

The Holmström–Milgrom model: a simplified and illustrated version[☆]

Jon Lundesgaard*

Hedmark College, P.O. Box 104, N-2451 Rena, Norway

Accepted 30 September 1999

Abstract

In principal-agent theory, Holmström (1979, *Bell Journal of Economics*, 10, 74–91) offers the canonical model including a hidden one-shot action taken by an agent contracted to provide effort. Holmström's classical investigation led to an important body of applied literature. In investigating a specification in which effort is a sequence of actions, Holmström and Milgrom (1987, *Econometrica*, 55, 303–328) were able to provide a proof of the optimality of linear reward schemes (in the one-shot model, reward schemes are never linear). The sequence-of-actions model has a corresponding (static), very tractable companion. In this paper a simplified and illustrated version of this model is presented. © 2001 Elsevier Science Ltd. All rights reserved.

Keywords: Principal-agent theory; Moral hazard; Holmström–Milgrom model; Economics of information; Multitask theory

1. Introduction

Based on elements from Mirrlees (1975,1976), together with an important paper by Holmström (1979), the canonical model of hidden action was established. In principal-agent theory this is the classical model of moral hazard. Despite the fact that it is among the technically most demanding models in economic theory, it has been

[☆]The paper is an abbreviated version of a paper presented at the 14th Nordic Conference on Business Studies (14–16 August 1997 in Bodø, Norway) entitled "The Holmström–Milgrom Model: Applications and Extensions." A rather extensive survey included in Chapter III of Lundesgaard (1996) has been omitted from this shorter version. In our work on the present version, the comments of an anonymous referee and of the language editor have been most helpful.

* Tel.: + 47-62-43-04-51; fax: + 47-62-43-05-00.

E-mail address: j.lundesgaard@osir.hihm.no (J. Lundesgaard)

a direct source of a rich literature of applied work. The initial model has been a direct source of a rich literature of applied work. The initial model has been refined and extended in various directions. Among the most interesting of these is an extension by Holmström and Milgrom (1987) in which the (hidden) effort of the agent is represented in a manner that is less restricted. Instead of a one-shot action, effort is seen as a series of actions over a finite time interval. In introducing this more complex specification of effort a certain price is paid in that the model offered is less general as regards how uncertainty is introduced, and in the representation of the utility assessment of the parties involved. An assumption of the dynamic model is that results are observed continuously by the agent as they are generated over the time interval (i.e. not only as an aggregate at the end of the time interval). Only aggregates over the time interval are observed by the principal. The dynamic action specification has dramatic effects on the optimality of reward schemes. This is because, as well as the total of effort provided, the time profile of effort is also important. In the model investigated by Holmström and Milgrom a constant level of effort is optimal. This means that the agent is offered a reward that is optimally linear in the aggregate of outcomes produced. This is particularly interesting in that in practice we often observe less complex schemes, such as linear schemes. The prediction of the classical theory of hidden action is more complex schemes.¹ In addition, for the limit case of continuous action, there are interesting implications in terms of practical applicability. Thus, it is interesting that in the aggregates over a time interval the optimal linear model leads to a formal framework that is attractive due to its simplicity and tractability. This framework is a static version of the specification investigated by Holmström and Milgrom (1987), and it is simply referred to here as the Holmström–Milgrom model.

The history of the establishment and further development of hidden action theory is intriguing. It is a history that offers a fascinating account of analytical ingenuity, in interaction with practical relevance. In an admirable manner Hart and Holmström (1987) have surveyed important aspects of this history. Building on this survey, an introduction to the hidden effort problem and to Holmström and Milgrom (1987) was included in Chapter II of Lundesgaard (1996). In the following pages, and in a less specialised manner, a step-by-step outline of the Holmström–Milgrom model is presented. The model has been simplified as much as possible and a set of graphic illustrations has been utilised.

