THE JOURNAL OF FINANCE • VOL. XLIII, NO. 2 • JUNE 1988

Bond Covenants and Delegated Monitoring

MITCHELL BERLIN and JAN LOEYS*

ABSTRACT

This paper examines alternative contracting arrangements available to a firm seeking to finance an investment project. The authors consider the choice between loan contracts with covenants based on noisy indicators of the firm's financial health and loan contracts enforced by a monitoring specialist. In one interpretation, the specialist is a financial intermediary. The firm's choice is shown to depend upon the firm's credit rating, the accuracy of financial indicators of the firm's condition, the loss from premature liquidation of the firm's project, and the cost of monitoring.

IN RECENT YEARS, FINANCIAL theorists have made substantial progress in explaining the functions of financial intermediaries and complex financial contracts. One strand of the literature explains the existence of specialized monitoring institutions as the result of market failures in direct credit markets. For instance, financial intermediaries have been characterized as delegated monitors. (See, for example, Boyd and Prescott [3], Campbell and Kracaw [4], Chan [5], Diamond [6], and Gorton and Haubrich [7].) Another strand of the literature focuses on the role of contractual covenants as mechanisms to control agency problems between firm insiders and outside investors. (See, for instance, Ho and Singer [8], Kalay [9], and Smith and Warner [13].) Yet there has been little work on the marginal benefits and costs of using monitoring specialists when covenants are at least partially effective in controlling agency problems.

This paper presents a model of a firm's choice between two different contractual arrangements: loan contracts with covenants but no monitor and loan contracts enforced by a monitoring specialist. One possible interpretation, and the one that we shall emphasize for concreteness, is that the first type of contract is a bond and the second type is a bank loan. This links our work to the literature on financial intermediation. A second interpretation is that our monitor is a bond trustee or an auditing firm hired by a trustee to examine the firm's finances. In this interpretation, the model provides a formal analysis of the demand for bond clauses that permit the trustee to audit the firm on demand. The model portrays a credit market in which different types of firms make different contractual

* Federal Reserve Bank of Philadelphia and J. P. Morgan and Company, respectively. We have benefited from suggestions by John Boyd, Joe Haubrich, George Pennacchi, Tony Saunders, and an anonymous referee, who cannot, of course, be blamed for any mistakes. The views expressed do not necessarily represent the views of the Federal Reserve Bank of Philadelphia, the Federal Reserve System, or J. P. Morgan and Co. Research for this paper was completed while Jan Loeys was an economist at the Federal Reserve Bank of Philadelphia.

choices and in which changes in the information technology shift the margin of choice.

In our model, as in much of the literature on why banks exist, bondholders have inadequate incentives to monitor on their own. However, in place of the typical informational assumption that investors choose between becoming fully informed and remaining totally ignorant, investors are assumed to face a richer menu of choices about the quality of the information they can gather about firms. We maintain a distinction between routine accounting information that can often be gathered at low cost and detailed information about the firm's prospects, which can only be gathered by intensive and costly monitoring. In addition, we introduce observable differences between firms that can affect contractual choices. Firms have different histories, operate in different markets, and have projects likely to produce different streams of returns. When firms enter credit markets, market participants have substantial prior knowledge that permits them to distinguish different types of firms, and contractual choices will reflect this knowledge.

We permit individual investors to observe low-cost, but noisy, interim indicators of the firm's future ability to repay. These indicators might be routine financial ratios used to measure the firm's financial health or even the firm's ability to make coupon payments on time. Bond covenants are written as functions of these indicators, and the firm's inability to satisfy these covenants places the firm in default. However, since the covenants are based on imperfect information, default policies based on these covenants will be inefficient. Bond contracts will tend to be either too harsh or too lenient; harsh contracts lead to excessive default, while lenient contracts permit too many unpromising projects to mature. A detailed investigation of the firm's condition would provide more information and would thus avoid this inefficiency. However, such investigations are costly, and bondholders holding diversified portfolios have limited private incentives to monitor, even when monitoring is worthwhile for all investors taken together.

When bond covenants based on imperfect information would lead to too many mistakes, the services of a monitoring specialist—a "banker"—may be required. However, this specialist must be provided with incentives both to monitor and to choose an efficient default policy. Our model examines the tradeoff between the gains from an efficient default policy and the agency costs that arise when the task of monitoring is delegated. In particular, we examine these tradeoffs as a function of a few parameters that characterize the firm and the information technology. The most important of these are the firm's "credit rating", the predictive value of interim financial indicators of the firm's future health, the costs of a detailed investigation of the firm's financial condition, and the lost value from premature liquidation of projects.

The remainder of the paper is organized as follows. Section I presents the model. Section II describes the optimal contracts without monitoring. Section III shows that securityholders have inadequate incentives to monitor and presents the optimal contract with monitoring. Section IV describes the tradeoffs between the different types of contracts and presents comparative-statics results. The final section concludes and discusses possible extensions.

I. The Model[1]

Consider a market composed of N risk-neutral entrepreneurs who have a project but no wealth and M identical, risk-averse investors who have wealth but no projects. Throughout, we shall refer to the entrepreneurs as firms. All projects require the investment of a single unit of wealth in period 0 and yield final project returns in period 2, the final period. Assume that, together, investors have sufficient wealth to finance the projects of all firms, i.e.,

$$WM > N, \tag{1}$$

where W is the wealth of one investor. To keep the notation simple, the model is developed in terms of a single firm. Unless otherwise noted, different firms may have different parameter values.

