



Managerial Incentive Problems: A Dynamic Perspective

Bengt Holmstrom

The Review of Economic Studies, Vol. 66, No. 1, Special Issue: Contracts. (Jan., 1999), pp. 169-182.

Stable URL:

<http://links.jstor.org/sici?sici=0034-6527%28199901%2966%3A1%3C169%3AMIPADP%3E2.0.CO%3B2-9>

The Review of Economic Studies is currently published by The Review of Economic Studies Ltd..

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

Please contact the publisher regarding any further use of this work. Publisher contact information may be obtained at <http://www.jstor.org/journals/resl.html>.

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

JSTOR is an independent not-for-profit organization dedicated to creating and preserving a digital archive of scholarly journals. For more information regarding JSTOR, please contact support@jstor.org.

Managerial Incentive Problems: A Dynamic Perspective

BENGT HOLMSTRÖM
MIT

First version received August 1998; final version accepted September 1998 (Eds.)

The paper studies how a person's concern for a future career may influence his or her incentives to put in effort or make decisions on the job. In the model, the person's productive abilities are revealed over time through observations of performance. There are no explicit output-contingent contracts, but since the wage in each period is based on expected output and expected output depends on assessed ability, an "implicit contract" links today's performance to future wages. An incentive problem arises from the person's ability and desire to influence the learning process, and therefore the wage process, by taking unobserved actions that affect today's performance. The fundamental incongruity in preferences is between the individual's concern for human capital returns and the firm's concern for financial returns. The two need be only weakly related. It is shown that career motives can be beneficial as well as detrimental, depending on how well the two kinds of capital returns are aligned.

1. INTRODUCTION

It is well understood by now that informational externalities may place special demands on the organization of economic exchange. Simple price-mediated markets will frequently fail in the presence of asymmetric information. In that case more elaborate contractual arrangements have to be used as substitutes for the price system. Lately, considerable effort has been devoted to the analysis of contracting under incomplete information with the objective to understand the range of economic institutions that emerge in response to the failure of the price system.

The analysis of moral hazard has played a prominent role in this development.¹ Moral hazard problems arise when, for some reason or another, transacting parties cannot contract contingent on the delivery of the good. For instance, in buying labour services it may be that the amount of labour supplied is not directly observable, precluding a simple exchange of wage for labour. As a partial remedy to this problem, an imperfect, mutually observed signal about the supply of labour can be used as a proxy in the contract. Frequently, output is taken as such a proxy. The drawback is that output is often influenced by other factors than labour input, which induce undesirable risk into the contract. One is therefore faced with a tradeoff between allocating risk associated with incomplete observability and providing incentives for a proper supply of labour. Gaining insight into this tradeoff is important not only for understanding contracting in the small (*e.g.* managerial incentive schemes), but also because it is closely related to the fundamental tension between equity and efficiency in the society as a whole.

While our understanding of moral hazard has advanced a lot in past years, it is clear that much work remains. An important question that has received little attention until

1. For some recent work on moral hazard the reader is referred to Mirrlees (1976), Harris and Raviv (1979), Holmstrom (1979), Shavell (1979), and Grossman and Hart (1983).

very recently concerns the effect time has on incentives. Intuitively, time should have a beneficial impact on policing moral hazard, because it permits a longer series of observations and thereby more accurate inferences about unobservable behaviour. This intuition has been made precise in work by Radner (1981) and Rubinstein (1981), who show that explicit long-term contracts can be written, which reduce incentive costs to zero when there is no discounting. Fama (1980) reaches this same conclusion using a conceptually different approach. He argues that market forces alone will frequently remove moral hazard problems, because managers will be concerned about their reputations in the labour market. Thus, there will be no need to resolve incentive problems using explicit contracts, since markets already provide efficient implicit incentive contracts.

The purpose of this paper is to investigate in some more detail Fama's rather provocative but interesting idea that career concerns induce efficient managerial behaviour. Since Fama does not provide an explicit model of moral hazard, I start by formalizing his intuition. In the first part of the paper I present a model, based on that in Harris and Holmstrom (1982), which permits an explicit analysis of the manager's decision to supply labour. Under some narrow assumptions I show that Fama's conclusion is correct. In general, however, it is not. Risk-aversion and discounting place obvious limitations on the market's ability to police incentives adequately. More interesting therefore is my analysis of transient learning effects and non-linearities in technology, which both lead to inefficiencies even when there is no discounting and the manager is risk-neutral.

In the second part of the paper I consider the implications of reputation on managerial risk-taking. I argue that so far there has been no good explanation for why there should be an incentive problem with risk-taking in the first place, although this is clearly perceived to be an important issue in the real world. Using some simple examples I show then how a basic incongruity in risk preferences between the manager and the firm arises from the manager's career concerns. Although I do not analyse how the problem should be resolved optimally, my analysis opens a new and promising direction for research on this question. Since managerial risk-taking problems appear specifically in a dynamic setting, this shows that, contrary to common intuition, time need not always be a blessing when it comes to incentive issues. It can create problems as well.

