

COWLES FOUNDATION FOR RESEARCH IN ECONOMICS

AT YALE UNIVERSITY

Box 2125, Yale Station
New Haven, Connecticut 06520

COWLES FOUNDATION DISCUSSION PAPER NO. 729

Note: Cowles Foundation Discussion Papers are preliminary materials circulated to stimulate discussion and critical comment. Requests for single copies of a Paper will be filled by the Cowles Foundation within the limits of the supply. References in publications to Discussion Papers (other than mere acknowledgement by a writer that he has access to such unpublished material) should be cleared with the author to protect the tentative character of these papers.

MANAGERIAL INCENTIVES

and

CAPITAL MANAGEMENT

by

Bengt Holmstrom

and

Joan E. Ricart i Costa

November 1984

MANAGERIAL INCENTIVES

AND

CAPITAL MANAGEMENT

by

Bengt Holmstrom*

and

Joan E. Ricart i Costa**

October 1983

March 1984, first revision

October 1984, second revision

*School of Organization and Management and the Cowles Foundation, Yale University.

**I.E.S.E., Navarra University, Spain.

This research was partially supported by grants from NSF and the Sloan Foundation.

ABSTRACT

In Holmstrom (1982) an example is given, which shows that a manager's concern for the value of his human capital will lead to a natural incongruity in risk-preferences between himself and the owners, even when no effort considerations are involved. In this paper we present a formal model of this channel of incongruity based on learning about managerial talent. We also explore the nature of an optimal incentive contract in the case where the manager may withhold but not misrepresent information about investment returns.

The optimal contract is an option on the manager's human capital value with a possible bonus for investing. The optimal investment rule accepts fewer investments than under full information, which is to say the hurdle rate exceeds the cost of capital--a commonly observed real world feature. Another phenomenon the model helps explain is the extensive use of capital budgeting and rationing schemes in place of linear or non-linear price decentralization, which are shown to be less efficient modes of allocation.

1. Introduction

Moral hazard--the problem of controlling unobservable actions by subordinates--has been extensively studied in recent years. Most of the research has focused on the case where action can be interpreted as effort --for good reasons it seems. Sincerity of labor input is inherently unobservable; an effort interpretation permits further restrictions on the model formulation; and effort, being a costly input, provides a natural source of incongruity in preferences between superiors and subordinates.

In a managerial context, however, effort is only part of the overall incentive problem. It is likely that many executives believe their managers are industrious enough--what they worry more about is how good these managers are at making decisions.¹ Part of this concern relates to the managers' willingness or unwillingness to take risks in short and long run production and investment decisions. Extensive capital budgeting and financial control procedures are concrete signs that the control of investment incentives is an important issue for the firm. To the extent that these incentive problems cannot be fully removed by monitoring--as marginalist thinking would readily suggest--they have allocational consequences not only for the firm but for the economy as a whole. The consequences could be significant in view of the economic impact managerial decisions have. It seems important therefore to take a more careful look at what, if anything, may go wrong in decentralized management.

Why would managers be inclined to act against the interests of shareholders and superiors when it comes to investment decisions? Two early papers in agency theory, Wilson (1969) and Ross (1973), suggested that incentive problems may arise as a consequence of attempts to utilize the

manager's risk absorption capacity. Optimal risk-sharing--barring exceptional cases--leads to incongruities in risk-preferences and a second-best solution would therefore trade off some of the benefits of risk-sharing against better investment incentives. To us, this trade-off seems negligible. The value of the manager as a risk-carrier is minimal in a firm of even modest size, particularly if the firm is publicly held. Consequently, an almost costless solution to the incentive problem could be obtained by paying the manager a flat wage and asking him to act in the best interests of the owners.

Another, more plausible source of incongruity in risk-preferences stems from costly (and unobservable) effort choice. In order to induce effort, the manager has to carry some risk and this leads him to value risky projects differently than his owners or superiors. Grossman and Hart's (1982) generalized principal-agent model can formally account for this channel of incongruity, but as their analysis strongly suggests, little can be said without being much more specific about the production technology. Models with highly simplified technological structures are analyzed in Holmstrom and Weiss (1982) and Lambert (1984).

Our purpose is to explore a third reason for why managers might not act in the best interests of their employer. The reason is simply this. Managers realize that performance will be read as a signal about their future potential. This implies that a manager's opportunity wage will vary with investment outcomes and make the investment choice a risky undertaking from a human capital point of view. In other words, investments (or any other production decisions for that matter) lead to two risky capital streams--one human and one financial. If no explicit incentive structures are set up, the manager is only concerned with the former

while owners are concerned with the latter. To bridge the preference gap, some form of financial incentives are needed.

The argument above was first suggested in Holmstrom (1982). It was made partly in response to Fama's (1980) claim that reputational concerns will police rather well managerial misbehavior even when no explicit contracts are in effect. While it is true that reputation will be a force in inducing effort (how well, depends on several factors; see Holmstrom, 1982), it is also true that it may create divergent interests in other dimensions of decision-making as we have just suggested. It is worth stressing that time, which we usually think of as alleviating incentive problems, is the source of problems in the investment setting we are going to study.

Our purpose here is to propose a formal model in which incongruities in risk-preferences of the kind just mentioned can be analyzed. Our model is one of learning. Managerial ability is an uncertain factor, which gets revealed over time through performance. To make matters analytically tractable, we assume that performance and investment decisions are public information so that everybody shares beliefs about managerial ability both initially and after outcomes have been observed. The manager's expertise lies in spotting potentially valuable projects. He has to propose investments to a superior, who makes the ultimate decision. The manager cannot misrepresent information about a project (this assumption is made to preserve symmetry in beliefs), but he can propose wrong projects or refrain from proposing projects that would increase the value of the firm. The question is what the optimal incentive contract should look like in this informational setting and what investment decisions it leads to.

A key assumption is that employers are unable to insure fully the manager's human capital risk, because involuntary servitude is prohibited.

If a manager turns out to be good, a commensurate wage contract has to be offered or else the manager will leave the firm. This is similar to Harris and Holmstrom (1982). We show then that an optimal contract offers the manager a downward rigid wage schedule. Unlike the Harris-Holmstrom result, the wage may be increased even if the manager's market value does not require an increase. The increase acts as an incentive for bringing proposals for consideration by higher management. With respect to investment decisions we show that hurdle rates for project acceptance are above the cost of capital (appropriately defined). There may be capital rationing and the stylized capital budgeting procedure built into the model is superior to any non-linear scheme of capital pricing. Indeed, the capital budgeting procedure is essential in securing the high degree of insurance being offered to the manager in the optimal contract.

The outline of the paper is as follows. The next section gives an example that illustrates the nature of the incongruity in risk-preferences when there is learning about ability. Section 3 defines the model and provides a preliminary analysis of the value of learning. Section 4 performs the analysis of optimal contracts and investment rules. Section 5 applies the general results to the earlier investment example. Section 6 discusses extensions. The final section offers a summary and a look at future directions of research.

2. Market Outcomes without Contracting--An Example

We begin with an example that illustrates how reputation concerns can lead to divergent investment preferences. In the example the manager is not offered any kind of contingent contract and his payment is therefore equal to his expected marginal product in each period. This assumption is made to rationalize the need for contracts, which we investigate subsequently.

There are two periods, $t = 1, 2$. Both periods are technologically identical and stochastically independent. In each period an investment can be made, which yields a publicly observed net payoff y_t at the end of the period. The manager will decide on whether or not to invest. Before deciding he can observe a signal s_t which is a predictor of the net payoff y_t . Specifically, the net payoff in period t can be written as the sum:

$$y_t = s_t + \epsilon_t,$$

where ϵ_t is a random factor. All random variables (s_1, s_2, ϵ_1 and ϵ_2) are assumed independent.

At the beginning of the period (when the manager is paid his fee) both s_t and ϵ_t are unknown. The possible values of s_t and ϵ_t are 0, -1 and +1, 0, respectively, so that y_t can be either -1, 0 or +1. The distribution of s_t puts equal weights on the two outcomes, while the distribution of ϵ_t depends on the characteristics of the manager. In particular, $\Pr(\epsilon_t = 1) = \eta$, where η is a measure of managerial ability, unchanged between periods.

The manager's ability is uncertain. The uncertainty is encoded in

a subjective probability distribution over η , which we take for concreteness to be $\text{Beta}(\alpha_1, \beta_1)$ in period 1. For a Beta-distribution

$$E(\eta) \equiv \mu_1 = \frac{\alpha_1}{\alpha_1 + \beta_1},$$

and so the unconditional probability of $\varepsilon_1 = +1$ is μ_1 .