After contracts have been drawn up and funds transferred in period 0, each firm's project type is determined by an independent, random variable, $t \in \{t_G, t_B\}$, where t_G indicates a good project and t_B a bad project. In period 1, the decision is made either to permit the project to mature or to liquidate the project. If a project is allowed to mature until the second period, its discounted gross revenues are either X_g, with probability $f(t)$, or X_b (zero for all firms), with probability $1 - f(t)$. Thus, even if the project type becomes known, some uncertainty about the final-period returns remains. Good projects are more likely to yield positive returns than are bad projects, i.e., $f(t_G) > f(t_B)$, and all final project returns can be observed by anyone.

Alternatively, if a project of *either* type is liquidated in period 1, it yields a nonrandom, discounted liquidation value, X_d, in period 2. Alternatively, X_d might be the firm's value after a reorganization. We assume that the expected gross revenues from a good project that matures exceed the project's liquidation value and that the opposite is true of bad projects; i.e.,

$$f(t_G)X_g > X_d > f(t_B)X_g. \tag{2}$$

In addition to investing in projects, investors may invest in a storage technology yielding a discounted return of 1 in the final period. The return to storage exceeds the liquidation value of projects; i.e., $X_d < 1$. Inequality (1) implies that, in equilibrium, some investors will store at least some of their wealth.

The information structure of the model is as follows. When funds are invested in period 0, no firm knows the type of project that it will draw. In period 1, however, each firm learns whether its project is good or bad, while investors observe costlessly an indicator of each project's type. The indicator is independent across projects and is denoted by $y \in \{y_G, y_B\}$. Everyone knows both the conditional distribution of project types for each value of the indicator and the prior distribution of the indicator. Let $p(t_i \mid y_j)$ denote the conditional probability of t_i given y_j, and let $p(y_j)$ denote the prior probability of y_j. Given the simple structure of uncertainty, a useful parameterization of the probability distributions

[1] For a list of symbols used in this model, see Appendix A.

is as follows:

$$\begin{bmatrix} p(t_G \mid y_G) & p(t_G \mid y_B) \\ p(t_B \mid y_G) & p(t_B \mid y_B) \end{bmatrix} = \begin{bmatrix} \rho + \alpha(1-\rho) & \alpha(1-\rho) \\ 1-\rho-\alpha(1-\rho) & 1-\alpha(1-\rho) \end{bmatrix}$$

and (3)

$$\begin{bmatrix} p(y_G) \\ p(y_B) \end{bmatrix} = \begin{bmatrix} \alpha \\ 1-\alpha \end{bmatrix},$$

where $0 < \alpha < 1$ and $0 < \rho < 1$. In this parameterization, α is both the prior probability of a favorable indicator and the prior probability of a good project. ρ measures the informativeness of the indicator, with $\rho = 1$ corresponding to a perfectly informative indicator and $\rho = 0$ corresponding to a totally uninformative indicator.[2] The indicator may be interpreted as a balance-sheet ratio, perhaps the working capital-asset ratio, or the firm's ability to make a coupon payment. If this indicator is a very good predictor of the firm's ultimate ability to repay, then ρ is near one.

In addition to observing the costless but noisy indicator, investors can learn any firm's project type by monitoring the firm in period 1 at a cost, e, the disutility of effort required to monitor the firm. We assume that no one is privy to the outcome of another's investigation or, for that matter, to whether any investigation was performed at all.

There are two types of lending arrangements, *bonds* and *bank loans*. Bonds are claims on the firm's final project returns sold directly by firms to investors. Alternatively, a group of investors may hire another investor—a "banker"—to contract with a single borrowing firm on their behalf. (Henceforth, the term "lender" will be used when statements might apply equally to bondholders or bankers.)

Both bonds and bank loans contain two types of clauses, *covenants* and *payments*. A covenant confers upon lenders the right to place the firm in default for specified values of the indicator, thereby forcing the firm to liquidate the project.[3] It is the contractual right to intervene under well-specified circumstances, rather than fixed contractual payments in non-default states, that distinguishes our bank loans and security contracts as debt contracts.[4] While nothing in the model requires the contracts to have fixed contractual payments in non-default states, all of our contracts are feasible if this constraint is imposed. The covenant may be conditioned both upon the indicator and upon statements

[2] An indicator with higher ρ is more informative in Blackwell's [2] sense. In addition, if $t_G = 2$, $t_B = 1$, and $y_G = 2$, $y_B = 1$, then ρ is the correlation coefficient beween t and y.

[3] We view the breach of a covenant restriction as an indicator of the firm's financial condition, which is exogenous in our model. Some covenants—such as restrictions on dividend payouts and priority rules—are primarily means of controlling moral-hazard problems and are best viewed as mechanisms to constrain management and equityholders' actions. This aspect of covenants is emphasized by Smith and Warner [13].