2. WORK INCENTIVES

2.1. *The basic model*

I will start by presenting the simplest model of reputation formation, leaving embellishments for later sections. Consider the following scenario of a manager operating in a competitive labour market. The manager is endowed with labour, which he sells in the market in exchange for consumption. No contingent contracts can be made, so we may envision that the manager is paid for his services in advance. In a one-period world he would have no incentive to work. The same is true in a multi-period world if there were no uncertainty about the characteristics of the agent. In order that there be some returns to the manager for good performance, it must be that present performance acts as information about future performance. Logically, this requires uncertainty about some characteristic of the manager. It is natural to take this characteristic to be talent, though many alternatives would do as well.

Let η be a qualified measure of the manager's talent and assume initially that it is fixed and incompletely known to the manager and the market. The market and the manager share prior beliefs about η ; specifically, assume that this prior is normally distributed

with mean m_1 and precision (equal to the inverse of the variance) h_1 . Over time, learning about η will occur through the observation of the manager's output. In period t , this output is given by the technology

$$y_t = \eta + a_t + \varepsilon_t, \quad t = 1, 2, \dots \quad (1)$$

where $a_t \in [0, \infty]$ is the manager's labour input and ε_t is a stochastic noise term. To be able to make inferences about η from (1) requires a distribution on ε_t ; I take ε_t 's to be independent and normally distributed with zero mean and precision h_ε .

The manager is assumed to be risk neutral with preferences given by an atemporal, separable utility function

$$U(c, a) = \sum_{t=1}^{\infty} \beta^{t-1} [c_t - g(a_t)]. \quad (2)$$

Disutility of labour is measured by $g(\cdot)$, which is increasing and convex. It is assumed that $U(\cdot, \cdot)$ is publicly known.

In order to decide how much labour to supply, the manager has to calculate the impact of present output on future wages. On the other hand, the dependence of future wages on past output is a function of the manager's decision rule. Consequently, the decision rule and the wage functions are determined simultaneously in equilibrium. In general, this interaction may be quite complicated, but for the simple technology considered here, an explicit solution is easily obtained.

Let $y^t = (y_1, \dots, y_t)$ be the history of outputs up to time t . This information is assumed known to the market and used as a basis for wage payments. Let $w_t(y^{t-1})$ be the wage in period t and $a_t(y^{t-1})$ be the manager's labour supply in the same period, both functions of the history. A competitive market, neutral to risk, will set

$$w_t(y^{t-1}) = E[y_t | y^{t-1}] = E[\eta | y^{t-1}] + a_t(y^{t-1}). \quad (3)$$

This determines the wage in period t given that the manager's decision rule is known. On the other hand, given (3), the manager's decision rule solves²

$$\max_{\{a_t(\cdot)\}} \sum_{t=1}^{\infty} \beta^{t-1} [E w_t(y^{t-1}) - E g(a_t(y^{t-1}))]. \quad (4)$$

The solution to (4) together with (3) determines equilibrium.

Notice that even though the market is not able to observe the manager's actions directly, it is able to infer them by solving (4). Therefore, observing y_t will in equilibrium be equivalent to observing the sequence

$$z_t \equiv \eta + \varepsilon_t = y_t - a_t^*(y^{t-1}), \quad (5)$$

where $a_t^*(y^{t-1})$ represents the equilibrium decision rule. Through the observation of the sequence $\{z_t\}$ the market learns about η . In fact, this learning process is well-known given the normality and independence assumptions. The posterior distributions of η will stay normal with means and precisions given by

$$m_{t+1} = \frac{h_t m_t + h_\varepsilon z_t}{h_t + h_\varepsilon} = \frac{h_1 m_1 + h_\varepsilon \sum_{s=1}^t z_s}{h_1 + t h_\varepsilon}, \quad (6)$$

$$h_{t+1} = h_t + h_\varepsilon = h_1 + t h_\varepsilon. \quad (7)$$

Observe that the mean process $\{m_t\}$ is a random walk with incremental variance that declines deterministically to zero. In the limit η will become fully known.

2. Since the manager is risk neutral and no contracts are considered, borrowing and saving can be ignored.

Using (6), (3) can be written as

$$w_t(y^{t-1}) = m_t(z^{t-1}) + a_t^*(y^{t-1}), \quad (8)$$

where $z^t = (z_1, \dots, z_t)$. Taking expectations in (8) (with actions fixed and non-contingent) yields

$$Ew_t(y^{t-1}) = \frac{h_1 m_1}{h_t} + \frac{h_\varepsilon}{h_t} \sum_{s=1}^{t-1} (m_1 + a_s - Ea_s^*(y^{s-1})) + Ea_t^*(y^{t-1}). \quad (9)$$

From (9) follows that for a non-stochastic equilibrium path of labour supply the marginal return to a_1 in period t will be $\alpha_t = h_\varepsilon/h_t$ independently of the past. The solution to (4) is then given by the first order conditions

$$\gamma_t \equiv \sum_{s=t}^{\infty} \beta^{s-t} \alpha_s = g'(a_t^*). \quad (10)$$

Obviously, γ_t is a declining sequence, and since the sum in (10) converges (because $\alpha_s \rightarrow 0$), $\gamma_t \rightarrow 0$. Consequently, the equilibrium sequence of labour inputs is declining and goes asymptotically towards zero as $t \rightarrow \infty$.