We will assume that all relevant parties (the manager and the labor market) share initial beliefs about the random variables including the manager's ability. Furthermore, we assume that if an investment is made, the value of s_t will become public knowledge at the end of the period. Thus, after an investment in period 1, y_1 and s_1 will become known before period 2 starts. Therefore ε_1 can be inferred and all parties will update their beliefs about the manager's ability in the same way. Since Beta is a conjugate family of distributions for Bernoulli trials (see DeGroot, 1970), we have that η at the beginning of the second period will be $\text{Beta}(\alpha_2, \beta_2)$, where

$$(1) \quad \begin{cases} \alpha_2 = \alpha_1 + 1, & \beta_2 = \beta_1, & \text{if } \varepsilon_1 = +1 \\ \alpha_2 = \alpha_1, & \beta_2 = \beta_1 + 1, & \text{if } \varepsilon_1 = 0. \end{cases}$$

Of course, if no investment is made, $y_1 = 0$ and ε_1 will not be possible to infer. Since s_1 does not contain information about η , nobody changes beliefs about η and $\alpha_2 = \alpha_1$, $\beta_2 = \beta_1$ in that case. In all cases $E(\eta) \equiv \mu_2 = \alpha_2 / (\alpha_2 + \beta_2)$ in the second period.

Let us analyze the market outcome recursively. In the second period, since it is the last in the manager's career, he has no reason not to make the socially preferred decision, which is to invest if $s_2 = 0$ and not to

invest if $s_2 = -1$. Notice that this decision rule does not depend on current beliefs about η , because if $s_2 = 0$ then $E(y_2) = \mu_2 > 0$ and if $s_2 = -1$ then $E(y_2) = -1 + \mu_2 < 0$, whatever μ_2 is. The value of the manager, before s_2 is observed, is $\frac{1}{2}\mu_2$ with this optimal rule. Denoting his wage in the second period by w_2 , we have that $w_2 = \frac{1}{2}\mu_2$, because the manager is paid his expected marginal product.

Now go back to the first period. The manager is paid a fee w_1 first. Then he observes s_1 and has to decide on the course of action. If he does not invest he knows that he will be paid $w_2 = \frac{1}{2}\mu_2 = \frac{1}{2}\mu_1$, because beliefs will not change. If he invests he faces a lottery, because μ_2 will depend on the investment outcome (or rather the inferred value of ϵ_1). A high investment outcome (implying $\epsilon_1 = +1$) will give him $w_2^+ \equiv (\alpha_1 + 1)/2(\alpha_1 + \beta_1 + 1)$ and a low investment outcome will give him $w_2^- \equiv \alpha_1/2(\alpha_1 + \beta_1 + 1)$. The expected value of this lottery is:

$$E(w_2) = \mu_1 w_2^+ + (1 - \mu_1) w_2^- = \frac{1}{2}\mu_1.$$

We find then that the option to invest is never attractive to a risk-averse manager. Whatever the value of s_1 , the manager will face the same income lottery if he invests and this lottery will have an expected value equal to what he would get without investing. Consequently, no investments will ever be undertaken in the first period and the marginal product and first-period wage will equal zero.

Our conclusion is not dependent on the two-period structure. With any finite number of periods it is easy to see (by backwards induction) that the manager will only be willing to invest in the very last period. What our conclusion is sensitive to, however, is the specific information

structure of the example. This is useful to understand and we will illustrate it by a slight modification of the distribution of s_t .

Let us change the example so that s_t is uniformly distributed between -1 and 0 . This changes the socially optimal decision rule to the following: invest if and only if $s_t \geq -\mu_t$, where $\mu_t = \alpha_t / (\alpha_t + \beta_t)$ is the expected value of η in period t , with η being distributed as $\text{Beta}(\alpha_t, \beta_t)$ as before. Consequently, the value of the manager in the second period and his fee will be:

$$\int_{-\mu_2}^0 (\mu_2 + s_2) ds_2 = \frac{1}{2} \mu_2^2.$$

As before,

$$(2) \quad E(\mu_2 | s_1) = \mu_1, \text{ for any } s_1,$$

because of the martingale property of conditional expectations. Jensen's inequality implies therefore:

$$(3) \quad E(w_2 | s_1) = E\left(\frac{1}{2} \mu_2^2 | s_1\right) > \frac{1}{2} \mu_1^2.$$

The last term is the manager's fee in the second period if he does not make an investment, while the middle term is the expected value of the manager's fee in the second period if he does make an investment. The reason the latter is higher is because information is obtained about ability if an investment is made. This is valuable (unlike in the preceding case), because the optimal rule now depends on assessed ability.

We conclude from (3) that a manager may want to invest in the first period if he is not too risk averse. Unfortunately, he will either invest

for all s_1 or for no s_1 . This is so, because the distribution of μ_2 is independent of s_1 (cf. (2)), given our assumption that s_1 is publicly observed. Consequently, the manager is either paid $w_1 = 0$ (if he is of the type that never invests) or $w_1 = \mu_1 - \frac{1}{2}$ (if he always invests). In either case his fee is less than he could obtain if a socially optimal decision rule were followed.

Both variations of the example make the point that there may be a significant incongruity in preferences due to the manager's career concerns. Without contracting the manager only worries about the human capital return of the investment, while the owners only look at financial returns. The divergent interests are so severe in the example that the manager's private information about s_1 is left entirely unexploited. In order to capitalize on his expertise, some form of contingent contracting is necessary.²

We will return to the optimal contracting solution of this example in Section 5.

3. A General Model

For notational simplicity we assume there are only two periods, $t = 1, 2$. The periods are technologically identical and stochastically independent. In each period a publicly observed action $a_t \in A$ is taken, which yields a net payoff

$$(4) \quad y_t = y_t(a_t, \eta_t, \theta_t) ,$$

where η_t is a quantified measure of managerial ability and θ_t is a random state of nature. Ability is measured so that y_t is increasing in η_t . There is always the opportunity of no action, a_0 , which yields a return that is independent of θ_t and we may as well assume that the net payoff in that case is identically zero.

The owner is risk neutral and the manager is risk averse with an atemporal utility function over consumption given by

$$(5) \quad u(c_1, c_2) = u(c_1) + \beta u(c_2) .$$

Both parties use the same discount factor $0 < \beta < 1$. The manager cannot borrow or save and will therefore consume his income in each period.

The manager gets to observe a signal, $s_t \in S$, about the payoff prospects before a_t is determined. This rationalizes his managerial position as a scout in search of valuable investment opportunities. Having observed s_t the manager is expected to report his information to the owner, who will make the final decision a_t . The manager has two options in reporting. He can either conceal his information by claiming that he did not observe anything worthwhile to invest in; in effect, he can veto investments by not suggesting any. Alternatively, he can present the true

s_t -value which we may interpret as an investment proposal. We assume that the validity of the information in the investment proposal can be verified by the owner. Thus, the manager may withhold but not misrepresent his information. This is of significance for retaining symmetry in the learning process.

The manager and the owner share beliefs about ability, signals and states. These beliefs are encoded in the following probability assessments.

- $p_1(\eta_1)$ = the prior distribution of η_1
- $r(s_t)$ = the stationary distribution over s_t
- $m(y_t, x_t | a_t, s_t, \eta_t)$ = the stationary conditional distribution of (y_t, x_t) given a_t , η_t and s_t , where x_t represents a vector of other observable variables.

The vector x_t may be interpreted as any other information that has bearing on the ability evaluation. One may observe how competitors perform, what the general economic circumstances are, etc. To the extent that these pieces of information are related to the manager's performance, this will influence the updated beliefs about ability.

Our formulation does not include all the primitive probability assessments for notational simplicity. For this reason we note the following. Ability (η) and state (θ) may be dependent. The signal (s) and the state (θ) are generally dependent as reflected in m . But the signal (s) and ability (η) are independent. The reason for this last assumption is that we do not want the manager to be better informed about his ability in case he decides not to report his signal to the owner. This assumption implies that ability has a purely productive function-- it is a measure of management skills rather than forecasting skills.

There is an outside labor market in which the manager can find alternative employment. For simplicity, we start with the assumption that this market has the same technological opportunities as the firm in which the manager is presently employed. Moreover, we assume that the market shares the same information as the owner, i.e., it starts with identical priors and gets to observe y_t , x_t , a_t and also s_t in case the signal is reported and an investment is made. The information about s_t we can also think of as revealed ex post together with the investment outcome. This is not so implausible, particularly since we really only need to assume that enough is revealed to make the same inference about ability from output as the manager and the owner does. Plenty of information about circumstances surrounding the investment outcome will generally be available ex post and that may well be sufficient for knowing s_t .

To sum up, the sequence of events within each (identical) period is as follows. First the manager is paid a wage, which he consumes instantly. (This wage need not be his marginal product in the period, since he can be on a contract.) Next, the manager observes a signal, which he either reports truthfully to the owner or withholds from reporting. The owner decides on what action to take. Once the action is taken, the payoff and the other observable values are revealed. All parties, including the market, observe the signal, the action, the outcome and the additional data if an investment is made.

The Value of Learning

The key feature of our model is learning about ability. When an action is chosen in the first period it selects a distribution over outcomes (y,x) , which will be used to update beliefs about ability. In statistical language, the first period action determines an experiment,

which is completely described by the family of conditional distribution $\{m(\cdot|a_1, s_1, \eta_1)\}$ with η_1 as the unknown parameter.

