies choose relatively focused strategies.

Jeremy C. Stein
E52-448, Sloan School of Management
Massachusetts Institute of Technology
50 Memorial Drive
Cambridge, MA 02142-1347
and NBER

I. Introduction

This paper analyzes the process by which internal capital markets channel resources to different uses inside a company. I develop a model which has two key features: First, there is a company-wide scarcity of funds--i.e., corporate headquarters faces an aggregate credit constraint and may not be able to finance all positive-NPV projects. Second, headquarters is actively engaged in reallocating these scarce funds across projects. For example, the cashflow generated by one division's activities may be taken and spent on investment in another division. Or alternatively, one division's assets may be used as collateral to help raise financing that is then diverted to the other division. Thus some projects receive less total financing than they would in a stand-alone setting, while others receive more. Simply put, individual projects must compete for the scarce funds, and headquarters' job is to pick the winners and losers in this competition.

The model captures the spirit 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.) A similar theme is sounded by Donaldson (1984), who writes: "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 towards B.

This investment-interdependence view certainly seems to jibe with how managers think about some of the costs and benefits of integration. When Kodak announced in May of 1994 that it was putting up for sale its Sterling Winthrop drug unit and two other businesses (representing a total of approximately 23% of revenues) chairman George Fisher explained the move by stating: "Kodak's future calls for nothing less than total commitment to imaging...we were passing up too many opportunities (to make) small investments because we were so strapped for cash. If we attempted to retain both health and imaging, we would have short-changed both." (Wall Street Journal, May 4, 1994)

Recent empirical work by Lamont (1994) provides further support for the notion of investment interdependence. Lamont documents what might be termed "intra-company liquidity spillovers" in the U.S. oil industry: when major oil companies' cashflows were hard-hit by the oil price decline of 1986, they cut investment across the board, including investment in their non-oil-related divisions. Most strikingly, many of these cuts occurred in the petrochemicals divisions of these oil companies. This is particularly noteworthy because the appeal of investing in petrochemicals should rise when oil prices fall, as oil is a major input in the production

process. And indeed, stand-alone petrochemical firms (i.e., those not affiliated with oil companies) did significantly increase investment over the same period.

All of this makes it sound as if internal capital markets allocate financing to individual projects or divisions in a way that is fundamentally different from, say, the way that bank lenders allocate credit to their customers. For example, one can easily imagine a single bank lending to both pure-play oil producers and petrochemical producers. And yet it seems much less likely that such a bank would sharply cut lending to the petrochemical producers when its oil clients had a bad year.¹ Our basic intuition about banks suggests that even if they underlend on average as compared to a first-best perfect-capital-markets benchmark, credit is allocated more according to the absolute merit of projects, and less according to the relative merit. Or said differently, the business of banking would seem to be more about increasing the total funding available to worthy projects, and less about dividing up a fixed pool of resources.

Yet at the same time, it is not immediately obvious on theoretical grounds why a bank and an internal capital market should be so different. Clearly, both represent a form of financial intermediation. Arguably, both arise out of the same primitive agency or information problems, which make it difficult for certain types of investment projects to finance themselves directly in the arm's length capital market. And in both cases, there is some presumption that these intermediaries are an efficient vehicle for acquiring information about the projects, thereby

¹Of course, one can think of examples where there is some degree of spillover across different bank clients. If a poorly capitalized bank suffers large losses on its real estate portfolio, it may become constrained by capital regulations and therefore lend less to its non-real estate related clients. (See, e.g., Bernanke and Lown 1991.) However, these spillovers would appear to be much less common and less pronounced in magnitude than those occurring in internal capital markets. A primary goal of this paper is to explore the underlying reasons for this apparent difference.

alleviating agency/information-type frictions and facilitating efficient investment.

So why does financing via an internal capital market work so differently than bank lending? In this paper, I follow recent work by Gertner, Scharfstein and Stein (1994), henceforth GSS, and focus on the different control rights associated with the two forms of finance. In particular, in an internal capital market, corporate headquarters has residual rights of control--in the sense of Grossman and Hart (1986) and Hart and Moore (1990)--over the projects it finances, while a bank lender does not. I argue here that this has two important consequences: corporate headquarters has both more incentive and more authority than a bank to engage in "winner-picking", the practice of shifting funds from one project to another.

While this paper shares with GSS the emphasis on headquarters' control rights in an internal capital market, it differs in one very fundamental way. In the GSS model, headquarters was assumed to have unlimited financial resources available. This implies that the issues that are of central concern here--the allocation of a fixed pool of funds, and the accompanying phenomenon of winner-picking--could never arise in the GSS model.² In contrast, in this paper, headquarters is assumed to have no pre-existing wealth of its own. Moreover, headquarters faces credit constraints that arise endogenously out of an agency problem between itself and the outside capital market. Taken together, these two features deliver the fixed-resources attribute that is a necessary starting point for thinking about winner-picking.