[4] Equityholders, of course, have open-ended rights of intervention as legal owners of the firm. However, outside equityholders, with limited holdings in a particular firm, will have limited incentives to bear the substantial costs of monitoring firm insiders. In fact, this is one reason why bond contracts that place clear, legally binding constraints on the firm may be preferred to equity.

made by lenders about the project's type.[5] A covenant can be written as a function:

$$d(t_i', y_j) \in \{0, 1\} \quad \text{for} \quad i, j = G, B. \tag{4}$$

The first argument of the function, t_i', is a statement made by a lender that the project is of type t_i. When $d(\cdot) = 1$, the project is permitted to mature, and, when $d(\cdot) = 0$, the firm is placed in default and the project is liquidated.

All payments are made by the firm in the final period when project returns are realized. Thus, payments may depend upon realized revenue as well as indicators and statements. For notational convenience only, payments made to the banker and to the bank's investors are written as if they were made directly by the borrowing firm. As long as final revenues are observable, there is no loss of generality. An arrangement in which the banker distributed revenues to investors (or vice versa) would yield an identical distribution of revenues among the different claimants. Thus, the total payment made by the firm to investors (other than the banker) can be written:

$$r_k(t_i', y_j) \quad \text{for} \quad i, j = G, B \quad \text{and} \quad k = g, b, d. \tag{5}$$

This is the payment made to investors when final revenues are X_k, the indicator is y_j, and a lender makes a statement t_i' about the project. Similarly, payments to the banker can be written:

$$b_k(t_i', y_j) \quad \text{for} \quad i, j = G, B \quad \text{and} \quad k = g, b, d. \tag{6}$$

Since firms have no wealth, all payments must be made out of project revenues. Thus, a feasible contract requires that

$$X_k \geq r_k(t_i', y_j) + b_k(t_i', y_j) \quad \text{for any} \quad i, j, k, \tag{7}$$

and, when the firm is placed in default (i.e., $k = d$), this holds with equality. In addition, all payments must be non-negative:

$$r_k(t_i', y_j) \geq 0 \quad \text{and} \quad b_k(t_i', y_j) \geq 0 \quad \text{for any} \quad i, j, k. \tag{8}$$

Together, (7) and (8) imply that no agent will have negative final-period consumption.

Consider now the different agents' objective functions. Assume that any statements about the type of project are truthful. Later, incentive-compatibility constraints will be imposed to make this true. A firm's expected profits at period 1, after it learns that it has project type t_i and after y_j has been observed, are

$$\pi_1^f(t_i, y_j) = d(t_i, y_j)\{f(t_i)[X_g - r_g(t_i, y_j) - b_g(t_i, y_j)]\}$$
$$+ (1 - d(t_i, y_j))\{X_d - r_d(t_i, y_j) - b_d(t_i, y_j)\}. \tag{9}$$

At period 0, the firm's expected profits are

$$\pi_0^f = \textstyle\sum_i \sum_j \pi_1^f(t_i, y_j) p(t_i \mid y_j) p(y_j). \tag{10}$$

[5] In our model, truthful revelations by the borrowing firm can only be achieved by a contract that treats firms with the same indicator identically, i.e., a pooling contract. Thus, statements by firms can provide no information that is not contained in the indicator.

Now consider an investor who has placed w dollars in a firm or a bank (that has invested in a firm). w is also the investor's share of the firm's outstanding securities because each project requires a single unit of wealth. Since investors are risk averse, we assume that each invests in many firms or banks and holds a well-diversified portfolio. Thus, investors will evaluate returns as if they were risk neutral. The investor's expected payments from a firm of unknown type when y_j has been observed are

$$\pi_1^I(y_j) = w\{\textstyle\sum_i[d(t_i, y_j)f(t_i)r_g(t_i, y_j) + (1 - d(t_i, y_j))r_d(t_i, y_j)]p(t_i \mid y_j)\}, \quad (11)$$

while expected payments at period 0, before the indicator has been observed, are

$$\pi_0^I = \textstyle\sum_j \pi_1^I(y_j)p(y_j). \quad (12)$$

II. Optimal Contracts without Monitoring

When no monitoring occurs, lenders know only the public indicator. Formally, a contract without monitoring can be viewed as a pooling contract, with terms depending solely upon the value of the indicator. In fact, when no monitoring occurs, only bonds will be used since there is no function for a banker (whose sole purpose in our model is to act as a specialized monitor). The no-monitoring assumption implies that contractual terms will depend only upon the indicator, or that

$$d(t_G, y_j) = d(t_B, y_j) \quad \text{for all} \quad j \quad (13)$$

and

$$r_k(t_G, y_j) = r_k(t_B, y_j) \quad \text{for all} \quad j, k. \quad (14)$$

In a competitive market in which all investors have access to a storage technology, the bond contract must maximize the firm's expected profits subject to investors' willingness to supply funds. In Appendix B, we describe market equilibrium when there are many ex ante distinguishable firms and show that, for a firm that can pay a higher return than storage, the optimal bond contract must satisfy the following problem:

$$\max_{d(\cdot), r(\cdot)} \quad \pi_0^f \quad (15)$$

subject to

$$\pi_0^I \geq w, \quad (16)$$

and (7), (8), (13), and (14).