Let $h_1 = (y_1, x_1, a_1, s_1)$ denote the history of events in period 1. After h_1 is observed, all parties will revise their assessments of the manager's ability to a posterior distribution denoted \hat{p}_1 . As a function of h_1 , \hat{p}_1 is itself random. It is well known from statistics that the distribution over p_1 is an equivalent and complete description of the experiment about ability. Moreover, the stochastic process, taking priors to posteriors, will form a martingale, i.e.

$$(6) \quad E[\hat{p}_1 | a_1, s_1] = p_1,$$

where the expectation is taken over x_1 and y_1 conditional on the first-period action and signal. The more variability there is in \hat{p}_1 (in the sense of a mean-preserving spread), the more informative the experiment is for learning about ability (Blackwell (1953)).

Let p_2 denote the probability assessment over the manager's ability at the beginning of period 2. If $\eta_1 = \eta_2$, as was the case in the introductory example, then $p_2 = \hat{p}_1$. However, one could well imagine that ability (or productivity) changes over time and in particular that the change is a function of the decisions made in earlier periods. For instance, involvement in a project may increase the manager's experience in a way that enhances his future productivity. We will include this as a possibility by assuming that

$$(7) \quad p_2(h_1) \succeq_1 \hat{p}_1(h_1), \text{ for all } h_1.$$

Here \succeq_1 refers to (weak) first-order stochastic dominance. The assumption in (7) allows ability to increase either deterministically or

stochastically. Combining (6) and (7) gives:

$$(8) \quad E[p_2 | a_1, s_1] \geq_1 E[\hat{p}_1 | a_1, s_1] = p_1 .$$

Note that the two expectations are averages of distributions and therefore compared by stochastic dominance.³

We proceed to determine the manager's opportunity wage and the value of learning. Let

$$(9) \quad \bar{y}(a, s, p) = \int ym(y, x | a, s, \eta) p(\eta) dy dx d\eta .$$

This represents the expected value of output given an action a , a signal s and a manager whose ability is assessed by p . Subscripts are omitted because the formula applies to either period. Let $\alpha : S \rightarrow A$ be a decision rule, which describes that the action $\alpha(s) \in A$ is chosen if $s \in S$ is the signal observed. Define

$$Y(\alpha, p) = \int \bar{y}(\alpha(s), s, p) r(s) ds ,$$

as the expected output before the signal is observed of a manager p when the decision rule α is used. The decision rule that maximizes expected output is given by

$$(10) \quad \alpha_p^* \equiv \operatorname{argmax}_{\alpha} Y(\alpha, p) .$$

With this decision rule the one-period value of a manager with ability p is

$$(11) \quad z(p) \equiv Y(\alpha_p^*, p) .$$

It is important to note that \bar{y} is a linear function of p and

so is $Y(\alpha, p)$. It follows that $z(p)$ is a convex function, because it is the maximum (by (10)) of a set of linear functionals.

Let the manager's marginal product in period t be z_t . It depends, of course, on the decision rule and p_t . By (10) we will always have $z_t \leq z(p_t)$. For the second period we know, however, that the manager has no incentive to deviate from the optimal rule (10), because it is the last period of his career (recall the same argument in the introductory example). Consequently, $z_2 = z(p_2)$. Combining the convexity of $z(\cdot)$ with (6) and (8) we find that the expected value of the manager's marginal product in the second period, conditional on any first-period action a_1 and signal s_1 , satisfies:

$$(12) \quad E[z_2 | a_1, s_1] \geq E[z(\hat{p}_1) | a_1, s_1] \geq z(p_1) .$$

Let the value of learning be defined as the difference in (12):

$$(13) \quad \begin{aligned} V_{p_1}(a_1, s_1) &\equiv E[z_2 | a_1, s_1] - z(p_1) \\ &= E[z(p_2) - z(\hat{p}_1) | a_1, s_1] + E[z(\hat{p}_1) - z(p_1) | a_1, s_1] . \end{aligned}$$

In (13) we have separated out the two components of value increase. The first one measures the direct increase in human capital due to learning by the manager. The second part measures the indirect increase in human capital due to learning about the manager. As the introductory example made clear, the value of learning about the manager depends on how sensitive the decision rule α^* is to the ability assessment p . The more α^* responds to p , the more convex will $z(\cdot)$ be and the more it is worth to know ability accurately. On the other hand, if α^* is independent of p , then there are no gains to information about η .

The returns to learning also increase with the variability in p_2 . This is precisely the reason why more informative experiments (in Blackwell's sense) have higher value. On the other hand, keeping the experiment (i.e., the investment) the same, one would think that learning would be more valuable the less is known initially about the manager. We do not know whether this is true in general, but it holds for some distributions. For instance, the interested reader can check that the beta distribution (used in the example) has this property. For a younger manager with the same average ability as an older one, investments will be more informative (leading to higher variability in p_2) and hence more valuable from a learning point of view. There is an offsetting factor in the overall value of the investment, however. If the payoff function (8) is concave in ability, then the higher variance in η for the younger manager makes his average financial product smaller. Indeed, firms reluctantly place unproven skills in responsible positions (where payoff is sensitive to ability) no matter how great the potential is, while for less responsible tasks younger managers may be preferred because of learning.

4. Optimal Contracting

A contract δ , for the two-period model, is a pair (ω, α) where $\omega = (w_1, w_2(h_1))$ is the wage contract and $\alpha = \alpha_1(s_1)$ is a decision rule for the first period. The second period decision rule will always be the one maximizing expected output $(\alpha_{p_2}^*)$ and we therefore omit it from the contract.

A contract will be agreed upon at the very beginning of period 1. The first period wage is paid before s_1 is observed so it is uncontingent. The second period wage will be paid before s_2 but after h_1 is known, so it can be contingent on h_1 . For the time being we also allow the decision rule to be agreed upon in a binding manner at the time of contracting.

Define

$$(13) \quad \begin{cases} U(\omega, a_1, s_1) \equiv u(w_1) + \beta E[u(w_2) | a_1, s_1] \\ V(\omega, a_1, s_1) \equiv E[y_1 + \beta z_2 | a_1, s_1] - w_1 - \beta E[w_2 | a_1, s_1] . \end{cases}$$

The expressions in (13) give the manager's expected utility and the owner's expected profit as functions of the action a_1 , the signal s_1 , and a wage contract ω . We have tried to simplify notation by dropping arguments and omitting explicit references to distributions. Recall, however, that w_2 is a function of $h_1 = (y_1, x_1, a_1, s_1)$ and so is $z_2 = z(p_2)$, because p_2 depends on h_1 via the updating rule. Thus, with a_1 and s_1 fixed, all expectations are taken with respect to the conditional distribution of (y_1, x_1) , i.e., $m(\cdot, \cdot | a_1, s_1, \eta_1)$.

We will proceed by solving a sequence of optimization programs, which correspond to increasingly constrained notions of feasible contracts.

As a bench mark case we start with what we call first-best, where the ideal Pareto optimal contract is designed.

A. First-Best

The contract that maximizes the gains from cooperation solves:

$$(A) \quad \max_{\delta} \int U(\omega, \alpha(s_1), s_1) r(s_1) ds_1 ,$$

$$\text{s.t. } \int V(\omega, \alpha(s_1), s_1) r(s_1) ds_1 \geq 0 .$$

The only constraint here is that the owner makes zero profits ex ante.

If not, the manager could receive a better initial offer from the market, which we recall has the same technological opportunities. Note, however, that once the manager accepts a contract he cannot renege on it later.

The manager commits himself irrevocably to the firm. Furthermore, in (A) we are assuming that the manager will report s_1 as observed, though as we will see he will have no incentive not to, given the optimal contract.

Proposition 1. The first-best solution satisfies:

- (a) $w_1 = w_2(h_1) = w$; a constant wage;
- (b) $\alpha(s_1) = \operatorname{argmax}_{a_1} E[y_1 + \beta z_2 | a_1, s_1]$.

Proof. Whenever $\alpha(\cdot)$ is being used, it is optimal to insure the manager fully; hence (a). Given that fact, the optimal $\alpha(\cdot)$ satisfies (b).

Q.E.D.

Part (b) of the proposition states that the optimal decision rule maximizes the expected net present value of the firm's total capital, $v = y_1 + \beta z_2$, which is the discounted sum of its financial return y_1 and the human capital return z_2 . Of course, if learning has no value,

then the expected value of human capital is independent of the decision rule and the traditional conclusion obtains that expected one-period profits should be maximized. In general, though, learning is of value and then the relevant objective is to maximize total capital.

Note that because of (a), the manager would indeed have no incentive to withhold his signal information.

B. Indenture Prohibited; Signal Public

A manager whose performance in the first period is such that his second-period value z_2 exceeds the constant contractual wage w , would be able to obtain a better offer from the market. Since legal constraints de facto prevent him from committing himself to the firm in advance, he must be expected to leave the firm if he is paid less than his marginal product in the second period. Adding this constraint to program (A) leads to:

$$(B) \quad \max_{\delta} \int U(w, \alpha(s_1), s_1) r(s_1) ds_1 ,$$

$$\text{s.t. (i) } \int V(w, \alpha(s_1), s_1) r(s_1) ds_1 \geq 0 .$$

$$\text{(ii) } w_2(h_1) \geq z_2 = z(p_2(h_1)) , \quad \forall h_1 .$$

In (ii) we have made explicit reference to the fact that p_2 , hence z_2 , depends on h_1 ; $z(p)$ was defined before as the second-period value of a manager who uses the decision rule which maximizes the one-period output.