In some cases, the credit constraint facing headquarters will take an extreme form:

²Harris and Raviv (1994) also model the internal capital allocation process in a setting where, as in GSS, headquarters faces no financial constraints. Unlike in either GSS or this paper however, Harris and Raviv do not stress control rights in the Grossman-Hart(1986)/Hart-Moore (1990) sense.

headquarters will not be able 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 in an internal capital market, 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: I argue below that since a bank is less able to engage in winner-picking, it can never create value unless it eases credit constraints on average for the projects it finances.

In practical terms, the sharpest distinction between a bank and an internal capital market probably has to do with the number and type of projects that each finances. Standard models of financial intermediation such as Diamond (1984) suggest that banks should lend to large numbers of borrowers, and that ideally these borrowers ought to be in a variety of unrelated businesses. In contrast, in my model, an internal capital market can create value while financing only a small number of projects. Moreover, in some cases, the winner-picking function of an internal capital market actually works better when the company chooses a focused strategy of financing closely related, rather than unrelated investment projects.

The argument in favor of a focused strategy flows directly from the observation that 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 relativistic ranking, it does not matter if headquarters makes absolute errors in evaluating the merits of projects, as long as these errors are correlated across the projects.³

³There is a close analogy between this logic and that underlying the literature on tournaments, or yardstick competition --see, e.g., Lazear and Rosen (1981), Green and Stokey (1983), Nalebuff and Stiglitz (1983) and Shleifer (1985). In both cases, the crucial ingredient in the story is the existence of errors that are correlated across projects/managers.

As a concrete example, consider a situation where a company adopts a focused strategy, and finds itself dividing up resources between two divisions that both produce toys. In some cases, headquarters will misjudge the absolute merit of investing in either toy division--e.g., they might be overly optimistic about the state of toy demand during the Christmas selling season. But if the total resources that are to be committed are fixed, these sorts of mistakes do not matter. All that matters is that headquarters correctly assess which division produces the better toys. In contrast, if the company were to choose a strategy of unrelated diversification, it might find itself dividing up resources between, say, one toy division and one biotechnology division. In this case, it would indeed misallocate resources if it misjudged the demand for toys.

The remainder of the paper is organized as follows. Section II lays out a simple 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 III shows how these investment distortions can be partially alleviated when multiple projects are located inside a larger corporate entity, with corporate headquarters acting as an intermediary between the projects and the external market. This section illustrates the basic winner-picking role of an internal capital market, and also discusses at some length the contrast with bank lenders. In Section IV the model is generalized to address the question of whether focus or unrelated diversification is the better strategy for maximizing the value created by an internal capital market. Section V concludes.

II. 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.⁴ 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 other wealth, nor does he have any managerial ability. This implies that the project will also require the services of a both professional manager--who also has no wealth of his own--as well as funding from a group of financiers. Throughout the paper, 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 1 or 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 cashflow y_1 if an amount 1 was invested, and returns y_2 if an amount 2 was invested. The cashflows 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 cashflows. Furthermore, it is assumed that $1 < y_1 < y_2 < 2$. In other words, from the perspective of the

⁴By "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.

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.

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 will 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 cashflow 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 show 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 1) always provide 1 unit of financing, or 2) always provide 2 units of financing.

In the former scenario, the expected net cashflows (i.e., net of the initial investment) on the project are $(p\theta + (1-p))y_1 - 1$. By virtue of our earlier assumptions, 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 cashflow 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 whatever is left over.

In the latter scenario, where investment is 2 in all states, the expected net cashflows are $(p\theta + (1-p))y_2 - 2$. Given our earlier assumptions, this quantity will be negative if the probability p of state G occurring is sufficiently small. From this point on, I will assume that this is the case. This implies that it is infeasible for managers to raise 2 units of financing--they

cannot credibly promise to fully repay the financiers, no matter how they structure the financing arrangement, because they simply do not generate enough cashflows on average.

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.

C. The limits of revelation schemes

The inefficiency described above results because managers' private information is not exploited, and investment is hence not conditioned on the state. This raises the question of whether one might be able to use an incentive scheme to elicit manager's 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.

To give revelation schemes their best shot, I assume that managers are asked to make an announcement of their private information before investment takes place, and that both this announcement and the ultimate realization of the state can be costlessly verified. This means that it is contractually feasible to give managers an incentive payment of the form $w(a,r)$, where a is their announcement, and r is the realized state. The only constraint is that $w(a,r) \geq 0$, since managers have no wealth of their own to begin with.