The interpretation of this result is straightforward. As long as ability is unknown there are returns to supplying labour, because output will influence perceptions about ability. Indeed, labour is a substitute for ability. By increasing its supply, the manager can potentially bias the process of inference in his favour. Of course, in equilibrium this will not happen, because the market will know what effort level to expect and adjust the output measure accordingly (see (5)). In other words, the manager cannot fool the market. Yet, he is trapped in supplying the equilibrium level that is expected of him, because, as in a rat race, a lower supply of labour will bias the evaluation procedure against him.

Furthermore, the returns to labour supply are bigger the more there is uncertainty about ability, as can be seen from (10). Early in the process, when there is less information, the market puts more weight on the most recent output observation when revising its beliefs about η . Eventually, η is revealed almost completely and new observations will have very little impact on beliefs. In the limit, therefore, there are no returns to trying to influence output and labour supply goes to zero.

2.2. The stationary case

The results above, of course, bear little relationship to efficient labour supply. Efficiency would require that $a_t = \bar{a}$ for all t , where \bar{a} is defined through

$$g'(\bar{a}) = 1. \quad (11)$$

The problem is that reputation formation is valuable only temporarily. To get a permanent reputation effect one must prevent η from becoming fully known. This is accomplished by assuming that ability is not fixed, but fluctuates over time. For instance, let ability progress according to the following process

$$\eta_{t+1} = \eta_t + \delta_t, \quad (12)$$

where δ_t are independent and normally distributed with mean zero and precision h_δ .

The learning process will change in a slight, but important way. As before,

$$m_{t+1} = \mu_t m_t + (1 - \mu_t) z_t, \quad (13)$$

where

$$\mu_t = \frac{h_t}{h_t + h_\varepsilon}. \tag{14}$$

However, h_{t+1} will be different. Let \hat{h}_t be the precision on η_{t+1} before observing y_{t+1} . We have, as before

$$\hat{h}_t = h_t + h_\varepsilon. \tag{15}$$

From (12) follows (by independence)

$$\frac{1}{h_{t+1}} = \frac{1}{\hat{h}_t} + \frac{1}{h_\delta},$$

which with (15) gives

$$h_{t+1} = \frac{(h_t + h_\varepsilon)h_\delta}{h_t + h_\varepsilon + h_\delta}. \tag{16}$$

Thus, h_t will still progress deterministically, but will not go to infinity with t (as before), because the δ -shocks keep adding uncertainty. Instead, h_t will approach a stationary state h^* in which learning through output observations is just enough to offset the periodic increase in uncertainty from the δ -shocks.

It is somewhat easier to express the stationary state in terms of μ_t s, which, of course, are in a one-to-one correspondence with h_t s through (14). Simple algebra gives the following recursion for the μ_t s

$$\mu_{t+1} = \frac{1}{2 + r - \mu_t}, \tag{17}$$

where

$$r \equiv \frac{h_\varepsilon}{h_\delta} = \frac{\sigma_\delta^2}{\sigma_\varepsilon^2}. \tag{18}$$

Stationarity requires $\mu_{t+1} = \mu_t = \mu^*$. Solving for μ^* from (17) yields

$$\mu^* = 1 + \frac{1}{2}r - \sqrt{\frac{1}{4}r^2 + r}. \tag{19}$$

Notice that $0 < \mu^* < 1$. If $r = 0$, so that ε has high variance relative to δ , then $\mu^* = 1$. In that case, the updating of m_t occurs slowly (see (13)). The reverse holds true if $r = 1$.

In terms of μ^* , the stationary level of the precision, h^* is (using (14) and (19))

$$h^* = \frac{h_\varepsilon \mu^*}{1 - \mu^*}. \tag{20}$$

This settles the stationary learning process. Next, consider the ramifications on incentives. Following the earlier reasoning, the optimal labour supply, a_t^* , is given by

$$\gamma_t \equiv (1 - \mu_t) \sum_{s=t+1}^{\infty} \beta^{s-t} [\prod_{i=t+1}^s \mu_i] = g'(a_t^*). \tag{21}$$

In the stationary state $\mu_s = \mu^*$. Substituting this into (21) implies that the stationary labour supply, a^* , satisfies

$$\frac{\beta(1 - \mu^*)}{1 - \mu^* \beta} = g'(a^*). \tag{22}$$

Notice that the left-hand side is between 0 and 1, so $a^* < \bar{a}$, the efficient level of labour supply. From (22) we also reach Fama's major conclusion: if $\beta = 1$, then $g'(a^*) = 1$, which means that the stationary state is efficient. It is rather striking that this occurs as soon as we add any amount of noise in the η -process. With $\beta = 1$, efficient labour supply is independent of the degree of this noise even though the noiseless case leads to no labour supply as was shown in the previous section. This discontinuity disappears as soon as $\beta < 1$. Then a small variance of δ_t relative to ε_t implies a μ^* close to 1 and a stationary labour supply close to 0.