We are still assuming that the manager reports his signal (s_1) truthfully. One case, consistent with our general formulation, is that no s_1 is observed. We could interpret a_1 in this case as a job assignment rather than an investment decision. Alternatively, of course, we

could think of s_1 as a public signal.

Proposition 2. The optimal solution without indenture satisfies:

$$(a) \quad w_1 \leq z_1 \leq z(p_1) ,$$

$$(b) \quad w_2(h_1) = \max\{w_1, z_2\} , \quad \forall h_1 , \quad \text{where } z_2 = z(p_2(h_1)) ,$$

$$(c) \quad \alpha(s_1) = \underset{a_1}{\operatorname{argmax}} \{E[y_1 + \beta z_2 | a_1, s_1] + H_{w_1}(a_1, s_1)\} , \quad \forall s_1 , \quad \text{where}$$

$$(15) \quad H_{w_1}(a_1, s_1) \equiv \frac{\beta}{u'(w_1)} E[u(w_2) - u'(w_1)w_2 | a_1, s_1] .$$

Proof. Let λ be the multiplier of B(i) and maximize the Lagrangian subject to (ii). Pointwise optimization of wages yields $\lambda = u'(w_1) = u'(w_2(h_1))$ if $w_2(h_1) \geq z_2$; otherwise $w_2(h_1) = z_2$ by (ii). From this (b) follows. The zero profit constraint (which determined λ and therefore w_1) implies (a). Therefore (a) and (b) must be satisfied for any decision rule α . Maximizing with respect to $\alpha(\cdot)$, and noting that $\lambda = u'(w_1)$, we obtain (c).⁴

Q.E.D.

The inequality in (a) will be strict if any learning at all takes place. The reason is that (b) obviously prevents the owner from making profits in the second period; if z_2 varies he will in fact make expected losses. To compensate for this loss the manager pays an insurance premium $z_1 - w_1 > 0$ in the first period by accepting a wage below his marginal product.

The optimal wage policy is familiar from before (see e.g., Harris-Holmstrom, 1982)). The risk neutral owner insures the manager to the extent possible, which means he guarantees the manager at least w_1 for the second period. If the manager's marginal product z_2 exceeds w_1 in the

second period the owner pays the marginal product and earns zero profits. Departing slightly from the terminology in Harris-Holmstrom we call this wage policy downward rigid (since it is nondecreasing in time) and tight (because the owner never raises the wage more than necessary to retain the manager).

An important feature of the optimal policy is that it makes the wage depend on output only indirectly through z_2 . As far as the wage is concerned, z_2 is a sufficient statistic for h_1 . If the family of likelihood functions $\{m(\cdot, \cdot | a_1, s_1, \eta_1)\}$, characterizing the first-period experiment about ability, satisfies the monotone likelihood ratio property (familiar from statistics)--which is natural to assume--then z_2 and hence wage will be an increasing function of output. But beyond this, little can be said about the relationship between wage and output without being explicit about distributions.

This ambiguity, also common in moral hazard models, may appear disturbing from a positive point of view. But in fact it is not, because the wage is tied to the value of the firm in a straightforward way. In the simplest interpretation of the model the manager is the only employee of the firm (or alternatively, the firm cannot function without him). Hence, the value of the manager, z_2 , is also the value of the firm. It follows that the downward rigid wage structure is a stock option of the whole firm. In general, of course, the manager will not be the only input factor, but to the extent his human capital value is aligned with the total value of the firm, a stock option (of a fraction of the firm) will be a good proxy for the wage schedule in Proposition 2. This interpretation seems both realistic and important in view of common compensation practice.

Next, consider the optimal decision rule. As (c) makes clear, the

optimal rule will differ from the first-best rule in which the total value of the firm was maximized. Because of constraint (ii) the manager is forced to bear some residual risk and a premium has to be charged accordingly. This premium, and hence the optimal rule, depends on the manager's utility function. It is determined by the appended term $H_{w_1}(a_1, s_1)$ in (15), which can be interpreted more easily by approximating utility with the first two terms of a Taylor series expanded around w_1 . We have:

$$(16) \quad H_{w_1}(a_1, s_1) \simeq \frac{\beta}{u'(w_1)} E \left[u(w_1) - u'(w_1)w_1 + \frac{1}{2}u''(w_1)(w_2 - w_1)^2 \right].$$

Substituting (16) into (c) tells us that $\alpha(s_1)$ is determined by maximizing

$$(17) \quad E[y_1 + \beta z_2 | a_1, s_1] - \frac{\beta R(w_1)}{2} E[(w_2 - w_1)^2 | a_1, s_1],$$

where $R(w_1)$ is the Arrow-Pratt measure of absolute risk aversion.

Expression (17) shows that for small risks the cost of risk is proportional to the manager's absolute risk aversion. The relevant measure of risk is the variance in w_2 (w_1 is fixed). Since $w_2 = \max(w_1, z_2)$ it is only human capital risk that is of concern. This risk is bigger the more informative the selected decision is in terms of ability. Learning about ability is bad in this regard.

Learning is good in another regard. It increases on average the value of human capital to the extent information about ability improves future decision-making. This offsetting factor is embedded in the difference $E[z_2 | a_1, s_1] - z(p_1)$, which is nonnegative by (12). That learning has both a cost and a benefit prevents unambiguous comparisons with decision-making under complete information about ability. A project may be

worth more or less than the value of its financial return stream, depending on whether learning has a net positive or negative value. In general, one can only say that the optimal decision rule will be biased in favor of financial returns when compared to the first-best decision rule, because human capital risk is costly. This will be illustrated shortly in the solution to the 0-1 investment problem.

A final remark about the optimal wage policy is in order. While its structure is the same as in the case without any decisions (see Harris and Holmstrom), one appealing feature is lost. Without decisions it turns out that the optimal wage policy is independent of the manager's utility function. In particular, w_1 is uniquely determined by the firm's zero profit constraint. However, once decision-making is introduced we can see from (17) that $\alpha(\cdot)$ is sensitive to the manager's utility function and consequently w_1 will depend on it as well.

C. Indenture Prohibited; Signal Private

Next we turn to the main case, which incorporates the simplest incentive problem. The manager observes s_t privately but he may choose not to report it. This way he can veto any investment project and it assures him a wage equal to $z(p_1)$ in the second period. If he chooses to report the signal, however, the truthfulness of the report can be verified. The problem to solve is:

$$\begin{aligned}
 \text{(C)} \quad & \max_{\delta} \int U(w, \alpha(s_1), s_1) r(s_1) ds_1, \\
 \text{s.t. (i)} \quad & \int V(w, \alpha(s_1), s_1) r(s_1) ds_1 \geq 0, \\
 & \text{(ii) } w_2(h_1) \geq z_2 = z(p_2(h_1)), \quad \forall h_1, \\
 & \text{(iii) } U(w, \alpha(s_1), s_1) \geq u(w_1) + \beta u(z(p_1)), \quad \forall s_1.
 \end{aligned}$$

The last constraint makes reporting desirable for all signal outcomes because the manager's veto can be built into $\alpha(\cdot)$. Note also that (iii) is actually independent of β and w_1 once the expression is simplified. The solution to problem (C) is given by:

Proposition 3. The optimal solution without indenture and privately observed signals satisfies:

- (a) $w_1 \leq z_1 \leq z(p_1)$,
- (b) $w_2(h_1) = \max\{w_1, b(s_1), z_2\}$, $\forall h_1$, where $z_2 = z(p_2(h_1))$,
- (c) $\alpha(s_1) = \underset{a_1}{\operatorname{argmax}}\{E[y_1 + \beta z_2 | a_1, s_1] + H_{w_1}(a_1, s_1)\}$,

where $H_{w_1}(\cdot, \cdot)$ is as in (15) and $b(s_1)$ is the solution of

$$(18) \quad u(z(p_1)) = E[u(\max\{b(s_1), z_2\}) | \alpha(s_1), s_1].$$

Proof. Fix a decision rule α . It is easy to check that the proposed solution is feasible for α . (Note that the wage depends on α .) Define

$$(19) \quad c(s_1) = \max\{w_1, b(s_1)\}$$

and the following multipliers:

$$(20) \quad \mu(s_1) = \frac{u'(w_1)}{u'(c(s_1))} - 1 \geq 0 \quad \text{for (iii),}$$

$$(21) \quad \lambda = u'(w_1) > 0 \quad \text{for (i),}$$

$$(22) \quad \gamma(h_1) = \frac{u'(w_1)}{u'(c(s_1))} [u'(c(s_1)) - u'(w_2(h_1))] \geq 0 \quad \text{for (ii),}$$

where w_1 , $b(s_1)$ and $w_2(h_1)$ are as given in the proposition.