There are two incentive constraints that must be satisfied. First, a manager must be willing to announce that the state is B when this is the case:

$$w(B,B) + sy_1 \geq w(G,B) + sy_2 \quad (1)$$

Second, the manager must be willing to announce that the state is G when this is the case:

$$w(G,G) + s\theta y_2 \geq w(B,G) + s\theta y_1 \quad (2)$$

It is easy to see that of these two constraints, only (1) will be binding. The lowest-cost way of satisfying the constraints is to set $w(G,G) = w(G,B) = w(B,G) = 0$, and to set $w(B,B) = s(y_2 - y_1)$. This has an intuitive interpretation: the manager has to be given a bonus for voluntarily choosing to invest only 1 instead of 2 in state B, and this bonus has to exactly offset the added private benefits that he passes up by doing so.

Thus in an ex ante sense, the 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. This 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. In order to simplify the exposition, I will for the remainder of the paper focus my attention on such cases, where revelation schemes are not used. This returns us to the informationless credit-rationing solution discussed above, where the level of investment is always 1, regardless of the state.

III. 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 paper

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 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.

Following GSS, one of the aspects of an internal capital market that I emphasize here has to do with control rights. Specifically, in an internal capital market, 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.⁵

In GSS, we also assumed that corporate headquarters was itself completely free of credit constraints--i.e., it always had unlimited funds to invest in the projects in question. Here, I go to the other polar extreme. I assume that headquarters itself has no resources of its own. Thus like project management, it has to raise any funds it wishes to invest from the outside financial market. Moreover, headquarters, like project management, is also subject to agency problems--

⁵In 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) cashflows. 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.

its interests diverge from those of the outside market. In particular, headquarters will also have a preference for maximizing the resources under its control.

At first blush, it might appear that these sorts of assumptions would rule out a useful role for internal capital markets. Financing a project via an internal capital market would seem to involve the creation of a second layer of agency (i.e., headquarters) without any obvious compensating benefits. And indeed, as we will see momentarily, this logic is correct in the single-project case. However, when headquarters is given control over multiple projects simultaneously, an internal capital market can in some situations add value.

The specifics are as follows. First, headquarters has the ability to monitor the projects it is responsible for. Monitoring involves a small cost of c per project, and it gives headquarters a (possibly noisy) signal of the state of the world. Second, its control rights allow headquarters to effectively appropriate from project management a fraction $\phi < 1$ of the private benefit. Thus, for example, if publicly verifiable cashflows are y , and total private benefits are sy , headquarters now gets to keep ϕsy , leaving only $(1-\phi)sy$ for project management.

This latter assumption is essentially identical to one made in GSS. There, we argued that by reducing project managers' ex post rents, headquarters would also reduce their ex ante incentives to put forth effort. This might in turn reduce investment output. To capture this here in a simple reduced-form way, I assume that when a project is overseen by headquarters, all cashflows 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 cashflows 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

cashflows 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 final assumption to be made is relevant only in the case where headquarters controls 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 non-owner 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.⁶ This assumption is very much in the spirit of Hart and Moore

⁶See Gertner (1986) for a different perspective on the significance of such exclusivity in an internal capital market.

(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)

In sum, there are four important assumptions being made about corporate headquarters: First, it has no resources of its own. Second, it can acquire information about projects' ex ante prospects. Third, its control rights allow it to capture a share of the private benefits associated with any project it oversees; however, this comes at the cost of diluting project managers' incentives. And fourth, its control rights also allow it to redistribute resources across projects.

B. The sub-optimality of an internal capital market in a one-project setting

Given the assumptions, it is clear that an internal capital market is always inferior to the external market in a one-project setting. First of all, the dilution of project managers' incentives in an internal market 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 offset. However, headquarters will be no more inclined to communicate truthfully with the outside market than will project-level managers.

To see this, assume that headquarters has sunk the monitoring cost c , and has observed 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.⁷

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.

It should be noted, though, that the conclusions here are very different from those in GSS. In that paper, an internal capital market could in many circumstances add value even when it financed only a single project. Again, this is in large part because in GSS, we assumed away credit constraints at the level of corporate headquarters.⁸

C. The role of headquarters with multiple projects and fixed resources: picking winners (and sticking losers)

In order for an internal capital market 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 case. Suppose that there are just two projects, i and j, each of which is exactly as described in Section 2 above. The only difference between the two projects is that

⁷This 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 cashflows. 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.