The general implications of (22) can be summarized by the following:

Proposition 1. *The stationary level of labour supply a^* is never greater than the efficient level of labour supply \bar{a} . It is equal to \bar{a} if $\beta = 1$ and $\sigma_\varepsilon^2, \sigma_\delta^2 > 0$. It is closer to \bar{a} the bigger is β , the higher is σ_δ^2 and the lower is σ_ε^2 .*

In words, the comparative statics results tell us that reputation will work more effectively if the ability process is more stochastic or if the observations on outputs are more accurate. Both features will speed up learning and move forward the returns from labour investments, reducing the negative effects of discounting.

2.3. Transient effects

Proposition 1 tells us how incentives depend on the discount rate and the degree of noise in output and ability. Next I will consider incentives before a stationary state is reached. This involves exploring the convergence to the stationary state, which in itself is important if the results in the previous section are to be taken seriously.

Again it is easiest to work with the μ_t s. The dynamics of μ_t is given by (17). From (17) follows that μ_{t+1} is an increasing function of μ_t and from (19) follows that there is exactly one stationary state within the interval (0, 1). These facts are recorded in Figure 1.

From Figure 1 it is seen that if one starts with a value $\mu_1 < \mu^*$, μ_t will converge from below to μ^* and if one starts with $\mu_1 > \mu^*$, μ_t will converge from above to μ^* . The system is therefore stable. From the definition of μ_t (equation (14)) it follows that the stability can be cast in terms of h_t as well. If $h_1 < h^*$, $h_t \uparrow h^*$, and if $h_1 > h^*$, $h_t \downarrow h^*$.

The dynamics of a_t^* will follow by studying (21). Let me show first that γ_1 is a decreasing function of μ_1 . The coefficient γ_1 is the sum of the terms $\beta(1 - \mu_1)$, $\beta^2(1 - \mu_1)\mu_2$, $\beta^3(1 - \mu_1)\mu_2\mu_3$, etc. If each term is decreasing in μ_1 , the same is obviously true for γ_1 . This step is proved by induction. Suppose $b_s(\mu_1) \equiv (1 - \mu_1)\mu_2\mu_3 \dots \mu_s$ is decreasing in μ_1 and consider $b_{s+1}(\mu_1)$. One can write

$$b_{s+1}(\mu_1) = \frac{1 - \mu_1}{1 - \mu_2} \mu_2 b_s(\mu_2) = \frac{1 - \mu_1}{1 + r - \mu_1} b_s(\mu_2),$$

by using (17), (18) and the definition of μ_t . By the inductive hypothesis, $b_s(\cdot)$ is decreasing. Since μ_2 is increasing in μ_1 by (17), it follows that $b_{s+1}(\mu_1)$ is decreasing in μ_1 . Consequently, γ_1 is decreasing as a function of μ_1 .

It follows, by the definitions of γ_t and μ_t , that, $\{\gamma_t\}$ is a decreasing (increasing) sequence if $\{\mu_t\}$ is an increasing (decreasing) sequence. Recalling then that $\mu_t \uparrow (\downarrow) \mu^*$ if $h_1 < (>) h^*$, I have established the following stability result.

Proposition 2. *The sequence of optimal labour supply $\{a_t^*\}$ will converge monotonically to the stationary state a^* . If the initial precision of information about ability, h_1 , is*




FIGURE 1

less than the stationary precision level h^* , the convergence of a_t^* is from above. Conversely, $h_1 > h^*$ implies $a_t^* \uparrow a^*$.

The convergence result is illustrated in Figure 2. With the interpretation of η as ability, it seems clear that $h_1 < h^*$ is the common case. Normally, we expect that the precision of information about ability increases as time goes on. The picture shows that in that case young people will overinvest in labour supply because the returns from building a reputation are highest when the market information is most diffuse.

This seems to accord nicely with casual empiricism (including introspection). There is some scientific evidence as well. Medoff and Abraham (1981) conducted a study where they measured the productivity of different age groups in various job categories. Though the evidence was not overwhelmingly strong, the study pointed towards the fact that young people are more productive. If one believes that equally able people are, roughly at least, placed in the same jobs, their findings imply that young people supply more labour.

To the extent that convergence to a stationary state is slow, which again will be the case if output is noisy relative to shocks in ability, the analysis above shows that there may be a substantial transient inefficiency even when there is no discounting.

2.4. Scale economies

Next I turn to changes in technology. It is clear that the linearity in (1) and (12) is essential for efficiency. To show this in general seems both messy and uninteresting so I will only discuss the matter via some illuminating examples. To reduce complexity, assume that




FIGURE 2

there is no noise in the observation of output, *i.e.* let $\varepsilon_t \equiv 0$. Nothing pathological is introduced in this way. It merely implies that all returns from labour supply accrue in the next period since $\mu^* = 0$. Notice that with the earlier used linear technology, efficiency obtains in all periods in this special case.

Now, suppose output is given by

$$y_t = f(\eta_t) + a_t. \tag{23}$$

I leave a_t outside $f(\cdot)$, because then efficiency simply requires that $a_t = \bar{a}$ in all periods. Instead of interpreting $f(\cdot)$ as a production function, one can view (23) as a way of making the learning process non-linear (and output non-symmetrically distributed). Because of this, there is no *a priori* reason to assume $f(\cdot)$ is concave.