By using the above multipliers and premultiplying each equation

by its unconditional probability (and β when necessary) the unconstrained problem can be written as

$$(23) \quad \text{Max}_{\delta'} \left\{ \int [(1 + \mu(s_1))U(w', \alpha'(s_1), s_1) - \mu(s_1)(u(w_1') + \beta u(z(p_1))) + \lambda V(w', \alpha'(s_1), s_1) + BE[\gamma(h_1)(w_2'(h_1) - z_2) | \alpha'(s_1), s_1)]r(s_1)]r(s_1)ds \right\},$$

which is concave in wages. Therefore the unique maximum is obtained by the first-order conditions:

$$(24) \quad \lambda = u'(w_1),$$

$$(25) \quad (1 + \mu(s_1))u'(w_2(h_1)) - \lambda + \gamma(h_1) = 0,$$

which are satisfied by our proposed solution.

Next we check that the complementary slackness conditions are satisfied. If (ii) is not binding, $c(s_1) > z_2$ and therefore $w_2(h_1) = c(s_1)$ which implies $\gamma(h_1) = 0$. If (iii) is not binding for s_1 , we must have $w_1 = c(s_1) > b(s_1)$ and therefore $\mu(s_1) = 0$. Furthermore all multipliers are nonnegative and (i) is always binding (to obtain w_1).

Therefore there exists a set of multipliers that lead to a Lagrangian optimum satisfying the complementary slackness conditions. Thus by weak duality the solution is optimal for any given decision rule α .

The decision rule in (C) is obtained by maximizing the unconstrained problem (23) with respect to α' (which can be done pointwise). Finally, (a) is a direct consequence of the zero-profit condition and (C)(ii).

Q.E.D.

The proposition can be understood as follows. Start with the solution to program (B) as given by Proposition 2. This ignores constraint

(iii). If the manager is willing to report s_1 for all its realizations, given the optimal downward rigid and tight wage policy (that is, if he always prefers $\alpha(s_1)$ to no investment) then we have a solution to (C) because (iii) will not be binding. Indeed, one might think this will always be the case because the manager's wage is insured downwards. That is incorrect. Since $w_1 < z(p_1)$ in general, a sufficiently risk averse manager could very well value the guaranteed wage increase $z(p_1) - w_1 > 0$ more than the uncertain increase $\max\{0, z(p_2) - w_1\}$. If this is the case for some s_1 for which the owner would prefer to invest, then the optimal way to induce a report from the manager is to add a bonus to his wage guarantee, which raises it to $b(s_1)$. The minimum required bonus is given by (18). Such a bonus is the cheapest incentive device because the value of money is highest at the lowest utility level. Also, there is no way of threatening the manager if he does not report, because he can quit and be assured $z(p_1)$. It is a familiar phenomenon (see, e.g., Becker and Stigler (1974)) that bribing is the second-best alternative when punishments are infeasible.

As the wage guarantee is raised, investment becomes less desirable for the owner and more desirable for the manager. In a model with a continuum of signals (as in the example of Section 2), the marginal signal value will be such that both the manager and the owner are indifferent between investing or not. Of course, since the owner never gains from the human capital return, it must be that $E[y_1|s_1] > 0$ whenever bonuses are paid.

The need for a bonus will make the manager's wage more variable than if the signal s_1 were publicly observed. This will bias the decision-rule further towards emphasizing a project's financial returns. In other

respects our earlier discussion of program (B) applies to (C) and there is no need to repeat it here. Instead, let us make two worthwhile observations about the general structure of (C).

First, we wish to stress the significance of the reporting procedure in the model. The owner could not insure the manager to the extent he does in Proposition 3, unless he retained some control over the decision process. A downward rigid wage schedule makes the manager more eager to take risks; if no bonuses are needed for inducing investments, the manager will in fact prefer to take more risk than the owner. Without reporting, such incentives would have to be counter-acted by forcing the manager to carry additional risk. With reporting, however, the problem is dealt with by rationing: the owner can veto undesirable projects. This is clearly more efficient than using investment and/or outcome contingent payment schedules, because maximum insurance is assured. That way the second-best decision rule will come as close as possible to the first-best rule. Our reporting procedure can naturally be interpreted as investment budgeting of the kind observed in firms. Our model suggests that investment budgeting, including non-price rationing of capital, is rational in view of the need to provide management with desired protection against human capital risks as well as a fair evaluation of performance.

The second thing worth noting about the solution to (C) is that in the end it only depends on how the various decision rules $\alpha(\cdot)$ map into joint distributions over the financial return y_1 and the human capital return z_2 . In other words, if we had started with joint conditional distributions $f(y_1, z_2 | \alpha(\cdot))$ --one for each decision rule $\alpha(\cdot)$ --and a marginal distribution $r(s_1)$ for s_1 , we would have had all the necessary information for computing the optimal scheme characterized in

Proposition 3. The finer structure of the information and learning processes that we discussed, is only relevant for the economic interpretations. This observation is useful for discussing extensions later on. Also, it makes precise the sense in which our model of management is about the control of two capital streams, one financial and one human. Since the insurance market is incomplete (due to restricted indenture), one cannot combine the two streams into a single financial measure (the total value of capital). The optimal solution will generally be sensitive to both dimensions separately.

5. The 0-1 Investment Case

As an illustration of the foregoing analysis, we will consider the solution to the contracting problem of the introductory example in Section 2. We concentrate on the case where the signal s_t has a uniform distribution on $(-1,0)$.

Recall that the value of the manager in the second period is $z_2 = z(p_2) = \frac{1}{2}\mu_2^2$, where $\mu_2 = \alpha_2/(\alpha_2 + \beta_2)$ is the probability of a successful investment. There are only three possible values for μ_2 , given μ_1 . If no investment is made then $\mu_2 = \mu_1$. If an investment is made, then $\mu_2 = \alpha_1/(\gamma_1 + 1)$ if a success occurs and $\mu_2 = (\alpha_1 + 1)/(\gamma_1 + 1)$ if a failure occurs, where $\gamma_1 = \alpha_1 + \beta_1$. Note that μ_2 does not depend on the value of s_1 , because $\gamma_1 = s_1 + \epsilon_1$ and s_1 is publicly observed so that ϵ_1 can be inferred. For notational simplicity, we let z_2^0 , z_2^+ and z_2^- stand for the three possible values of the manager in the second period (with obvious meanings).

The value of learning from an investment is according to (12):

$$(26) \quad V_{p_1} = E[z_2 | s_1] - z_2^0 = \frac{1}{2}\text{Var}(\mu_2) = \frac{\mu_1(1 - \mu_1)}{(\gamma_1 + 1)^2} > 0.$$

This value is independent of s_1 . As a slight aside we can note that $\gamma_t = \alpha_t + \beta_t$ measures the precision of information about ability. From the updating rules of the Beta-distribution (see (1)) we have that $\gamma_{t+1} = \gamma_t + 1$. Thus, the precision increases linearly with the number of observations. From (26) we see that the value of learning is inversely related to the precision of knowledge and goes to zero as γ_t goes to infinity, which it does as the number of observations grows large. An

economic interpretation is that the value of learning is greater for young than for old managers (with the same average ability), which we made reference to already earlier.

From (26) it follows that the first-best decision rule $\alpha^*(s_1)$, would accept a project if and only if:

$$(27) \quad s_1 + \mu_1 + \beta \cdot V_{p_1} \geq 0 .$$

This rule maximizes the total discounted value of capital. Evidently, it can be restated as: invest if and only if $s_1 \geq s_1^*$, where s_1^* is the hurdle rate that solves (27) as an equality.

The second-best decision rule depends on the wage structure. We know that $w_2 = \max\{w_1, b(s_1), z_2\}$. In the special case we are considering, $b(s_1)$ is independent of s_1 , because z_2 is. Thus, $b(s_1) = b_1$, where b_1 is the solution of (see (18)):

$$(28) \quad u(z_2^0) = \mu_1 u(z_2^+) + (1 - \mu_1) u(b_1) .$$

Using part (C) of Proposition 3, it follows that the second-best rule $\alpha(s_1)$, accepts a project if and only if:

$$(29) \quad s_1 + \mu_1 + \beta \cdot V_{p_1} + \frac{\beta}{u'(w_1)} E[u(w_2) - u'(w_1)w_2] \\ \geq \frac{\beta}{u'(w_1)} [u(z_2^0) - u'(w_1)z_2^0] .$$

Again, the rule takes the form: accept if and only if $s_1 \geq \bar{s}_1$, where \bar{s}_1 is the hurdle rate that solves (29) as an equality.

We can make the following comparison between the hurdle rates s_1^* and \bar{s}_1 :

Proposition 4: If the manager is strictly risk-averse then $\bar{s}_1 > s_1^*$, i.e. the hurdle rate is higher in the second-best than the first-best solution.

Proof: By comparing (27) and (29) it suffices to show that:

$$(30) \quad E[u(w_2) - u'(w_1)w_2] - [u(z_2^0) - u'(w_1)z_2^0] \geq 0 .$$

Define the function

$$g_{w_1}(x) \equiv u(x) - u'(w_1)x .$$

This function is strictly concave since u is. It is increasing for $x < w_1$ and decreasing for $x \geq w_1$. By the definition of b_1 ,

$$u(z_2^0) \leq E[u(\max\{z_2, b_1\})] ,$$

where z_2 , of course, depends on the investment outcome. Since $w_2 = \max\{w_1, z_2, b_1\}$ if an investment is made,

$$u(z_2^0) \leq E[u(w_2)] < u(E[w_2]) ,$$

by Jensen's inequality. Therefore, $z_2^0 < E[w_2]$.