⁸Similarly, 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.

their productivity shocks are not perfectly correlated. Indeed, the two 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 two 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 case. Given the equilibrium behavior that emerges in the model, outside investors always can be relied on to provide at least 2 units of funding. Moreover, there will in general be some circumstances in which headquarters will be able to raise more than the sum of what the individual projects could raise on their own. However, if one can document a net benefit to internal capital allocation without a change in overall resources available, this will provide a simple and intuitive lower bound for how much an internal capital market can improve things in the more general case.

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 must assume that this is in fact 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, once it has sunk the cost

of monitoring, 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 cashflows. Thus in attempting to maximize its private benefits, headquarters will also pick the mix of investment that maximizes verifiable cashflow value.

It then remains to verify that headquarters will indeed decide to sink the monitoring cost of c per project. However, if this cost is small relative to the increase in private benefits that accrues to headquarters from making the right decisions, it will be strictly optimal for headquarters to monitor. Henceforth, I assume that c satisfies this sort of condition, so that headquarters always engages in monitoring.

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 have to face credit constraints from the outside market--it still has some incentive to channel the funds it does obtain towards "winners" and away from "losers". (One 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, while 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:

$$EM = 2(y_1(p\theta + (1-p)) - 1) \quad (3)$$

Equation (3) reflects the fact that in an external market, each of the two projects always invests exactly 1. In contrast, in an internal market, the total expected net output, denoted by IM, is given by:

$$IM = 2(1-p)^2ky_1 + 2p^2k\theta y_1 + 2p(1-p)ky_2 - 2 \quad (4)$$

Equation (4) differs from (3) in two ways: First, there is the effort-dilution factor k , which tends to make IM less than EM. Second, when the two 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.⁹

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 towards projects that are winners and away from those that are losers.

⁹Note that the monitoring cost c does not appear in (4). This is because it does not affect the net verifiable cashflows available to the founder. When c rises incrementally, this simply eats into the rents enjoyed by headquarters.

Of course, I have rigged things somewhat by making headquarters' private benefits directly proportional to project cashflows. 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 seeks to maximize "corporate wealth".¹⁰ And it is exactly the sort of premise that has led Jensen (1993) and others to argue that plentiful internal cashflow is more likely to create a divergence between managerial and shareholder interests when companies have relatively few attractive investment opportunities.¹¹

Clearly, the assumption that private benefits and verifiable cashflows 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 cashflows. 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

¹⁰Donaldson'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).

¹¹Louis Kaplow has suggested an alternative formulation that leads to similar results. Suppose that private benefits are proportional to the amount invested rather than the cashflow produced, so that based on private benefits alone, managers care purely about the scale of investment, and are completely indifferent as 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 towards overinvestment.

funding status even when its cashflows do not merit this, simply because it offers disproportionately high private benefits to headquarters.¹²

D. Discussion

1. 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 credit constraints are therefore very binding, it will be more important 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 quite primitive, this might make it easy for managers to siphon off private benefits--i.e., the parameter s would be generally high for most projects. This in turn would make it hard for individual projects to attract the first-best level of financing. In such circumstances, internal capital markets could potentially play a very valuable role.

To take one example, headquarters in an Eastern European internal market might find

¹²Such 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 IV below.

itself overseeing the funding of a large, old "smokestack" plant as well as several less tangible, but more innovative and higher-NPV development projects. Now it may be that the old smokestack plant has sufficient current cashflow and collateral value that on a stand-alone basis, it could continue to reinvest substantially. But at the same time, given the overall scarcity of resources, it might be better to divert this funding to the more innovative development projects, which would have difficulty raising enough themselves if left as stand-alones. This is a case where headquarters' ability to engage in winner-picking might be particularly useful.

Of course, a corollary to this line of reasoning is that as external capital markets develop over time--i.e., the auditing 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.

2. Intra-company liquidity spillovers

As noted in the Introduction, one motivation for the theory developed here is the empirical work of Lamont (1994), who documents intra-company liquidity spillovers in the oil industry. However, the model sketched above does not literally deliver predictions about such spillovers. This is because in the context of the one-period model, all resources available to headquarters come from new external finance, and there is no scope for internal liquidity as a source of funds.

Fortunately, it requires only a very minor reinterpretation to think of the model as being about intra-company liquidity spillovers. In a multi-period setting, headquarters will have at its disposal both: 1) any new external funds it can raise against its portfolio of projects; as well as

2) any internal (verifiable) cashflow that these projects have already generated and that has not yet been paid out to the outside investors under the terms of the financial contract.¹³ Assuming that headquarters has the same reallocation authority with respect to internally generated cashflows that it does with externally raised funds, the liquidity-spillovers implication will follow.