Since there are two values of the indicator, there are four possible default policies, only two of which can be optimal. One contract, the *lenient bond contract* (S_L) never places the firm in default. Another contract, the *harsh bond contract* (S_H), places the firm is default only when the bad indicator, y_B, is observed. Contracts that place the firm in default when a good indicator is observed cannot be optimal for the following reasons. One that imposes default only when the good indicator is observed is never optimal because a good indicator is good news about the project; if it is optimal to impose default when $y = y_G$, then the same must be true when $y = y_B$. (See Milgrom [10] for a definition of "good news".) However, default will not always be imposed because, by assumption, storage strictly dominates any contract that always places the firm in default.

For both the lenient and harsh bond contracts, the firm's maximum expected profits can be derived by solving (16) as an equality and substituting into (15). Letting $\pi^f(S)$ denote the firm's maximum expected profits at period 0 under contract type S, we have

$$\pi^f(S_L) = p(y_G)[f(t_G)X_g p(t_G \mid y_G) + f(t_B)X_g p(t_B \mid y_G)] + p(y_B)[f(t_G)X_g p(t_G \mid y_B) + f(t_B)X_g p(t_B \mid y_B)] - 1 \tag{17}$$

and

$$\pi^f(S_H) = p(y_G)[f(t_G)X_g p(t_G \mid y_G) + f(t_B)X_g p(t_B \mid y_G)] + p(y_B)X_d - 1. \tag{18}$$

To understand these contracts, it is helpful to compare them with the first best, the optimal liquidation policy if the project type were known. The first-best policy maximizes total contractual surplus, which by (2) involves liquidating bad projects and continuing good projects. The S_L contract, however, does not liquidate bad projects, while the S_H contract liquidates good projects when the indicator is misleadingly unfavorable. Subtracting (18) from (17) and using (3) results in an expression for the relative profitability of the two contracts. For expositional purposes, we make the simplifying assumption that $f(t_B)$ is very close (but not equal) to one and that $f(t_G)$ is very close (but not equal) to zero for all firms. While nothing important hinges upon this assumption—essentially the same results have been derived for the more general case—it permits an uncluttered presentation of results.[6] Thus,

$$\begin{aligned}\pi^f(S_L) - \pi^f(S_H) &= p(y_B)[p(t_G \mid y_B)X_g - X_d] \\ &= (1 - \alpha)[\alpha(1 - \rho)X_g - X_d].\end{aligned} \tag{19}$$

This difference is positive and the S_L contract is optimal when α is high and ρ is low. When α, the prior probability of a good project, is high, it is optimal to ignore even unfavorable information. When ρ is low, the indicator is an imprecise predictor of the firm's future prospects, and bad information is discounted. The difference is also positive when X_g is large and X_d is small, that is, when the liquidation value of the project is small relative to the value of the project should it succeed. This will be the case, for example, if the project involves capital goods that have no ready secondary market. When this is true, the cost of liquidating too seldom is small and the S_L contract is optimal.

III. Optimal Contracts with Monitoring

In this section, after showing that diversified bondholders will, in general, monitor too little, we derive the optimal incentive contract for a delegated monitor and

[6] If $f(t_B) = 0$, then the firm with a bad project could always be induced to state truthfully its project type and separation could be achieved without monitoring. The simplification creates a bias toward harsh contracts; in the general case, good projects are not quite so good and bad projects are not quite so bad. In the simple case, the "degree of goodness" is captured by varying X_g and X_d parametrically.

analyze the agency costs of contractual schemes with monitoring. At the end of the section, alternative interpretations of these contractual arrangements are discussed.

Since there are two values of the indicator, monitoring might occur for either or both of the two indicators. However, monitoring when the indicator is good can be valuable only when both the prior probability of a good project and the informativeness of the indicator are quite low, precisely those values for which storage is likely to dominate investment in the firm under any type of contract. Thus, we focus on contracts in which monitoring takes place only when the indicator is unfavorable. Such contracts permit a more efficient liquidation policy while economizing on both monitoring and agency costs.

A necessary condition for monitoring to be valuable is that the expected gross revenues from an efficient liquidation policy less monitoring costs exceed the expected gross revenues from both (i) liquidating all projects for y_B and (ii) continuing all projects; i.e.,

$$\text{i. } p(t_G \mid y_B)X_g + p(t_B \mid y_B)X_d - e \geq X_d,$$

$$\text{ii. } p(t_G \mid y_B)X_g + p(t_B \mid y_B)X_d - e \geq p(t_G \mid y_B)X_g, \tag{20}$$

which can be rewritten:

$$\text{i. } p(t_G \mid y_B)[X_g - X_d] \geq e,$$

$$\text{ii. } [1 - p(t_G \mid y_B)]X_d \geq e. \tag{21}$$

Securityholders, however, will have inadequate incentives to monitor when each has limited wealth invested in any one firm. Even if we abstract from the real-world problems of achieving bondholder consensus, any individual bondholder will weigh the private gains from monitoring—measured by the left-hand sides of (21.i) and (21.ii) multiplied by w—against the full costs of monitoring. When w is small, no investor will monitor even when (21) holds. We assume, henceforth, that bondholders hold a well-diversified portfolio and will not choose to monitor.[7]

We now examine the case where a delegated monitor—a banker—is provided with a compensation scheme that induces an efficient liquidation policy for y_B. The banker, like other investors, is too risk averse to invest an appreciable amount of wealth in the project he or she monitors without bearing excessive risk. Thus, the banker invests in the market portfolio and must be provided with a payment scheme to induce him or her to monitor efficiently on behalf of other investors. We make the standard assumptions that the banker's objection func-