Let η_0 be the ability level inferred from the last observation. The manager's wage today is $w_1 = Ef(\eta_1) + \bar{a}$ under the assumption that $a_1 = \bar{a}$. The question is, will he choose $a_1 = \bar{a}$? To answer this, the returns from a_t have to be calculated. They will come from w_2 only, since $\mu^* = 0$. For w_2 we have the expression

$$w_2 = Ef(\eta_2) + a_2 = Ef(\hat{\eta}_1 + \delta_1) + a_2 = E[f(f^{-1}(f(\eta_0 + \delta_0) + a_1 - \bar{a}) + \delta_1)] + a_2. \tag{24}$$

Here $\hat{\eta}_1$ is the ability level that the market infers from y_1 , by computing $f^{-1}(y_1 - \bar{a})$. If $a_1 \neq \bar{a}$, then $y_1 = f(\eta_0 + \delta_0) + a_1 - \bar{a}$ and $\hat{\eta}_1 \neq \eta_1$. The expectation in (24) is taken over δ_0 and δ_1 under the assumption that the manager knows no more about his ability than the market when choosing a_1 . The marginal benefit from a_1 at $a_1 = \bar{a}$ is then

$$E[f'(\eta_0 + \delta_0 + \delta_1)(f^{-1})'(f(\eta_0 + \delta_0))] = E\left[\frac{f'(\eta_0 + \delta_0 + \delta_1)}{f'(\eta_0 + \delta_0)}\right]. \tag{25}$$

Obviously, this expression will generally differ from 1 as efficiency would require. For instance, if $f'(\cdot)$ is convex (*i.e.* $f'''(\cdot) > 0$), then it is strictly greater than 1 (by Jensen's inequality). Thus (strong) convexity points to oversupply of labour. The reverse holds for (strong) concavity.

Another, perhaps more natural, example of non-linearity is the following

$$y_t = a_t \eta_t. \quad (26)$$

If $g(a_t) = \frac{1}{2}a_t^2$, then efficiency requires $a_t = \eta_t$. With this decision rule the marginal returns to today's effort can be easily calculated to be $\eta_t^2 + \sigma_\delta^2$. The marginal return from output is, however, η_t according to (26). Thus there will be overinvestment in labour when ability is perceived to be high and underinvestment when ability is perceived to be low. Labour input will vary more than efficiency would dictate.

A third class of cases with inefficient outcomes arises when job matching is introduced. Suppose managers are matched to jobs according to perceived ability. If output is linear in ability in each task, then optimal matching of persons to tasks will yield overall returns to ability which are convex (see Rosen (1982) for more on this point). This convexity will result in proportionately larger returns to labour from reputation than are the actual returns from production. Since this case is formally very similar to the previous example I omit a more detailed argument. The idea can perhaps be most easily grasped if we think of the returns from labour in a pure signalling model of schooling. In that case there is no productive value from students working hard for better grades. Yet, students do work hard, because of reputation effects, even though it is entirely wasteful from a social point of view.

The general point illustrated by the examples above, is, of course, that the returns from signalling need not be closely aligned with the returns to present output, unless the technology is linear.

2.5. Discussion

Fama has argued that in a dynamic perspective reputation effects will frequently be sufficient to police moral hazard problems without recourse to explicit output based contracts. The exercises above were conducted to explore the generality of such a statement. Although anything but general themselves, they suggest, to me at least, that quite restrictive conditions have to be imposed to reach efficiency.

The mere observation that a number of factors reduce the efficiency of market incentives is of limited interest. After all, there is plenty of empirical evidence that explicit incentive schemes as well as implicit wage structures are important in the real world. Furthermore, the most obvious reason for a need to contract has so far gone unmentioned: risk-aversion. The market incentives discussed above do not protect the manager at all against risk and as such they are clearly suboptimal.

Thus, there is little reason to doubt that contracts will play an important part in a fuller analysis of dynamic moral hazard. The value of the present analysis rests with the faith that even when contracts are included, some of the qualitative conclusions reached here will remain true; in particular, that the need for incentives which increase labour supply, is small in the early stages of a manager's career and in the situations where returns to ability are convex.

3. INCENTIVES FOR RISK TAKING

Providing work incentives is only part of the managerial incentive problem. To secure proper behaviour in the choice of investments is equally important. Firms frequently express a concern over the way their management takes risks. Some think their managers take too much risk; but perhaps more commonly managers, particularly the younger ones, are seen as overly risk-averse.

Wilson (1968) and Ross (1973) have addressed the problem of designing reward schemes which induce correct incentives for risk taking. I would argue, however, that their models do not capture the essential aspect of the problem. The reason is that in their models an incentive problem arises only as a consequence of attempts to utilize the manager's risk absorption capacity. This may be relevant in small, closely held firms. But in a firm of even modest size or in a publicly held corporation, gains from having the manager carry some risk are certainly negligible. The apparent solution (in their models at least) is to offer the manager a constant wage and ask him to act in the firm's best interest. This will yield an outcome that for all practical purposes is efficient.³

Thus in the Wilson–Ross model, there really is no incentive problem in the first place. So what can account for the common concern?