According to the simplest version of the model, any liquidity spillovers of the sort documented by Lamont (1994) would be interpreted as efficiency-enhancing. However, if one generalizes the model to allow for private benefits that are not always strictly proportional to cashflows, the interpretation would be less restrictive: ex ante the potential for intra-company liquidity transfers would be seen as valuable on average, but in any specific case, such a transfer might also simply reflect headquarters' desire to subsidize a pet project.

3. The relaxation of external credit constraints and the optimal scale of an internal capital market

In the analysis of the above section, it was assumed that headquarters has no more total resources to work with than the individual projects would if they were financed as stand-alones in the external market. And importantly, headquarters can still add value, by doing a better job of allocating the fixed amount of resources. A natural next question to ask is if headquarters might, by virtue of its greater allocative efficiency, be able to raise more total resources from

¹³It is easy to construct examples where outside investors optimally allow verifiable cashflows to be retained and reinvested. If outside investors can predict that these funds will be put to use in a way that is positive-NPV from their perspective, allowing the funds to be retained is isomorphic to advancing new external financing at the initial date in the model.

the outside market than would the individual project managers.

There are two related reasons why the answer to this question might be yes. First, even if it oversees just two projects, there are some parameter values for which headquarters can raise three units of funding even when the individual projects can just raise one each. Intuitively, this is because of the independence of project outcomes. Given that it has the right incentives to allocate funds across projects, headquarters will only overinvest with three units when both projects are in state B simultaneously. This occurs with probability $(1-p)^2$. In contrast, in the stand-alone setting, if there is more than one unit of financing per project, overinvestment occurs with probability $(1-p)$.

Second, and more generally, credit constraints could in principle be relaxed through a policy of large-scale diversification. To see how, consider an extreme example where headquarters oversees a very large number of statistically independent projects, say 1000. 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. This in turn implies that outside investors can count on being repaid even if they invest $1000(1+p)$, rather than just 1000. Investment will therefore converge to the first-best level, and credit constraints will no longer be binding.¹⁴

In this example, large-scale diversification performs a function similar to that in Diamond's (1984) model of a bank--it allows outside investors to have perfect certainty about

¹⁴Although 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 .

the aggregate distribution of investment output, even if they know nothing about the state of the individual projects. Given this perfect certainty, they are willing to invest more.

This line of argument suggests that the model as developed thus far--with just 2 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. There are a number of possibilities in this regard.

Perhaps the most simple and direct way to close the model would be to acknowledge that headquarters has limited capacity for overseeing multiple projects. In a model such as this, one has to take fairly seriously the notion of headquarters as a single agent, or something quite close to it. It would be inappropriate to just assume--in the absence of an underlying theoretical justification--that headquarters can supervise an unlimited number of projects, without either compromising its monitoring ability or altering its incentives. Thus as a very crude first cut, one might simply posit that the per-project monitoring cost c is an increasing convex function of the number of projects overseen by headquarters. In this setting, the 2-project case that was analyzed above would correspond to a monitoring technology that becomes very costly once pushed beyond a couple of projects.¹⁵

¹⁵A 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. An alternative way to endogenize the optimal scale of an internal capital market would be to appeal to the sorts of arguments made by Meyer, Milgrom and Roberts (1992) and Rotemberg and Saloner (1994). Both of these papers point out reasons why it may be problematic for an owner such as headquarters to engage in large-scale diversification; they are discussed in more detail in Section IV.D below.

4. The distinction between headquarters and a bank lender

The various analogies to Diamond's (1984) model underscore the question that was raised in the Introduction, and that has been lurking beneath the surface throughout: What is it about the model of this paper that makes it specifically about the internal capital allocation process? Or said differently, 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.

In broad terms, the distinction between headquarters and a bank lender can be stated as follows: The primary role of a bank is to bring more (or at least cheaper) funding to bear than the individual projects could raise on their own. Indeed, in Diamond's setting, this is all a bank does. 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, the main function of headquarters is to reallocate a fixed pool of funds among winners and losers. Headquarters may also be able to ease credit constraints, as was just illustrated, but unlike with a bank, this is not at all necessary for it to create value.¹⁶

¹⁶Boyd 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.

Of course, a bank simply cannot engage in the sort of winner-picking that headquarters does in my model, because, as noted above, it does not have the authority to. With fixed resources, the flip side of winner-picking is what may be termed "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, pre-emptively secede from the joint-financing scheme.