2. Some preliminaries to the Holmström–Milgrom model

In motivating and illustrating models it is common practice to relate the models to some sort of standard stories told. Thus, in moral hazard theory and the discussion of

¹The prediction of the classical theory is even stronger. In one of the most remarkable papers in the whole of this literature, namely in a student paper by Gjesdal (1976), it is demonstrated that linearity is contrary to the classical model.

hidden effort a standard story concerns a principal who is the owner of a technology producing certain outcomes, subject to exogenous uncertainty. Under additive uncertainty, for instance, the structure of this technology is represented by

$$x = f(a) + \varepsilon, \quad (1)$$

in which x represents outcomes. The a priori known and deterministic structural element of this technology is represented by $f(a)$, while a is the level of effort provided. Element ε is random, and in the Holmström–Milgrom model this stochastic element is random normal with zero mean. In the representation

$$x = a + \varepsilon, \quad (2)$$

outcome x is a noisy observation of effort a , or x is seen as a product that is one-to-one random normal in effort provided. Moreover, and alternatively, (2) is seen as the normalised version of the technology element of a model including (1). Henceforth, and for the sake of simplicity, our discussion is based upon (2).

The situation becomes interesting if the principal is unable to provide the needed effort himself. The problem then is to decide what sort of contract to write with the party who provides effort. Under full certainty, establishing a contract for the provision of a specific amount of effort by the other party (the agent) is a trivial matter. Things are different, however, if the principal cannot observe the amount of effort provided. In other words, it is not possible to write a contract to which the parties can commit themselves, with reference to some desired amount of effort. Hence, as an alternative to a contract regarding effort (observed directly under full certainty), the parties may establish their business relation on the basis of an observation correlated with effort provided. The result is a contract built on observations of a noisy or random kind. Otherwise, effort could have been inferred with full certainty from what is observed (so that the problem is eliminated). In (2), for instance, it is registered that outcomes x are random and correlated with effort a . Thus, if jointly observed and contractible, x in (2) provides an opportunity for the parties involved to establish a business relation. It is this business relation that is the object of the following exposition.

It is noted that the value of an observation, consisting only of realisations of the random result x , depends on the variance in the error term of the random element. In this connection, and more formally, Holmström and Milgrom use an interpretation of informativeness based on the *signal-to-noise ratio*.² For the more generally formulated technology $x = f(a) + \varepsilon$, the signal-to-noise ratio is $(df(a)/da)^2/\sigma^2$. For (2) we have

² The interpretation has been used in several published applications, such as Holmström and Milgrom (1994). In the classical model the likelihood ratio is important to an understanding of informativeness, that is to say, an understanding related to the problem of the testing of hypotheses in statistical theory. In the Holmström–Milgrom model, the understanding of information corresponds to that of estimation. Thus, consider an effort level a^* and the deviation Δa from this level. Outcomes are generated by $x = f(a^* + \Delta a) + \varepsilon \approx f(a^*) + f'(a^*)\Delta a + \varepsilon$. We then have the transformation $(x - f(a^*))/f'(a^*) = \Delta a + \varepsilon/f'(a^*)$. On the right-hand side and in taking advantage of the transformation of the original signal x in the observation of Δa , it is seen that the precision of the transformation in estimating Δa is $(f'(a^*))^2/\sigma^2$.

$df(a)/da = 1$, which gives the signal-to-noise ratio $1/\sigma^2$, often referred to as the precision of x (as an estimate of effort).

Further, it is registered that if effort a is provided at no costs to the agent, there is no problem. If the agent is indifferent between effort and no effort, he would have a weak preference to perform. Thus, for there to be a problem, it is important that the effort of the agent is costly. The costs of effort provided are represented by

$$c = c(a), \text{ in which } c'(a) > 0 \text{ and } c''(a) > 0 \quad (3)$$

(the costs of effort are strictly convex). A consequence of costly effort is that the principal-agent contract has to motivate the agent to provide effort. Again, the contract is based on a jointly observed x , and in the Holmström–Milgrom model this contract is optimally linear in x , so that the rewards offered are