[7] If there are substantial transactions costs of assembling and maintaining a diversified portfolio, then investors might be unable to diversify fully. Our model effectively assumes that these transactions costs are negligible. The reduction in brokerage fees and the rapidly expanding menu of mutual funds available to the typical investor have dramatically reduced transactions costs in recent years. We are hesitant, however, to view a mutual fund manager as a delegated monitor in our sense. The incentives for mutual fund managers to engage in detailed monitoring of the firms in the fund's portfolio are quite limited. A more typical behavior is to sell securities at the first sign of bad news.

tion,

$$U(b_k + W) - e,$$

is separable in income and monitoring effort and that $U' > 0$ and $U'' < 0$.

Incentive compatibility requires that the banker have an incentive both to monitor and to reveal truthfully the outcome of his or her investigation. It is easy to show that, if the banker has an incentive to monitor, then he or she will truthfully reveal the outcome of his or her investigation. Thus, the binding incentive-compatibility conditions are that the banker's expected utility from an efficient liquidation policy is greater than the expected utility from (i) a policy of liquidating all projects and (ii) a policy of permitting all projects to continue. These two conditions can be written:

$$\text{i. } p(t_G \mid y_B)U[b_g(t_G, y_B) + W] + p(t_B \mid y_B)U[b_d(t_B, y_B) + W] - e \geq U[b_d(t_B, y_B) + W],$$

$$\text{ii. } p(t_G \mid y_B)U[b_g(t_G, y_b) + W] + p(t_B \mid y_B)U[b_d(t_B, y_B) + W] - e \geq p(t_G \mid y_B)U[b_g(t_G, y_B) + W] + p(t_B \mid y_B)U[W].^8 \quad (22)$$

Letting the function G denote the inverse of U, treating both as equalities, and solving for b_g and b_d, we have

$$\text{i. } b_g(t_G, y_B) = G\left(\frac{e}{p(t_G \mid y_B)(1 - p(t_G \mid y_B))} + U[W]\right) - W,$$

$$\text{ii. } b_d(t_B, y_B) = G\left(\frac{e}{1 - p(t_G \mid y_B)} + U[W]\right) - W, \quad (23)$$

while the banker's compensation for y_G is zero since no monitoring occurs when the indicator is favorable. Note that $b_g > b_d$. If this were not true, the banker would always choose the certain return from a policy of liquidating all projects over the uncertain return from monitoring.

Let B denote the contract in which the banker monitors and chooses an efficient liquidation policy for y_B. The optimal bank contract chooses a set of payments to maximize

$$\begin{aligned}\pi_0^f = p(y_G)&[p(t_G \mid y_G)\{X_g - r_g(t_G, y_G) - b_g(t_G, y_G)\}]\\ &+ p(y_B)[p(t_G \mid y_B)\{X_g - r_g(t_G, y_B) - b_g(t_G, y_B)\}\\ &+ p(t_B \mid y_B)\{X_d - r_d(t_B, y_B) - b_d(t_B, y_B)\}] \quad (24)\end{aligned}$$

subject to

$$\begin{aligned}\pi_0^I = p(y_G)&p(t_G \mid y_G)r_g(t_G, y_G)\\ &+ p(y_B)[p(t_G \mid y_B)r_g(t_G, y_B) + p(t_B \mid y_B)r_d(t_B, y_B)] \geq 1 \quad (25)\end{aligned}$$

[8] Unlike Diamond [6], we do not permit the monitor to be subjected to "nonpecuniary" penalties. Although the time costs of bankruptcy for firm management and the loss of reputation are, no doubt, important, it is very doubtful that these costs can be designed as part of an optimal contract as in Diamond. Rather, they should be modeled explicitly as part of the legal or market environment.

and (23). Define $z \equiv p(t_G | y_B) = \alpha(1 - \rho)$. This variable is inversely related to the informativeness of bad information about the project. When z is low, a bad indicator is reliable evidence that the project is bad. Letting (25) be satisfied with equality and substituting this and (23) into (24), the firm's maximum expected profits are

$$\pi^f(B) = \alpha(\rho + z)X_g + (1 - \alpha)[zX_g + (1 - z)X_d - C(e, z)] - 1, \quad (26)$$

where

$$C(e, z) \equiv zG(ez^{-1}(1 - z)^{-1} + U[W]) + (1 - z)G(e(1 - z)^{-1} + U[W]) - W.$$

C denotes the total expected payments required to make the banker choose efficient liquidation decisions. If the outcome of the banker's investigation were observable, then it would suffice to pay him or her the utility cost of monitoring, e, whenever he or she monitored. However, since monitoring is unobservable, the banker must be compensated in excess of the actual costs of monitoring to induce him or her both to monitor and to make efficient liquidation decisions. This excess is the agency cost of hiring a delegated monitor.

The behavior of C is as follows. First, C is increasing in e since $G' > 0$. Second, for any strictly increasing utility function, C grows infinitely large as z approaches zero or one. The intuition is straightforward. When z is very low, the probability of finding a good project is very low. Unless the banker's payment for a successful project (b_g) is very large, he or she will choose to liquidate all projects. When z is very high, the banker will find a policy of allowing all projects to mature to be very attractive. Thus, the banker's payment when the project is liquidated (b_d) must be very large to induce him or her to monitor. Since $b_g > b_d$, so must the banker's payment be should the project succeed.