Moreover, given the authority to do so, the whole reason that headquarters actually chooses to monitor and then makes the right allocation decisions is that its control rights also enable it to capture some of the private benefits associated with doing so. Even if a bank lender did somehow have the authority to engage in resource reallocation, it would have weaker incentives to pick the right projects, because it could not capture any private benefits. Thus, for example, even if the monitoring costs c were made trivially small, a bank endowed with the same monitoring technology as headquarters would endogenously choose not to monitor.¹⁷

Perhaps the starkest distinction between the two theories has to do with their respective implications for the optimal type of "portfolio" to be held by the intermediary. First, with

¹⁷The notion that headquarters' control rights give it a greater incentive to acquire information than a bank lender also shows up in GSS and in Aghion and Tirole (1994). In the context of the current model, it is important to remember that we are always talking about choosing between two investment strategies--i.e., invest either 2 in a state-G project or 1 each in a state-G and a state-B project--which are both positive-NPV in terms of verifiable cashflows. Thus a bank lender would not lose money no matter what choice it makes. This is why, absent private benefits, a bank would have weaker incentives to pick the better mix of investment, and would be unwilling to expend any resources on project evaluation.

regard to the size of the portfolio, a Diamond-style model has the clear implication that a bank ought to lend to a large number of borrowers, because diversification is so central to its ability to create value. In contrast, the value of corporate headquarters does not hinge solely on the sort of large-scale diversification effects that make outside investors more willing to provide financing. Thus one might expect that internal capital markets could play a useful role even when they finance only a few separate divisions, and hence when they do not do much to relax overall credit constraints.¹⁸

Second, with regard to the composition of the portfolio, there is again a very clear prediction for banks: as much as is feasible, they should seek to lend to unrelated borrowers. Again, the logic is one of diversification--this is the best way to reduce the ex post variance of outcomes.

In an internal capital market, the implications for portfolio composition are somewhat subtler. In general, opposing forces are at work. On the one hand, as with banks, there is an "ease the credit constraint" effect of the sort described above. This effect tends to favor a strategy of headquarters investing in unrelated lines of business. On the other hand, if one holds the degree of credit constraint fixed, the desire to allocate a given amount of resources more efficiently can in some cases actually lead companies to choose focused strategies, in which they invest in related lines of business.

In order to understand why focused strategies might make more sense for corporate headquarters, it is necessary to generalize the model somewhat. This task is taken up next.

¹⁸In other words, internal capital markets can be valuable even when one explicitly recognizes that limited capacity may prevent headquarters from overseeing large numbers of projects simultaneously.

IV. Focus vs. Unrelated Diversification as Corporate Strategies

To keep the formal analysis simple, I will focus on the case considered in Section III.C above, where there are 2 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 associated with diversification. The only remaining question is whether by choosing a focused strategy, headquarters can do a better job in the pure winner-picking dimension.

A. A trivial argument in favor of diversification: uncorrelated project outcomes

Even keeping the total resources available to headquarters fixed, there is another reason why a diversification strategy might make sense. Suppose the model is generalized slightly to allow project outcomes to be somewhat correlated, while maintaining the assumption that headquarters can predict these outcomes without error. Then if one thinks of diversification as doing nothing more than making project outcomes less correlated, it follows immediately that the value of an internal capital market is an increasing function of the degree of diversification.

The intuition is most easily seen by considering 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. Thus focus reduces the value of an internal capital market.

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 where there remains a good deal of independence across projects. For example, consider a pharmaceutical company that invests in 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. 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. To the extent that headquarters' job is to assess the relative, rather than absolute merits of projects, no harm will come from errors in judgement, so long as these errors are correlated across the projects. And as I argue in more detail below, 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 $1/2 < q < 1$, 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: $\text{prob}(L^i / G^i, G^j, L^j) = (1-q)(1+\alpha) > \text{prob}(L^i / G^i) = (1-q)$. Analogously, $\text{prob}(H^i / B^i, B^j, H^j) = (1-q)(1+\alpha) > \text{prob}(H^i / B^i) = (1-q)$. Here α is a parameter between 0 and 1 that I would like to interpret as a measure of the degree of focus--the more focused is the strategy, the greater is α .¹⁹

The rationale behind such an interpretation is as follows. When headquarters monitors two closely related projects, whatever biases or misjudgments it makes about one project are likely to also show up in its assessment of the other project. For example, in the drug-

¹⁹As mentioned in footnote 2 above, this sort of information structure--with correlated prediction errors--is also a necessary ingredient in tournament-style models. It is also a key feature of the herding model of Scharfstein and Stein (1990).

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 much less likely that the same misjudgment will be made across the two 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.

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 $\text{prob}(L^i / B^i, G^j, L^j) = \text{prob}(L^i / B^i) = q$, and analogously that $\text{prob}(H^i / G^i, B^j, H^j) = \text{prob}(H^i / G^i) = q$. 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. Table 1 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 Table 1.