I think a major reason for incongruity in risk preferences stems from the manager's career concerns.⁴ A large part of managerial talent relates to projecting investment returns and choosing the good prospects. If talent is not fully known, investment decisions become tests that provide information about talent. Perceptions about talent, in turn, determine the manager's future opportunity wage and this is what makes investments risky from the manager's perspective even if income is not explicitly tied to profits. The solution suggested for the Wilson–Ross model (a constant income) is not feasible, because a manager whose ability is perceived high will be bid away (see Harris and Holmstrom (1982)). I will elaborate on this idea in two examples below.

3.1. *An incongruity in risk preferences*

Consider a manager who is in charge of choosing investment projects for a risk-neutral firm. He may be either talented or not. Talent is associated with the likelihood that investments are successful. Presently, the probability that he is talented is assessed to be η by the firm as well as the manager.

Investments can either fail or succeed. Let y_- be the payoff if a project fails and y_+ if it succeeds. The likelihood that a project succeeds is l_T if the manager is talented and l_N if he is not. Obviously, $l_T > l_N$. The overall probability of success is then

$$p = l_T \eta + l_N (1 - \eta). \quad (27)$$

In this set-up, investment projects are characterized by the vector $I = (y_+, y_-, l_T, l_N)$ (or equivalently by the vector (y_+, y_-, l_T, p)). The pool of potential projects is a collection of such I s. The manager's expertise lies in observing this pool while others do not.

From the pool the manager will choose at most one project and propose it for investment.⁵ Such a proposal involves presenting the information I in a verifiable way to

3. A similar point is made in Ross (1977).

4. An alternative reason is that work incentives will require that the manager is paid as a function of firm output and this in turn induces a difference in preferences for risk. The model by Grossman and Hart (1983) can formally account for this possibility, but they do not explore the consequences of such incongruity.

5. Assuming that at most one project is selected is without loss of generality if y_- and y_+ are the same for all projects (and in a more general model with arbitrary investment outcomes).

his superiors who will make the final decision. Thus, potential incentive problems are not associated with misrepresenting information about a proposed project, but with the possibility that the proposed project is not the best available alternative from the firm's perspective.

I now show that hiding information will indeed be a problem. Let η_+ (η_-) be the probability that the manager is talented given that the investment succeeds (fails). By Bayes' rule

$$\eta_+ = \frac{l_T \eta}{p}, \quad \eta_- = \frac{(1-l_T)\eta}{1-p}. \quad (28)$$

The manager's opportunity wage will be a function of the updated assessments above. What the exact relationship is depends on the exact specification of the investment pool. Shortly, I will examine a case where the opportunity wage is linear in η , so let me proceed with this assumption. Without loss of generality, (28) then coincides with the payoff for the manager. The expected value of the manager's risk is therefore

$$p \frac{l_T \eta}{p} + (1-p) \frac{(1-l_T)\eta}{1-p} = \eta. \quad (29)$$

The fact that the expected value coincides with the prior probability of talent is actually more general. Since the manager's lottery forms a martingale with respect to beliefs, it will be true whenever payoffs are linear functions of the posteriors.

If the manager is risk-neutral he is indifferent between all projects. He can therefore be expected to propose the project which the firm prefers most. For a risk-averse manager things are different. The expected return from undertaking an investment is no higher than abstaining from investments altogether. Since investing carries risk it is then clear that the manager would not like to invest at all. He will have an incentive to claim that no worthwhile investment opportunity was present in the pool of potential investments. Under the informational assumptions made, such a claim cannot be invalidated.

The analysis above shows that career concerns induce a genuine incongruity in risk preferences between the firm and the manager. To emphasize this point, notice that the risk facing the manager is quite different from the risk that is of concern to the firm. A key variable for the manager is the likelihood of success l_T . The manager dislikes investments, which will reveal accurately whether he is a talented manager or not, since these investments make his income most risky. He prefers investments which leave him protected by exogenous reasons for investment failure. The firm, however, has no interest in l_T given p .⁶ Instead, it is mainly concerned with the actual payoffs (y_- , y_+) of the project and these again are irrelevant for the manager.

Evidently, the manager has to be given some stake in the real outcome if preferences are to be brought closer together. Giving him a share of the firm may not be the best strategy, however, since it carries both downside and upside risk. A stock option could be a more valuable incentive, since it removes the downside risk. This would be an interesting conclusion in view of the prominent role options have played in managerial incentive plans, but verification of its validity has to await a more careful analysis.

6. This may not be generally true if the firm finds value in learning the manager's ability for purposes of placement.

3.2. A "lemons" problem

My final example, an elaboration on the previous one, illustrates that if the manager cannot communicate investment risks in a verifiable way incentive problems get even more severe.