Third, C must be increasing in z for $z \geq 1/2$. As z increases, the banker is more likely to find a good project when the indicator is bad. Thus, the banker's liquidation payment must increase to induce him or her to monitor rather than to choose a policy of allowing all projects to mature. Also, for $z \geq 1/2$, b_g is increasing in z. Since both payments are increasing in z and since the probability of the banker receiving the larger payment is also increasing, total expected payments must be increasing. Finally, for typical utility functions such as the square root and logarithmic functions, C can be shown to be decreasing in z when z is small. For small values of z, b_g is decreasing in z, thus imparting a negative slope to C, which achieves a minimum at some value of z less than or equal to $1/2$.

What types of real-world contractual arrangements does our incentive-compatible scheme resemble? We consider two interpretations. One interpretation is that our monitor is a loan officer delegated to oversee part of a bank loan portfolio. This interpretation rationalizes the monitor's risk aversion and inability to invest an appreciable amount of wealth in the project.[9] The loan contract

[9] If the banker can invest an appreciable amount of his or her own wealth without bearing excessive risk, then the costs of inducing an efficient default policy are lowered because the banker's interests are, to some extent, harmonized with bank investors. If individual bankers can manage a very large, diversified portfolio, then even small amounts of wealth will suffice. (In the limit, a perfectly diversified portfolio could be managed efficiently with zero agency costs and no wealth posted.) If, however,

is of the following type: If the firm's financial condition as measured by the indicator is good, then the banker does not monitor and the firm pays a contractually agreed loan rate. (As noted above, this economizes on monitoring costs.) If the indicator is bad, then the banker retains the right to reorganize the firm, conditional upon the outcome of an investigation. Whenever the indicator is bad, the firm must pay a higher markup on its loan, reflecting the costs of monitoring: both direct costs and agency costs.

In a second interpretation, the monitor is not a banker but a bond trustee or an auditor hired by the trustee. In addition to distributing payments, trustees are often given the contractual right to audit the firm on demand. The audit, which is usually performed by an outside firm, is triggered by the firm's inability to satisfy promptly covenant requirements. In this interpretation, our model provides a formal analysis of the marginal benefits and costs of contractual clauses that give trustees the right to audit the firm.

While suggestive, both of these interpretations have problems. For the first interpretation to be fully consistent, the rise in the firm's markup must be paid out to the loan officer as part of a managerial compensation scheme. While de facto compensation schemes may resemble our incentive schedule, there is little evidence that formal compensation schemes take this form. (See Robert Morris Associates [12].) Also, the rise in the markup when the firm fails to satisfy covenant requirements will be the outcome of a bargaining game, perhaps constrained by implicit agreements. We are currently studying when an incentive-compatible monitoring scheme can be sustained without a binding written contract. Similarly, trustees and auditing firms are typically paid a fixed fee rather than a variable payment that depends upon the outcome of their investigations. It is clear that more work is required to bring the theoretical contracting scheme closer to those actually observed in financial markets.

IV. The Choice between Contracts with and without Monitoring

In this section, we show that the optimal contract choice depends upon the relative inefficiencies of excessively lenient or excessively harsh liquidation policies and upon the agency costs of hiring a banker to make efficient liquidation decisions. We also present comparative statics that show how contractual choice is affected by the model parameters.

Using expressions (17), (18), and (26), we can derive expressions for the differences in the firm's expected profits under the three contracts:

$$\text{i. } \pi^f(B) - \pi^f(S_L) = (1-\alpha)\{(1-z)X_d - C(e, z)\},$$

$$\text{ii. } \pi^f(B) - \pi^f(S_H) = (1-\alpha)\{z(X_g - X_d) - C(e, z)\},$$

$$\text{iii. } \pi^f(S_L) - \pi^f(S_H) = (1-\alpha)\{z(X_g - X_d) - (1-z)X_d\}. \tag{27}$$

monitors have limited spans of control, as in Boyd and Prescott [3] and Ramakrishnan and Thakor [11], then diversification will require a coalition of bankers. Such coalitions face difficult tradeoffs between risk sharing and free-rider problems. We feel that real-world limits on the ability to impose risk on bank managers are sufficiently important for the limiting case we consider to be of interest.

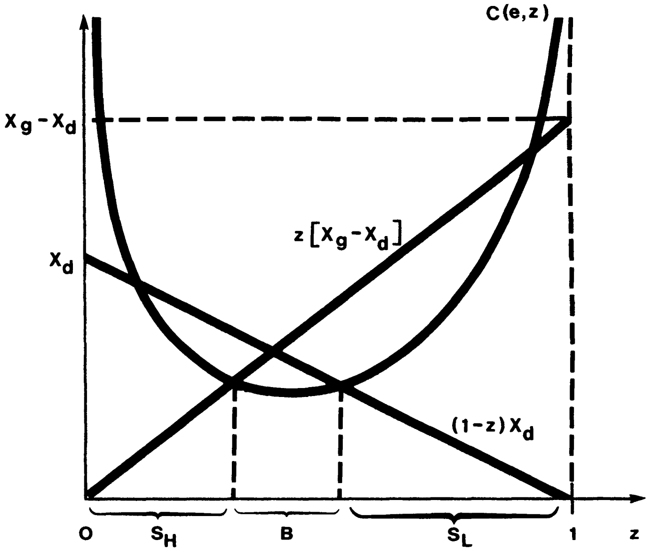

Figure 1. The Relationship between z and the Optimal Contract

A graphical analysis of these expressions is possible. There are three separate terms to consider: $(1-z)X_d$, $z(X_g - X_d)$, and C. Although it is possible for the B contract to be dominated throughout—if e is very large or $(X_g - X_d)$ is very small—we focus on those cases where banks dominate securities for some values of z; i.e., both (27.i) and (27.ii) are positive for some z.