Given the information in the table, 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 an inspection of the table. 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 absolutely no harm. Both projects each receive one unit of funding, which is the correct allocation of resources. In contrast, getting differing 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.²⁰

C. Pulling it together: the optimal degree of focus

Taken together, the arguments in the previous two subsections suggest that there are two opposing forces at work. On the one hand, to the extent that i) focus primarily 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 to the extent that i) there are errors in assessing projects; and ii) focus primarily increases the correlation of these errors, then a focused strategy will tend to make an internal capital market work better.

D. Other work on focus vs. unrelated diversification

There is a large literature on the relative merits of focus vs. unrelated diversification as corporate strategies, and I will not attempt to do it justice here. Rather, I will touch briefly on just a few recent contributions that seem particularly relevant. On the theoretical front, the ideas in this paper are related to those in Rotemberg and Saloner (1994), and Meyer, Milgrom and Roberts (1992). Both of these other papers are also set in an incomplete contracting framework, and both also build models that can be used as arguments in favor of focused strategies. However, what is meant by "focus" seems to differ somewhat across the various works.

²⁰Note 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.

For example, in Rotemberg and Saloner (1994), the main point is really that it can be better for a firm to be in fewer lines of business. By committing to what they call "narrowness" such a firm can better motivate its employees to pursue innovative activities in an incomplete contracting environment. In contrast, the notion of focus developed above has less to do with how many lines of business a firm is in per se. Rather, it emphasizes headquarters' ability to accurately assess the relative prospects of any given number of businesses.

Like Rotemberg and Saloner (1994), Meyer et al (1992) is also really more an argument for a company being involved in fewer lines of business. In particular, it can be desirable for a company to divest a poorly performing unit, when that unit is draining cash from other, more profitable operations. What is most interesting about their paper in relationship to this one is that they emphasize what might be thought of as the dark side of headquarters' ability to reallocate funds. They argue that the potential for such reallocations may lead project managers to engage in wasteful influence activities. Thus while their model is built on the same basic premise as the one in this paper--the idea that headquarters has the control rights that allow it to shift resources from one project to another--they cast this feature as a drawback of internal capital markets, while I have cast it as a potential advantage.

In principle, one could combine the insights of the two models to provide a more complete description of the optimal scope of an internal capital market. In Section III.D.2 above, I gave one simple reason why an internal capital market might optimally be restricted to only a small number of projects: headquarters cannot effectively monitor a large number. The Meyer et al (1992) logic suggests another potential constraint on firm size: the more projects headquarters oversees, the more resources it controls, and the greater the potential for wasteful

influence activities. Thus an optimal balance may involve having an internal capital market restrict its scope to a relatively small set of projects where the winner-picking benefits are so large (perhaps because there is a high degree of focus among the projects) as to exceed the influence costs.

On the empirical front, there is a large body of work, beginning with Rumelt (1974), that asks in one way or another whether diversified firms perform better or worse than their more focused counterparts. Some of this work looks at accounting-based measures of performance over time. The results here seem to be somewhat sensitive to the sample period and the control techniques used. (See Mueller 1987 for a survey.) Several more recent papers have employed stock market measures to capture investors' perceptions of value. Wernerfelt and Montgomery (1988), Lang and Stulz (1994) and Berger and Ofek (1995) all document a negative relationship between diversification and Tobin's q .²¹

In a related vein, Morck, Shleifer and Vishny (1990) show that unrelated acquisitions have been treated by the stock market as bad news, while Comment and Jarrell (1995) document that increases in focus have been seen as good news. Finally, Porter (1987), Ravenscraft and Scherer (1987) and Kaplan and Weisbach (1992) all find that firms are much more likely to divest unrelated, as opposed to related units that were obtained through prior acquisitions.

Overall, the message that emerges from this work is that the value consequences of unrelated diversification are at best highly suspect. At the same time, more focused strategies involving multiple units in the same or similar lines of business do not appear to do as badly.

²¹ For example, Berger and Ofek (1995) estimate that multi-segment firms trade at a discount of roughly 15% relative to single-segment firms. Moreover, this discount is more pronounced when the multiple segments are in different 2-digit SIC codes.

Although they hardly constitute a sharp test, these findings are broadly consistent with the spirit of the theoretical story sketched above.

V. Conclusions

One theme of this paper 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 paper is the optimal structure of a multi-level 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, the forces considered in this paper would seem to strictly favor 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 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.