Now, using the fact that $g_{w_1}(\cdot)$ is concave and decreasing for $x \geq w_1$, it follows by a second application of Jensen's inequality that

$$g_{w_1}(z_2^0) \geq g_{w_1}(E[w_2]) > E[g_{w_1}(w_2)] ,$$

which is the same as (30).

Q.E.D.

The result in Proposition 4 is intuitive. The first-best rule maximizes the total value of capital (financial plus human). At the critical value s_1^* , one is indifferent between investing or not, provided that the manager is offered full insurance. However, the second-best solution cannot offer full insurance, and hence a marginal investment becomes unprofitable. When a risk premium for human capital is added, it raises the critical value for investment above s_1^* . Fewer projects will be undertaken in second-best than in first-best and the rule is therefore biased towards discounting the value of human capital. One could also interpret Proposition 4 as saying that the internal cost of capital is above the market cost of capital, which is a common observation about firms. However, in making this statement one should notice that the return stream includes the human capital value.⁵

As we already noted in the general analysis, one cannot compare \bar{s}_1 with the optimal hurdle rate \hat{s}_1 for the one-period problem, which is given by:

$$(31) \quad \hat{s}_1 + \mu_1 = 0 .$$

Since $V_{p_1} > 0$, we have that $s_1^* < \hat{s}_1$, but $\bar{s}_1 \gtrless \hat{s}_1$ depending on the cost of inducing risk H_{w_1} on the manager and the value of learning, V_{p_1} .

We proceed to discuss the wage structure with the objective to show that a bonus is sometimes necessary and sometimes not. Start with the assumption that a bonus is not needed. Then $w_2 = \max\{w_1, z_2\}$. The first-period wage w_1 is determined by the zero-profit constraint. This constraint can be expressed as:

$$(32) \quad \int_{\bar{s}_1}^0 [(s_1 + \mu_1 + \beta\mu_1(1 - \mu_1))(z_2^- - w_1)] ds_1 - w_1 = 0 .$$

This expression can be explained as follows. If $s_1 < \bar{s}_1$ then $w_2 = z_2^0$ and the firm makes zero profits in the second period. Also, if the investment is a success then $w_2 = z_2^+$ and the same conclusion obtains. Thus, what remains is the case where an investment is made and it ends up as a failure. Then the firm pays $w_2 = w_1$ but earns only $z_2^- < w_1$ in the second period. The net value of the investment conditional on failure is the integrand in (32). Integrated over the values of s_1 for which an investment is undertaken, we have the expected benefit of the decision rule to the firm, which has to equal w_1 .

We see that the optimal cut-off value \bar{s}_1 depends on w_1 in (29); and reversely, w_1 depends on \bar{s}_1 in (32). By solving the two equations simultaneously, we get the values for \bar{s}_1 and w_1 determined. Now, if $w_1 > b_1$ as defined in (28), then no bonus is in fact necessary and we have the true solution to (C).

One case for which $w_1 > b_1$ is if the manager is risk neutral (or close to it).⁶ This follows directly from (28), which implies that $b_1 < z_2^-$, since a risk neutral manager always prefers to invest even without any insurance because $E[z_2] > z_2^0$. We have therefore $b_1 < z_2^- \leq w_1$ (by the zero-profit constraint (32)).

It is worth noting that when no bonus is required, there will be values of s_1 (close to \bar{s}_1) such that the firm loses from the investment even though $s_1 > \bar{s}_1$. The reason is the following. Since no bonus is necessary, the manager will prefer to invest for all values of s_1 (recall that z_2 is independent of s_1). If \bar{s}_1 were such that the firm also preferred to invest for some values below \bar{s}_1 , then it would be

Pareto improving to lower the hurdle rate. Hence, marginal investments must yield negative expected profits to the firm. The only reason the firm goes along with investing in these cases is that we allow the contract to specify a binding decision rule $\alpha(\cdot)$. Ex ante it is desirable to give the manager the right to invest even in some cases where it does not pay the firm to invest, because the manager captures the value from increased human capital. (We will discuss in the next section the case where commitments to invest cannot be made.)

Since the manager always wants to invest when no bonus is needed in the contract, the firm is effectively rationing the manager when $s_1 < \bar{s}_1$. We commented already earlier on the importance of using rationing as a means for offering maximal insurance.

Now, suppose $w_1 < b_1$ when one solves (w_1, \bar{s}_1) from (29) and (32). Then a bonus is needed to induce a report from the manager. The bonus is given by (28). Both w_1 and \bar{s}_1 will have to adjust as well. Since the manager is indifferent between investing or not, given the optimal bonus in (28), it is up to the firm to decide what projects to accept. The cut-off value \bar{s}_1 is therefore given by:

$$(33) \quad \bar{s}_1 + u_1 + \beta(1 - u_1)(z_2^- - b_1) = 0 .$$

As we mentioned before, the firm neither gains nor loses if no investment is made or if the investment is successful, which explains (33).

Note that (33) is not dependent on w_1 . Given b_1 from (28) and \bar{s}_1 from (33), we find w_1 from the zero-profit condition:

$$(34) \quad w_1 = \int_{\bar{s}_1}^0 [(s_1 + u_1 + \beta(1 - u_1)(z_2^- - b_1))] ds_1 .$$

Finally, $w_2 = \max\{b_1, z_2\}$, because $b_1 > w_1$.

There are cases for which a bonus is required. To see this, note that for a sufficiently risk-averse manager, $b_1 \simeq z_2^0$ in (28). Since w_1 in (34) is always strictly less than z_2^0 as

$$\int_{\underline{s}_1}^0 (s_1 + \mu_1) ds_1 \leq \int_{\hat{s}_1}^0 (s_1 + \mu_1) ds_1 = z_2^0,$$

and $(z_2^- - b_1) < 0$, we have that $w_1 < b_1$ in that case. This implies a need for a bonus.

To sum up our discussion, there are two possibilities. Either $w_1 \leq b_1$, $w_2 = \max\{b_1, z_2\}$, and all projects which give ex post positive expected profits will be accepted; or $w_1 > b_1$, $w_2 = \max\{w_1, z_2\}$ and some projects with ex post negative expected profits will be accepted. In the former case a non-contingent bonus is paid; in the latter case rationing occurs. The former case occurs with a sufficiently risk-averse manager; the latter case occurs with a manager sufficiently close to risk-neutrality.

6. Extensions

In this section we will discuss some simple extensions that are easy to understand based on the preceding analysis.

A. Multi-Period Returns

The intertemporal structure of our model is special in two ways: we dealt with only two periods and we assumed that the returns from an investment accrue in a single period. A full-fledged multi-period model could be written down and analyzed via dynamic programming in much the same way as the two-period model we have already dealt with. The main features of the optimal contract are preserved: downward wage rigidity, the occasional need for bonus payments, and the trade-off between added risk and added information from learning in the design of the optimal decision rule. For details, the interested reader is referred to Ricart i Costa (1984a); see also Harris and Holmstrom (1982).

The possibility that return streams extend over more than one period deserves some further comment, however, in light of a recent paper by Narayanan (1983). As any observer of common business practice could testify, firms often use the payback criterion as a second measure in evaluating investment proposals, despite the fact that net present value (NPV) should be sufficient according to received theory. Narayanan (1983) explains this by noting that managers may prefer investments with a lower NPV, if the returns come in sufficiently fast. The idea is that good early performance raises the manager's opportunity wage and is more valuable than later good performance, because of the downward rigid wage structure.

In light of our general analysis, we can understand this problem better as follows. If markets were complete, so that the manager could

be offered full insurance, then the right criterion to use in judging investment returns would be NPV applied to the total capital stream--that is, financial plus human capital returns. To the extent that real world decisions are based on the NPV of financial returns alone, it is plainly an incomplete measure. But even if the NPV were calculated based on total returns, it would be insufficient in an incomplete market, because the right risk premium would have to be adjusted in response to the human capital risk involved. While the payback criterion is only a proxy for how large the human capital risk is, it can certainly add to the information content of NPV, particularly one based on financial returns. In this way its use can be rationalized, not merely from the manager's perspective as Narayanan's examples would suggest, but also by Pareto optimality. One cannot conclude, however, that early paybacks are always better than late paybacks for any given financial net present value. This depends on whether the value of learning about the manager outweighs the costs of added risk. It appears though that empirically the value of learning often dominates, since early paybacks are considered to be desirable by most managers.

B. Firm Specific Capital

In our model, outside firms were assumed to have identical technological opportunities with the firm in which the manager operates. The market was also assumed to have the same information. Consequently, the value of the manager was identical in all companies and his human capital general rather than firm specific in the language of human capital theory.