Figure 1 shows a typical configuration for which this is true.[10] The S_H contract dominates for small z, and the S_L contract dominates for large z. For both small and large z, the agency costs of hiring a delegated monitor are very large and the marginal benefits from an efficient liquidation policy are very small; thus, one of the bond contracts will be optimal. The harsh contract will be optimal for z near zero because few good projects will be found when the indicator is bad. The lenient contract will be optimal for z near one because many of the projects with bad indicators will turn out to be good. For intermediate values of z, however, the marginal value of monitoring is large and, in turn, the expected costs of inducing efficient decision making by the banker are small. In this range, the bank contract will be chosen.

We now investigate how variations in the parameters affect the optimal contract choice. Figure 1 shows that, as long as investors have insufficient incentives to monitor, a decrease in the cost of monitoring increases the range over which bank loans dominate. This follows because $dC/de > 0$. However, as e continues to fall, it may become profitable for individual investors to monitor the firm themselves.[11] The same is true if the individual investor's share (w) of a firm's debt rises sufficiently.

[10] This configuration will hold if $(X_g - X_d)$ is not too small. For a more complete analysis, see Berlin and Loeys [1]. This paper also has a more detailed analysis of the comparative statics and examines the conditions under which monitoring for both values of the indicator will be optimal.

[11] Although for sufficiently small monitoring costs it may be profitable for individual investors to

Consider now the effect of changes in α or ρ. As α, the prior probability of a good project, increases, z increases. This implies that the S_L contract dominates over an increasing range. For firms that are very likely to have good projects, the incremental value of interim information is small unless the indicator is very informative. Thinking of α as a firm's credit rating, this says that a highly rated firm is more likely to secure finance using bonds with lenient covenants than a lower rated firm. Firms with lower credit ratings are likely to secure finance using bank loans or bonds with harsh covenants.

It is natural to interpret ρ as a measure of the stability of the relationship between the information upon which covenants are based and the firm's ultimate ability to repay. Thus, if the covenants specify that the firm must satisfy certain financial ratios, a large ρ indicates that these ratios bear a stable relationship to the firm's longer term project revenues. As ρ increases, either of the two contracts that are responsive to interim information, B and S_H, become relatively more attractive. For firms without very high credit ratings, it can be shown that, if $(X_g - X_d)$ is large, then the model predicts the use of bank contracts for intermediate values of ρ and the use of harsh bond contracts for high values of ρ.

Finally, Figure 1 shows that a rise in X_g or a decline in X_d enlarges the range over which lenient security contracts dominate and reduces the preference for either harsh security contracts or loan contracts. This result is no surprise. For a project with a low liquidation value, neither monitoring nor harsh covenant restrictions produce large gains because investors cannot gain much from early liquidation. Low liquidation values are likely to be found with firms that have invested mostly in firm-specific or intangible capital that does not have a secondary market.

V. Conclusion

We have examined two contractual alternatives available to a firm seeking to finance an investment. Debt contracts with covenants based on noisy indicators of the firm's financial condition tend to be either too harsh (too many good projects are liquidated) or too lenient (too many bad projects are permitted to mature). Hiring the services of a delegated monitor ensures a more efficient liquidation policy, but providing the monitor with proper incentives is costly. The firm's optimal choice reflects a tradeoff between the inefficiencies of rigid bond covenants and the agency costs of hiring a delegated monitor. This tradeoff depends upon a number of parameters characterizing the firm's production technology and the information technology.

Although our contractual choices have similarities to institutions in the real world, we are yet some distance from being able to reproduce the contractual forms actually observed in financial markets. Although we have interpreted the

monitor firms themselves, the actual decision to monitor has game-theoretic elements that we do not analyze. Also, in a world where individual securityholders do have incentives to monitor, a reasonable model must take account of conflicts of interest among claimants and the costs of forging consensus, which play no part in our model.

delegated monitor as a banker, to be fully convincing this interpretation requires an explicit analysis of the bank's liability side. Also, while the model treats all outcomes under the bank loan as if they arose from a binding written contract, we feel that renegotiation and implicit agreements are quite important features of delegated-monitoring arrangements. Our debt contracts without monitoring are more nearly recognizable as bonds, but the secondary market for financial assets plays no role in our model. It would be quite interesting to analyze the tradeoffs between complex covenants, which may be necessary to control agency problems, and the size of the secondary market for a firm's securities. In future work, we hope to address some of these problems.