Let the investment pool consist of only one project. The project characteristics are $I = (-1, +1, s, \frac{1}{2})$. The manager's only private information is the likelihood $l_T = s$ that the project succeeds if he is talented. One can view s as a signal about the likelihood of success, which is relevant only if the manager has talent. From the firm's point of view, the manager should invest if $s \geq \frac{1}{2}$, since the expected value, conditional on s , is $\eta(2s - 1)$.

The firm does not know what s is, but assesses a uniform distribution to it. *Ex ante*, the value of the manager (*i.e.* his information) is then easily seen to be $\frac{1}{4}\eta$. If the manager only lives for two periods, then no incentive problems arise in the second period and his opportunity wage will be $\frac{1}{4}\eta'$, where η' is the revised talent assessment.

The posterior beliefs about talent will depend on the manager's decision rule. Suppose beliefs are updated under the assumption that the manager invests if $z \equiv 2s - 1 > 0$. The posteriors on his talent will then be

$$\eta_+ = \frac{3\eta}{2 + \eta}, \quad \eta_- = \frac{\eta}{2 - \eta}. \quad (30)$$

Of course, if no investment is made the posterior is $\eta' = \eta$. Will a risk neutral manager actually use $z \geq 0$ as his investment criterion? Simple algebra shows that he will invest if

$$\frac{\frac{1}{2}(1 - \eta z)}{2 - \eta} + \frac{\frac{3}{2}(1 + \eta z)}{2 + \eta} > 1. \quad (31)$$

The left-hand side is increasing in z and the value at $z = 0$ is less than 1 (for $\eta \neq 0, 1$). Consequently, the manager will use as his cutoff rate some $\bar{z} > 0$. Thus, if a risk neutral manager is rewarded according to expected marginal product, computed based on the rule to invest if $z \geq 0$, *he will not conform to this rule*. He will take less risk, because of a concern for the negative talent evaluation that follows upon failure. More specifically, he realizes that the firm will update beliefs about talent conditional on the general knowledge that $\{z > 0\}$ obtained (since an investment was made), which puts him in an unfavourable position if z is actually close to 0.

It is natural to ask whether there is another cutoff value \bar{z} such that the manager wants to invest exactly when $z > \bar{z}$ given that he is paid his expected product in the second period and given that this expected product is calculated based on the updating rules that apply when $z > \bar{z}$ is the investment rule of the manager.

The updating rules for talent, conditional on investment when $z > \bar{z}$, are

$$\eta_+ = \frac{\eta(3 + \bar{z})}{2 + \eta + \eta\bar{z}}, \quad \eta_- = \frac{\eta(1 - \bar{z})}{2 - \eta - \eta\bar{z}}. \quad (32)$$

On the other hand the manager invests whenever z is such that

$$\eta_+(1 + \eta z) + \eta_-(1 - \eta z) \geq 2\eta. \quad (33)$$

Combining (32) and (33) gives the equilibrium condition for \bar{z}

$$\frac{(3 + \bar{z})(1 + \eta\bar{z})}{2 + \eta + \eta\bar{z}} + \frac{(1 - \bar{z})(1 - \eta\bar{z})}{2 - \eta - \eta\bar{z}} = 2. \quad (34)$$

Equation (34) can be shown to have no other solution in $(-1, +1)$ than $\bar{z} = 1$. As in Akerlof's (1970) "lemons" model, the only equilibrium is the degenerate one where no investments are made. Thus, if the manager cannot have his investment information validated it makes him more conservative. Even a risk neutral manager acts as if he is risk averse in this example.

3.3. Discussion

There are a number of reasons why the incentive problems described above may not be as severe as stated. For the same reasons as in Section 2.4, it could be that payoffs are convex in talent, reducing the aversion to risk-taking. The manager may also know more about his talent than the firm. An undervalued manager would then be willing to take risk in order to prove himself implying that risk-taking in itself would be a signal of talent. The same would be true if talented managers would receive higher signals on average than less talented managers.

Indeed, possibilities like these suggest a rich agenda for future research and indicate that modelling risk-taking from a dynamic perspective is a fruitful approach. I note in passing that such models may also help us understand the puzzle why investment procedures in firms are so detailed and centralized. As the latter example showed, it may have as much to do with securing a proper evaluation of managerial talent as it has to do with controlling what projects get selected.

4. CONCLUDING REMARKS

This paper has explored some ramifications of the thesis that managerial incentive problems are closely tied to learning about managerial ability. It implies a dynamic perspective on incentive issues. The paper has raised rather than answered questions, but the awareness of issues is a first and important step towards resolving problems.

Regarding work incentives I conclude that one can certainly not make any sweeping arguments about moral hazard problems disappearing in the long-run. Contracts will clearly play an important role still. The relevant question to address then is whether the insights we have gained from studying one-period models will be significantly changed when looking at multi-period models. This of course, will require an explicit dynamic of contracting.

Regarding investment incentives, I note that dynamics is what seems to raise the problem in the first place, so in this case time appears to hurt rather than help reduce incentive costs. Perhaps this is the most interesting aspect of dynamics in the context of managerial incentives.