With our present understanding of the problem it is easy to see how one would modify the analysis to incorporate differences in firm technologies, which would permit some firm-specific skills. Let z_2^m be the

market value of the manager (i.e. his maximal value in an alternative employment) and let z_2^f be his value in the present firm. Then the extension would specify a distribution $f(y_1, z_2^f, z_2^m | \alpha(\cdot))$ for each decision rule $\alpha(\cdot)$; or alternatively, one could merge the value z_2^f with y_1 into a single return stream in a model where return streams were over more than one period. Either way, the optimal wage contract would remain downward rigid, with z_2^m as the relevant lower bound in the Pareto problems (B) or (C). The cost of learning would only depend on variations in z_2^m , because variations in z_2^f could be fully insured by the firm. Consequently, the decision rule would be biased towards firm-specific returns as in the standard theory of human capital. Note, however, that this would not only occur because firms are unable to appropriate the returns from an increase in general human capital, but also because these returns impose risk on the manager. In this regard our analysis differs slightly from that of standard human capital theory.

An instance of firm-specific capital of a particular kind that is worth mentioning, arises when the firm (and the manager) is better informed about ability than the market at large. One would think this is rather common, given the extra monitoring opportunities that the present employer has relative to competitors. For instance, the market may not be able to observe the circumstances (in our model s_t) under which a particular investment was made and therefore is not able to update beliefs as accurately.

Differential information of this kind will lead to a rather intricate analysis of how the market goes about assessing the manager's value. The reason is that the market has to be wary of "winner's curse": the risk of hiring away only those managers whose value is less than what the market

bids for them. We will not go into the details of the complicated process by which equilibrium wages are determined in this situation since they are discussed in Ricart i Costa (1984b); see also Waldman (1984). We only point out that given the equilibrium market value z_2^m , the optimal wage structure within the firm will remain as stated in Proposition 3. No turnover will occur as long as the market does not have superior technological opportunities.

C. Investment Proposals Verifiable Ex Post

We assumed that when the manager proposes an investment, s_t can be verified by the firm. A slight relaxation of this assumption would permit the firm to verify s_t only ex post, i.e. together with the investment outcome y_t . In that case, the manager would have full discretion to accept or reject projects (at least within a fixed set), though he might be held responsible for bad investment decisions ex post.

We are accustomed to think that ex post verifiability is as good as ex ante verifiability; simply punish the manager if he made the wrong decision ex ante. But in our model (and we think in reality as well) punishments, sufficiently strong to deter all wrong decisions, are not readily available. The manager can always quit and accept the best market offer and that is the maximal threat of the firm. Thus, what is interesting here is that ex post vs. ex ante verifiability may make a difference.

The simplest case is one with a risk neutral manager whose ability is unchanged from period to period. If $z(p)$ is strictly convex then a risk neutral manager wants to invest no matter what the signal value is since the expected return from human capital is strictly positive. This cannot be prevented by punishments. The only alternative is to pay a

bonus, i.e., bribe the manager for telling the truth.⁷ The bribe is subtracted from the first-period wage and in the end no loss in efficiency is incurred, because of risk neutrality.

With risk aversion, losses will generally be incurred, but the structure of the solution remains the same. To see this we must solve the following program:

$$\begin{aligned}
 \text{(D)} \quad & \max_{\delta} \int U(w, \alpha(s_1), s_1) r(s_1) ds_1, \\
 \text{s.t.} \quad & \text{(i)} \quad \int V(w, \alpha(s_1), s_1) r(s_1) ds_1 \geq 0, \\
 & \text{(ii)} \quad w_2(h_1) \geq z_2 = z(p_2(h_1)), \quad \forall h_1, \\
 & \text{(iii)} \quad \alpha(s_1) \in \operatorname{argmax}_{a_1 \in A} E[u(w_2(h_1)) | a_1, s_1], \quad \forall s_1.
 \end{aligned}$$

Note that program (D) is as (C) except for the last constraint. We suppress the action, a_0 , as before. The manager can choose any action (or reveal the s_1 to induce it) since the signal will only ex post be observable by the firm. Hence, the manager has the option to invest even when he should not. He risks being punished if he does so and the project fails, but he is protected against punishments if the project succeeds. To present the solution for (D) we need to define a reservation utility, representing the maximum expected punishment for choosing the wrong action. This reservation utility is given by:

$$\text{(35)} \quad M(s_1) = \max_{a_1 \in A} E[u(z_2) | a_1, s_1].$$

In order to induce $\alpha(s_1)$ given s_1 , the manager must be assured an expected utility of at least $M(s_1)$, since he can always obtain it by his free choice of project and the market guaranteed wages.

We can also define the minimum bribe $\hat{b}(a_1, s_1)$ required to induce action a_1 when the signal is s_1 ; $\hat{b}(a_1, s_1)$ is the solution of:

$$(36) \quad E[u(\max\{\hat{b}(a_1, s_1), z_2\}) | a_1, s_1] = M(s_1) .$$

By using the same kind of arguments that we used in the proof of Proposition 3, one can prove that the solution of (D) differs from the solution in (C) only as follows:

Proposition 5. The optimal solution without indenture when the signal is observed ex post by the owner (and the market) is given by Proposition 3 with (b) replaced by:

$$\begin{aligned} (b') \quad w_2(h_1) &= \max\{w_1, b(s_1), z_2\} , \forall h_1 \text{ such that } a_1 = \alpha(s_1) , \\ w_2(h_1) &= z_2 , \forall h_1 \text{ such that } a_1 \neq \alpha(s_1) , \\ \text{where } b(s_1) &= \hat{b}(\alpha(s_1), s_1) . \end{aligned}$$

The main difference from our previous solution is that we must now specify wages for all possible actions even if they are not part of the decision rule. If the manager deviates from the desired decision rule he can only be penalized if the project fails. This is done optimally by paying him his market value whenever he deviates. In equilibrium it will still be the case that the manager is paid a downward rigid wage, because he will not deviate.

While simple models with symmetric information like this one cannot strictly speaking accommodate firings, paying the market wage may be interpreted as such. Thus, ex post verifiability, suitably embellished, provides a means by which firings might be introduced.

D. Self-Enforced Decision Rules

The last variation we discuss is one where the owner cannot commit himself to a decision rule. That is, we require the decision rule to be self-enforced, so that it is in the owner's interest to follow the decision rule once he has verified the signal presented by the manager.

If we interpret the manager's report as an investment proposal it may be particularly difficult to make the decision rule explicitly contingent on that information. (Note that net present value, for instance, is not a sufficient statistic for a proposal; information about human capital returns and risks also play a role.) The impact of this is that now the owner can also veto an investment after seeing s_1 . As a consequence no investment with negative expected financial returns will ever be approved because recall that the owner gets no share of the human capital returns (ex post). The corresponding program to be solved assuming ex ante verifiability (just for simplicity) is the same as (C) with the additional constraint:

$$(iv) \quad V(w, a(s_1), s_1) \geq \beta z(p_1) - w_1 - \beta E[w_2 | a_0, s_1] .$$

We call this program (E). Note that equation (iv) is independent of β and w_1 , since terms can be cancelled.

Let $d(s_1)$ be the maximum guaranteed wage the owner is willing to pay to go along with a particular action, given s_1 . This value $d(s_1)$ is defined as the solution to

$$(37) \quad V(w, a(s_1), s_1) + w_1 = 0 ,$$

where the wage contract is w_1 and $w_2(h_1) = \max\{z_2, d(s_1)\}$.

Now we can characterize the solution as follows.

Proposition 6. The solution for program (E) satisfies

$$w_2(h_1) = \max\{d(s_1), z_2\} \text{ for } d(s_1) < w_1$$

$$w_2(h_1) = \max\{w_1, b(s_1), z_2\} \text{ for } d(s_1) \geq w_1 .$$

Otherwise the solution stays as before.

The proof of the proposition follows the same arguments as before and is omitted. Note only the differences which are introduced with the additional constraints. Now both parties can block an investment, the manager by concealing information, the owner by his decision authority. The result is that less projects get accepted and, in particular, the hurdle rate on investments will always exceed the cost of (financial) capital (see (37)). Regarding the contract we may note that it still features a guaranteed wage as a function of the signal, but because it has to be self-enforced on both sides the guarantee is signal-dependent and may go below w_1 (when $d(s_1) < w_1$) in order to induce the owner to go along. Thus, we have a case where the wage need not be downward rigid.

7. CONCLUSION

Our main purpose has been to formalize the simple idea that managers may prefer different courses of action than owners or superiors, because of a concern for the consequences actions have on their future opportunities and careers. We feel this source of incongruity in preferences may well be more central to the managerial incentive problem than traditional arguments based on either risk sharing or effort considerations. Our argument allows managers to be industrious and turns instead on legal constraints on involuntary servitude, which limits the degree of income smoothing. The approach imposes discipline by assuming that managers only value their life time income stream. A concern for growth or other dimensions of performance become endogenous to the model.