Appendix A

List of Symbols

α: probability that a firm's project is of type t_G.
$b(\cdot)$: payment by firm to bank.
B: bank (loan) contract.
$d(\cdot)$: default covenant; $d(\cdot) = 0$: default; $d(\cdot) = 1$: no default.
e: monitoring costs.
$G(\cdot)$: inverse of $U(\cdot)$.
π_x^f: expectation at time x of firm's profits.
π_x^I: expectation at time x of investor's profits.
$\pi^f(S)$: firm's maximum expected profit if lending contract is of type S, where $S = B, S_H, S_L$.
$r(\cdot)$: payment by firm to investor (other than the bank).
ρ: informativeness of the indicator y.
S_L: lenient bond contract.
S_H: harsh bond contract.
t_i: type of project, where $i = G$ (good) or B (bad).
$U(\cdot)$: banker's utility of income.
X_k: period-2 revenues of a project, discounted to period 1, where $k = g, b, d$.
y_j: indicator about project's type, where $j = G, B$.
w: share of individual investor in total debt of a firm.
W: total wealth of an individual investor.
z: $p(t_G \mid y_B)$.

Appendix B

Firms are indexed by $n = 1, \cdots, N$. For simplicity, assume that each firm is different. Then n may also be used to denote the vector of observable characteristics of the nth firm:

$$n = (\alpha(n), \rho(n), X_g(n), X_d(n)).$$

We assume that firms offer contracts and that investors then allocate their wealth among firms and storage. Let w_{mn} be the wealth invested by investor m

in the nth firm, and let w_{mo} be the wealth placed by investor m in storage. Let $T(n)$ denote an incentive-compatible contract offered by the nth firm. (If this contract requires the efforts of a monitor, then it must satisfy the incentive-compatibility conditions described in the text.) The expected profits of the nth firm under a contract $T(n)$ are denoted by $\pi^f(n, T(n))$, and the contract offered by this firm, $T^*(n)$, must satisfy

$$T^*(n) = \text{argmax } \pi^f(n, T(n)) \tag{A1}$$

subject to

$$\textstyle\sum_m w_{mn} \geq 1. \tag{A2}$$

If this is infeasible, the firm does not operate. Let $\pi^I_m(n, T^*(n))$ denote the expected return per unit of funds invested by investor m in the nth firm. Since each investor can store his or her wealth,

$$\pi^I_m(n, T^*(n)) \geq 1 \tag{A3}$$

for any firm that operates. In equilibrium, contracts must yield equal returns across investors and across firms; i.e.,

$$\pi^I_m(n, T^*(n)) = \pi^I_{m'}(n', T^*(n')) \quad \text{for all} \quad m, m', n, n', \tag{A4}$$

where n and n' are firms that receive funds. Let F denote the set of firms that receive funds. By (1) and (A3), F will contain all firms that have a maximum expected return at least equal to one. Then the description of equilibrium is completed by a market-clearing condition,

$$\textstyle\sum_m(\sum_{n\in F} w_{mn} + w_{mo}) = Mw_m. \tag{A5}$$

Together, (A1) through (A5) describe equilibrium with heterogeneous firms.

REFERENCES

1. M. Berlin and J. Loeys. "The Choice between Bonds and Bank Loans." Working Paper No. 86-18, Federal Reserve Bank of Philadelphia, November 1986.
2. F. Blackwell. "Comparison of Experiments." In J. Neyman (ed.), *Proceedings of the Second Berkeley Symposium in Mathematical Statistics and Probability*. Berkeley: University of California Press, 1951, 93–102.
3. J. H. Boyd and E. C. Prescott. "Financial Intermediary Coalitions." *Journal of Economic Theory* 38 (April 1986), 211–32.
4. T. Campbell and W. Kracaw. "Information Production, Market Signalling, and the Theory of Financial Intermediation." *Journal of Finance* 35 (September 1980), 863–82.
4. Y. Chan. "On the Positive Role of Financial Intermediaries in the Allocation of Venture Capital in a Market with Imperfect Information." *Journal of Finance* 38 (December 1983), 1543–69.
6. D. W. Diamond. "Financial Intermediation and Delegated Monitoring." *Review of Economic Studies* 51 (July 1984), 393–414.
7. G. B. Gorton and J. G. Haubrich. "Bank Deregulation, Credit Markets, and the Control of Capital." Working Paper No. 8-86, Rodney White Center for Financial Research, February 1986.
8. T. H. Ho and R. F. Singer. "Bond Indenture Provisions and the Risk of Corporate Debt." *Journal of Financial Economics* 10 (December 1982), 375–406.
9. A. Kalay. "Stockholder-Bondholder Conflict and Dividend Constraints." *Journal of Financial Economics* 10 (July 1982), 211–33.

10. P. R. Milgrom. "Good News and Bad News: Representation Theorems and Applications." *Bell Journal of Economics* 12 (Autumn 1981), 380–91.
11. R. Ramakrishnan and A. Thakor. "Information Reliability and the Theory of Financial Intermediation." *Review of Economic Studies* 51 (July 1984), 415–32.
12. Robert Morris Associates. *Incentive Compensation Systems for Commerical Loan Officers*. Philadelphia: Robert Morris Associates, 1983.
13. C. W. Smith and J. B. Warner. "On Financial Contracting: An Analysis of Bond Covenants." *Journal of Financial Economics* 7 (June 1979), 117–61.