Acknowledgements. This paper was originally written in April 1982 for an unpublished volume in honour of the 60th birthday of Professor Lars Wahlbäck, Rector of the Swedish School of Economics and Business Administration in Helsinki, Finland. The research was supported by grants from the National Science Foundation and the Center for Advanced Studies in Managerial Economics at Northwestern University.

I wish to thank Milton Harris for conversations on the topic in 1982, the Editor, Patrick Bolton, for suggesting to publish the paper a decade and a half later and Jaime Ortega for proof-reading it. Barring minor technical corrections, I have made no changes to the original version even though some of the statements, particularly the third paragraph of the introduction, no longer reflect my current thinking. However, the abstract is new as there was none before. Also, I have added a few references to more recent papers that have adopted the same modeling approach to study reputation effects. This modeling approach, which Fudenberg and Tirole have given the apt name "signal jamming", differs from traditional signalling models in two respects: the agent's action is only observed with noise and hence the agent does not have to have any private information (other than about the actions taken).

REFERENCES

- AKERLOF, G. (1970), "The Market for Lemons: Quality Uncertainty and the Market Mechanism", *Quarterly Journal of Economics*, **84**, 488–500.
- FAMA, E. (1980), "Agency Problems and the Theory of the Firm", *Journal of Political Economy*, **88**, 288–307.
- GROSSMAN, S. and HART, O. (1983), "An Analysis of the Principal-Agent Problem", *Econometrica*, **51**, 7–45.
- HARRIS, M. and HOLMSTROM, B. (1982), "A Theory of Wage Dynamics", *Review of Economic Studies*, **49**, 315–333.
- HARRIS, M. and RAVIV, A. (1979), "Optimal Incentive Contracts with Imperfect Information", *Journal of Economic Theory*, **20**, 231–259.
- HOLMSTROM, B. (1979), "Moral Hazard and Observability", *Bell Journal of Economics*, **10**, 74–91.
- MEDOFF, J. and ABRAHAM, K. (1980), "Experience, Performance and Earnings", *Quarterly Journal of Economics*, **95**, 703–736.
- MIRLLEES, J. (1976), "The Optimal Structure of Incentives and Authority within an Organization", *Bell Journal of Economics*, **7**, 105–131.
- RADNER, R. (1981), "Monitoring Cooperative Agreements in a Repeated Principal-Agent Relationship", *Econometrica*, **49**, 1127–1148.
- ROSEN, S. (1982), "Authority, Control, and the Distribution of Earnings", *Bell Journal of Economics*, **13**, 311–323.
- ROSS, S. (1973), "The Economic Theory of Agency: The Principal's Problem", *American Economic Review*, **63**, 134–139.
- ROSS, S. (1977), "The Determination of Financial Structure: The Incentive Signaling Approach", *Bell Journal of Economics*, **8**, 23–40.
- RUBINSTEIN, A. (1981), "An Optimal Conviction Policy for Offenses That May Have Been Committed by Accident", in Brahm, Schotter and Schwodiauer (eds.) *Applied Game Theory* (Wein: Physica Verlag).
- SHAVELL, S. (1979), "Risk Sharing and Incentives in the Principal and Agent Relationship", *Bell Journal of Economics*, **10**, 55–73.
- WILSON, R. (1968), "The Structure of Incentives for Decentralization Under Uncertainty", *La Décision*, No. 171.

ADDITIONAL READING

- BORLAND, J. (1992), "Career Concerns: Incentives and Endogenous Learning in Labour Markets", *Journal of Economic Surveys*, **6**, 251–270.
- DEWATRIPONT, M., JEWITT, I. and TIROLE, J. (1999), "The Economics of Career Concerns, Part I: Comparing Information Structures", *Review of Economic Studies*, **66**, 183–201.
- FUDENBERG, D. and TIROLE, J. (1986), "A 'Signal-Jamming' Theory of Predation", *Rand Journal of Economics*, **17**, 366–376.
- GIBBONS, R. and MURPHY, K. (1992), "Optimal Incentive Contracts in the Presence of Career Concerns: Theory and Evidence", *Journal of Political Economy*, **100**, 468–505.
- HOLMSTROM, B. and RICART I COSTA, J. (1986), "Managerial Incentives and Capital Management", *Quarterly Journal of Economics*, **101**, 835–860.
- JEON, S. (1996), "Moral Hazard and Reputational Concerns in Teams: Implications for Organizational Choice", *International Journal of Industrial Organization*, **14**, 297–315.
- MEYER, M. and VICKERS, J. (1997), "Performance Comparisons and Dynamic Incentives", *Journal of Political Economy*, **105**, 547–581.
- MILGROM, P. and ROBERTS, J. (1988), "An Economic Approach to Influence Activities", *American Journal of Sociology*, **94**, Supplement, 154–179.
- NARAYANAN, M. P. (1985), "Managerial Incentives for Short-Term Results", *Journal of Finance*, **40**, 1469–1484.
- SCHARFSTEIN, D. and STEIN, J. (1990), "Herd Behavior and Investment", *American Economic Review*, **80**, 465–479.
- STEIN, J. (1989), "Efficient Capital Markets, Inefficient Firms: A Model of Myopic Corporate Behavior", *Quarterly Journal of Economics*, **104**, 655–669.