$$s(x) = \alpha x + \beta. \quad (4)$$

Hence, the agent is exposed to some uncertainty, and this uncertainty increases in the contractual parameter α . The other contractual parameter β assumes the character of a fixed (non-random) transfer. The random net of the principal is

$$x - s(x) = (1 - \alpha)x - \beta, \quad (5)$$

and if risk neutral for instance, the certainty equivalent (and expected net) of this party is $\pi_p = (1 - \alpha)a - \beta$. That is to say, the principal's assessment of element $(1 - \alpha)x$ of the contract is in no way affected by transfer element β . For there to be an interesting problem, the agent has to be risk-averse. Otherwise, the problem is to find an optimal trade-off between the expected outcome ($E[x] = a$) and the costs $c(a)$ of producing outcomes. The optimal solution to this problem is always $\alpha = 1$.

In Holmström and Milgrom (1987), exponential utility is attributed to the contracting parties. In other words, if they are exposed to uncertainty (and are risk-averse), the parties' absolute degree of risk aversion is a constant. Hence, for the agent we have

$$u_A(\alpha x + \beta - c(a)) = -\exp\{-r_A(\alpha x + \beta - c(a))\}. \quad (6)$$

The certainty equivalent of the agent is π_A , and due to the combination of exponential utility, normal uncertainty and linearity, it is possible to find the amount (the risk premium) that renders the agent indifferent between a fixed π_A and $\alpha x + \beta - c(a)$. The first step is to utilise an exponential transform of the expected utility.³ For the agent, the exponential transform is

$$\begin{aligned} E[u_A(\alpha x + \beta - c(a))] &= \int -\exp\{-r_A(\alpha x + \beta - c(a))\} f(x) dx \\ &= -\exp\{-r_A(\alpha a + \beta - c(a)) + r_A \alpha^2 \sigma^2 / 2\}. \end{aligned}$$

³ Using the exponential transform means that the expected (exponential) utility is written in a similar manner to that of the utility function, i.e. on the form of $-\exp\{\cdot\}$. This transform is obtained in going through the steps of integration that follow from the definition of expected utility.

The certainty equivalent π_A is found in

$$-\exp\{-r_A(\alpha a + \beta - c(a)) + r_A\alpha^2\sigma^2/2\} = -\exp\{-r_A\pi_A\}, \text{ such that}$$

$$\pi_A = \alpha a + \beta - c(a) - r_A\alpha^2\sigma^2/2,$$

in which $r_A\alpha^2\sigma^2/2$ is the risk premium of the agent. In the characterisation of π_A above, it is important that the agent can be brought to any level of π_A by fixing β appropriately. It is also important that this is free of effects upon the agent's share of expected outcomes, his effort costs, and his costs due to exposure to risk. Apart from the outcomes x , it can be noted that both parties are fully informed about preferences, about the technology involved, about the agent's cost function (the agent's effort costs are unobserved by the principal, however), about the distributional and parametric aspects of uncertainty, and about the agent's opportunity loss from taking the job.⁴

3. The problem of the agent seen as an “inner” problem

In Fig. 1, in a flow-chart type of illustration, the problems facing the contracting parties are shown. There are two problems, interlinked as a problem within the problem. In this complex of problems, we should start by looking at the “inner” problem (the solution of this problem is important for the other problem). The “inner” problem is that of the agent, and it concerns finding an optimal level of effort. The character of the problem facing this party can be seen from his certainty equivalent. In providing effort, he has to find the right trade-off between a linearly increasing fixed share of x on the one hand and increasing convex costs of effort on the other. The overall problem is then to find the optimal parameters α and β of the contract. A more powerful incentive, i.e. a higher α , leads to more effort. With more effort, the expected outcome increases, and so do the costs of effort. In addition, with more powerful incentives the costs of exposure to risk also rise (a higher risk premium for the agent). For any α chosen, using the contractual instrument β , it will be remembered that the agent can be brought to any level of π_A . Hence, in the incentive parameter α , the problem of the principal is to find the right trade-off between the net of effort provided (efficiency in the provision of effort) and increasing costs of exposure to risk (efficiency in risk-sharing). In other words, in the Holmström–Milgrom model, the focus is on efficiency. Distribution issues are less interesting.