The central element in our model is learning about managerial characteristics. This makes past performance a rational basis for forecasting future performance. The manager, recognizing that performance is used as a signal about productivity, tries to influence the evaluation process by his choice of actions. The potential source of incongruity in preferences stems from the dissonance between the value of performance as a signal of future productivity and its immediate financial return. This was highlighted in the model formulation by associating each action with two streams of capital, one human and one financial. Without any incentive scheme a manager would only be interested in administering his own human capital stream, while owners would only be interested in the financial capital stream. The socially desirable objective in contrast would look after the sum of the two capital streams and attempt to maximize its expected value.

How one goes about harmonizing the objectives of owners and management in an efficient way depends on the specific information structure. Does the manager have better information about his potential? Does he choose actions privately? What does the market observe relative to the present employer? We chose to consider the simplest possible case of incongruity in which the manager had some private information about investment opportunities. He could withhold this information, but not misrepresent it in case an investment was undertaken. The benefit of this structure was to retain symmetry in beliefs about productivity and thereby make the analysis more tractable. Still it provided a source of incongruity that yielded some interesting insights into the resolution of this kind of incentive problem.

Within our information structure we found a rationale for detailed capital budgeting and rationing. The manager's reward structure took the form of an option on the firm's value, provided that that was linked to his human capital value. The optimal decision rule maximized the total value of capital with a risk premium added for human capital risk. The risk premium implied a higher hurdle rate for investment projects than the cost of capital.

As a general insight we found that there is both a benefit and a cost in learning about managerial ability. The benefit is that the manager can be better placed and used in the future if more is known about him--in effect, his human capital increases (in expectation). The cost is that learning makes his opportunity wage riskier. This cannot be fully internalized because indenture is not permitted. Thus, the desired extent of learning will depend on how these two conflicting interests will be traded off given the manager's risk aversion. We noted earlier that

while the risk aspect is novel to our model, the general conflict is familiar from traditional human capital theoretic analysis, where firms cannot appropriate returns from investment in general skills. The essential distinction is that learning will cause further distortions in decision-making due to insurance concerns.

One of the general implications of our model, which derives directly from the view that managers are concerned about their human capital, is that potential conflicts of interest arise only in areas where human capital is heavily involved. Investment decisions belong to this category, particularly when they expand the manager's scope of authority. This may explain the great emphasis on capital budgeting. On the other hand, routine decision-making, even if financially substantial, may need little attention from owners and superiors. An example of the latter would be pricing decisions in bidding situations. It is somewhat paradoxical that managers may often be given complete authority on pricing decisions of great financial importance at the same time as relatively small investment decisions are scrutinized carefully. Our theory could accommodate these stylized observations by arguing that pricing decisions involve small changes in the manager's human capital, while investments do not.

While the general insights of our analysis - the presence of a conflict of interests, the trade-off between risk and return from learning and the need for control in areas in which human capital is sensitive to decision-making - are certainly robust, the specific conclusions about wage rigidity and the form of decision rules are likely to be rather dependent on the model structure. In particular, our analysis would change considerably with changes in informational assumptions. Some interesting and perhaps more realistic modelling work may lie ahead

in exploring alternative information structures. One possibility is that the manager can take discretionary actions, because his decisions cannot be observed. This would lead to an analysis of contracting that would be rather similar to standard moral hazard analysis. The other major informational change is to assume that the manager is better informed about his ability than his employer (or the market). This case is especially relevant in view of the fact that any forecasting expertise that a manager may possess must necessarily involve private information of a kind that leads to differences in beliefs about ability. In that situation an optimal incentive scheme would not merely try to motivate the manager to submit proposals, but would also attempt to reduce the information gap about ability via self-selection from a menu of possible schemes. A start on the problem has been made in Ricart i Costa (1984c).

It is fairly obvious that our results on the downward rigidity of wages will not obtain when there is asymmetric information either about action choice or ability. Firms would try to make inferences about the unobservables from investment outcomes, which would be reflected in incentive schemes that did not offer maximal insurance subject to the market constraints. Regarding investment choice under asymmetric information, we want to mention the intriguing possibility that managers may want to correlate their decision-making as a way of protecting themselves against excessive risk. A mistake that is common to all managers will be less costly for each one of them than a mistake that is unique in the industry or market. Of course, the value of being uniquely correct in a particular situation is correspondingly more valuable, but the opportunity and the risk do not cancel because of risk aversion. Bandwagon effects, where common (but rational) mistakes get aggravated,

inducing a higher degree of social risk, may be an important element of the competitive system, which has not been recognized to date.⁸

Indeed, the kind of model we have exhibited, with learning as a key ingredient in explaining managerial conduct, leads more generally to new questions about how managers, and consequently firms, compete for capital and customers. It suggests a competitive perspective of the "survival-of-the-fittest" variety. The central issue is whether, or rather to what extent, managerial concern for reputation, and ultimately survival, will drive behavior towards efficient decision-making and the market towards efficient matching of talent and technological opportunities. It is our view that this kind of competition needs to be explored in much more detail for a proper understanding of how well our free market system performs.

FOOTNOTES

1. This view is shared by Kaplan (1984, p. 405), who in his review of managerial accounting research, considers effort-based models highly inadequate for capturing the incentive issues that management faces.
2. Some readers have commented that the situation would be quite different if we allowed w_1 to depend on the outcome y_1 . This is true. We have excluded payment schemes of this form, because they represent contingent contracts. Our point is precisely that contingent contracts will be desirable. As will become clear shortly, one can do better than merely make w_1 contingent on y_1 . Given that fact, we will retain for analytical simplicity the assumption that w_1 has to be paid before y_1 is observed. In a model with many periods, the distinction between paying w_t before or after y_t is observed, becomes quite irrelevant.
3. Huberman (1983) has studied a model that is a special case of this general structure. In his model future ability (or productivity) is a deterministic function of the actions chosen in earlier periods. There is no uncertainty about ability at all.
4. Since (ii) may represent an infinite number of constraints, it is strictly speaking improper to apply Kuhn-Tucker conditions to the problem as we do in the proof. However, one can make the argument rigorous by invoking weak duality. This is done in the proof of Proposition 3, which makes the short-cut here excusable.
5. As we will see in section 6, the internal cost of capital will always be higher than the market cost of capital when applied to financial returns alone, provided that the decision rule cannot be agreed upon in advance.

6. In the risk neutral case there are many optimal contracting solutions. One is the downward rigid wage contract characterized in Proposition 3, while another is to pay the manager his marginal product in each period. However, the optimal decision rule is always given by $\bar{s}_1 = \hat{S}_1$.
7. The use of bribes when sufficient punishments are infeasible is familiar from earlier work on incentives; see Becker and Stigler (1974).
8. In a similar vein, one might expect that the degree to which responsibility for decision-making is shared among several managers in a firm, has an impact on the willingness to take risks. Kenneth Arrow alerted us to this consideration. He also pointed out that Japanese firms might have been able to take more risks (with higher expected payoffs) than U.S. firms, because of the way in which responsibility is shared in their managerial culture, and that this may partly explain their rapid growth. We may add that it is also an interesting theoretical question to study how one updates beliefs about ability of many managers from common decisions.

REFERENCES

- Becker, G. and G. Stigler (1974), "Law Enforcement and Compensation of Enforcers," Journal of Legal Studies, Vol. 3, No. 1.
- Blackwell, D. (1953), "Equivalent Comparison of Experiments," Annals of Mathematical Statistics, Vol. 24.
- DeGroot, M. (1970), *Optimal Statistical Decisions*, New York: McGraw Hill.
- Fama, E. (1980), "Agency Problems and the Theory of the Firm," Journal of Political Economy, Vol. 88, No. 2.
- Grossman, S. and O. Hart (1983), "An Analysis of the Principal-Agent Problem," Econometrica, January.
- Harris, M. and B. Holmstrom (1982), "A Theory of Wage Dynamics," Review of Economic Studies, July.
- Holmstrom, B. and L. Weiss (1982), "Managerial Incentives, Investment and Aggregate Implications: Scale Effects," TR No. , IMSSS, Stanford University. Forthcoming in Review of Economic Studies.
- Huberman, G. (1983), "The disciplinary Market: A Multi-Period Agency Problem," WP No. 110, University of Chicago.
- Kaplan, R. (1984), "The Evolution of Management Accounting," The Accounting Review, Vol. LIX No. 3, July.
- Lambert, R. (1984), "Executive Effort and Selection of Risky Projects," mimeo, Northwestern University.
- Narayanan, M.D. (1983), "Observability and the Payback Criterion," mimeo, University of Florida.
- Ricart i Costa, J. (1984a), "On Managerial Contracts," Ph.D. dissertation, Northwestern University.
- Ricart i Costa, J. (1984b), "Managerial Task Assignment and Promotions," CMSEMS Discussion Paper, Northwestern University.

- Ricard i Costa, J. (1984c), "Managerial Talent, Investment Decisions and Asymmetric Information," mimeo, Northwestern University.
- Ross, S. (1973), "The Economic Theory of Agency: The Principal's Problem," American Economic Review, Vol. 63.
- Waldman, M. (1984), "Job Assignments, Signalling and Efficiency," Rand Journal of Economics, Spring.
- Wilson, R. (1969), "The Structure of Incentives for Decentralization Under Uncertainty," La Decision, No. 171.