References

- Aghion, Philippe and Jean Tirole, "Formal and Real Authority in Organizations," working paper, 1994.
- Berger, Philip and Eli Ofek, "Diversification's Effect on Firm Value, Journal of Financial Economics, 37, 1995, 39-65.
- Bernanke, Ben S. and Cara S. Lown, "The Credit Crunch," Brooking Papers on Economic Activity, 1991, 205-39.
- Boyd, John H. and Edward C. Prescott, "Financial Intermediary-Coalitions," Journal of Economic Theory, 38, 1986, 211-232.
- Comment, Robert and Jarrell, Gregg A., "Corporate Focus and Stock Returns," Journal of Financial Economics, 37, 1995, 67-87.
- Diamond, Douglas, "Financial Intermediation and Delegated Monitoring," Review of Economic Studies, 51, 1984, 393-414.
- Donaldson, Gordon, Managing Corporate Wealth, Praeger Publishers, New York, 1984.
- Gertner, Robert H., "The Organization of Capital Market Transactions: Exclusive Contracts and Vertical Integration under Asymmetric Information," MIT working paper, 1986.
- Gertner, Robert H., David S. Scharfstein and Jeremy C. Stein, "Internal versus External Capital Markets," Quarterly Journal of Economics, 109, 1994, 1211-1230.
- Green, Jerry and Nancy Stokey, "A Comparison of Tournaments and Contests," Journal of Political Economy, 91, 1983, 349-64.
- Grossman, Sanford and Oliver Hart, "The Costs and Benefits of Ownership: A Theory of Vertical and Lateral Integration," Journal of Political Economy, 94, 1986, 691-719.
- Harris, Milton and Artur Raviv, "Capital Budgeting, Incentives and Information," Northwestern University working paper, 1994.
- Hart, Oliver and John Moore, "Property Rights and the Nature of the Firm," Journal of Political Economy, 98, 1990, 1119-58.
- Jensen, Michael, "The Modern Industrial Revolution, Exit, and the Failure of Internal Control Systems," Journal of Finance, 48, 1993, 831-80.
- Kaplan, Steven N. and Michael S. Weisbach, "The Success of Acquisitions: Evidence from Divestitures," Journal of Finance, 47, 1992, 107-138.

Lamont, Owen, "Cash Flow and Investment: Evidence from Internal Capital Markets," MIT working paper, 1993.

Lang, Larry H.P. and René Stulz, "Tobin's q , Corporate Diversification, and Firm Performance," Journal of Political Economy, 102, 1994, 1248-80.

Lazear, Edward and Sterwin Rosen, "Rank-Order Tournaments as Optimal Labor Contracts," Journal of Political Economy, 89, 1981, 841-64.

Morck, Randall, Andrei Shleifer and Robert W. Vishny, "Do Managerial Objectives Drive Bad Acquisitions?" Journal of Finance, 45, 1990, 31-48.

Mueller, Dennis C., The Corporation: Growth, Diversification, and Mergers, Harwood, Chur, Switzerland, 1987.

Meyer, Margaret, Paul Milgrom and John Roberts, "Organizational Prospects, Influence Costs, and Ownership Changes," Journal of Economics and Management Strategy, 1, 1992, 9-35.

Myers, Stewart C., "Still Searching for Optimal Capital Structure," Journal of Applied Corporate Finance, 6, 1993, 4-14.

Nalebuff, Barry and Joseph Stiglitz, "Prices and Incentives: Towards a General Theory of Compensation and Competition," Bell Journal of Economics, 14, 1983, 21-43.

Porter, Michael E., "From Competitive Advantage to Corporate Strategy," Harvard Business Review, 65, 1987, 43-59.

Ravenscraft, David J. and Frederic M. Scherer, Mergers, Sell-offs, and Economic Efficiency, Washington, DC, Brookings Institute, 1987.

Rotemberg, Julio and Garth Saloner, "Benefits of Narrow Business Strategies," American Economic Review, 84, 1994, 1330-49.

Rumelt, Richard P., Strategy, Structure and Economic Performance, Harvard University Press, Cambridge, MA, 1974.

Scharfstein, David S. and Jeremy C. Stein, "Herd Behavior and Investment," American Economic Review, 80, 1990, 465-79.

Shleifer, Andrei, "A Theory of Yardstick Competition," Rand Journal of Economics, 16, 1985, 319-27.

Wernerfelt, Birger and Cynthia A. Montgomery, "Tobin's q and the Importance of Focus in Firm Performance," American Economic Review, 78, 1988, 246-254.

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

Table 1
The Value of an Internal Capital Market
as a Function of the Degree of Focus

	<u>Outcome/Signal Configuration</u>	<u>Probability</u>	<u>Payoff: External Market</u>	<u>Payoff: Internal Market</u>
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	$(1-p)^2(q^2 + \alpha(1-q)^2)$	$2y_1$	$2y_1$