By first analysing the problem of the agent, analytical elements already introduced can be used. The agent's problem is to find an effort level a , such that π_A is maximised (for a given pair of $\{\alpha, \beta\}$). That is,

$$\text{Max}_a \alpha a + \beta - c(a) - r_A\alpha^2\sigma^2/2, \tag{7}$$

⁴ If the principle is the active party in the contracting process, such as in the role of a Stackelberg leader, the agent is in no need of knowing the principal's preferences.

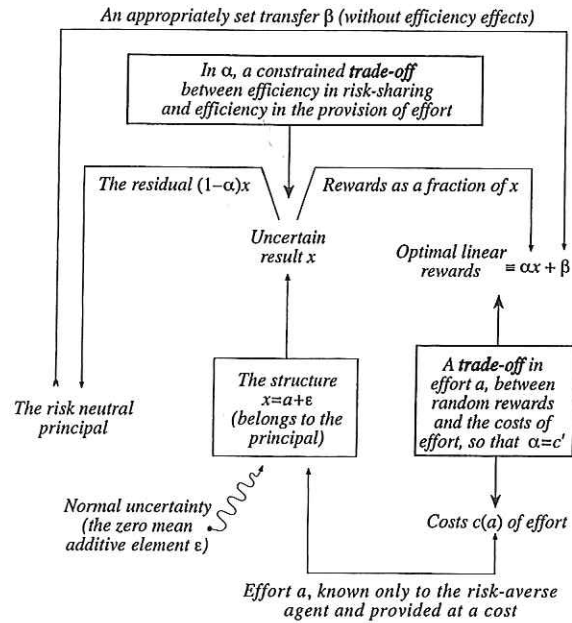


Fig. 1. The “inner” and the overall principal-agent problem.

giving the first-order condition

$$\alpha - c'(a) = 0 \quad \text{or} \quad \alpha = c'(a). \tag{8}$$

In (8), the effort level is determined, and the only parameter of importance in this context is α . The agent adjusts his effort level in such a way that the marginal benefits of effort (his marginal expected reward) are equal to his marginal cost of effort. Fig. 2, illustrates the problem of the agent in terms of the benefits and costs of effort and the corresponding marginal effects. The benchmark case $\alpha = 1$ is also included in the illustration.

The agent’s second-order condition is $-c''(a) < 0$, which in accordance with the assumption in (3) means that the cost function has to be strictly convex. Otherwise, as can be seen from the marginal effects shown in Fig. 2, the intersection point giving an optimal level of effort is not achieved. Further, the natural interpretation in Fig. 2 is that of $c(0) = 0$ and $c'(0) > 0$. If $c(0) > 0$, we can of course regard this as a set-up cost. Strictly speaking, for a cost function with $c'(0) > 0$, the agent’s first-order condition is $\alpha - c'(a) \leq 0$, and the optimal effort level is equal to zero if the inequality part is satisfied. Fig. 3 shows that in addition to $c'(0) > 0$, we could have $c'(0) = 0$ or $c'(0) < 0$. In the last case, effort is marginally pleasurable up to a certain level of effort.

Up to now the $c'(a)$ -curves have been convex. However, the curvature of $c'(a)$ is not important to the problem. What does matter is that $c'(a)$ is increasing (which means that $c(a)$ is convex). A linear $c'(a)$ means that $c''(a)$ is constant. As we are going to see, if




Fig. 2. The problem of the agent.

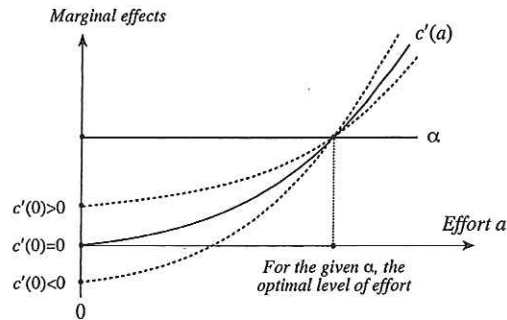


Fig. 3. Different assumptions regarding the marginal costs of effort.

this is the case, the understanding and determination of the incentive parameter α becomes simpler, which explains why this case is of some interest.

The effort level that gives an optimal inner solution from the agent's point of view can be found in (8). For each admissible value of the incentive parameter α , there is a level of effort that is optimal in this perspective. The slope characteristics of the agent's implied reaction function $a = a(\alpha)$ can be obtained from (8). Thus, starting from $\alpha = c'(a(\alpha))$, a marginal increase in α gives us $1 = c'' da/d\alpha$, such that

$da/d\alpha = 1/c''$. Reactions to incentives are positive, and they depend on the curvature of the agent's effort costs. A less curved cost function leads to stronger reactions to a marginal increase in the incentive parameter. The second-order derivative is generally a function of the effort level chosen, which means that (in general) reactions to incentives are going to vary as regards the levels of effort provided. For a constant second-order derivative, reactions to incentives are constant.

4. Finding the optimal incentives

The analysis of the second problem, i.e. the problem of finding the incentive parameter α of the optimal contract, is simpler than in the standard formulations of a principal-agent problem. First, instead of maximising the net of the principal, we may as well maximise the joint surplus of the agency (remembering that without any efficiency effects the agent reaches any level of his certainty equivalent by an appropriate fixing of β). Second, the incentive constraint is dropped, since the reaction function of the agent is inserted in the maximand. Thus, the joint surplus problem is

$$\underset{\alpha}{MAX} a(\alpha) - c(a(\alpha)) - r_A \alpha^2 \sigma^2 / 2, \quad (9)$$

which leads to the first-order condition

$$\frac{da}{d\alpha} - c' \frac{da}{d\alpha} - r_A \alpha \sigma^2 = 0, \quad \text{i.e.} \quad \frac{da}{d\alpha} - c' \frac{da}{d\alpha} = r_A \alpha \sigma^2. \quad (10)$$

The joint surplus in (9) is composed of two parts.⁵ The first part, $a(\alpha) - c(a(\alpha))$, is the expected net of effort provided. The second part, $r_A \alpha^2 \sigma^2 / 2$, gives the costs associated with exposing the agent to risk. In a contracting optimum, the marginal net of effort has to be equal to the marginal costs of exposing the agent to risk. Thus, alternatively, after an insertion of $\alpha = c'(a)$ and $da/d\alpha = 1/c''$ from the "inner" problem, we have

$$1 - \alpha = r_A \alpha \sigma^2 c''. \quad (11)$$

Instead of marginal effects in the incentive parameter α , the equilibrium condition (11) is appropriately interpreted as the marginal effects of an increased effort level a , on the part of the agent. The left-hand side is the marginal increase in the expected net $a - c(a)$ for a marginal increase in effort. The right-hand side is the agent's marginal risk premium for a marginal increase in effort. This can be seen by writing the risk premium as a function of effort, i.e. $r_A (\alpha(a))^2 \sigma^2 / 2$ ($\alpha = \alpha(a)$ is the inverse function of $a = a(\alpha)$). The interpretations connected with (10) and (11) are based upon regarding the maximand as composed of benefit and cost elements. In optimum these elements

⁵ The second-order condition of problem (9) is explored in Chapter II of Lundesgaard (1996)

are in a best possible manner traded off against each other, and the optimal α is obviously

$$\alpha = \frac{1}{1 + r_A \sigma^2 c''} \tag{12}$$

The incentive parameter increases in reductions in the agent’s risk aversion, and in the exogenously given uncertainty. In addition, α increases in stronger reactions to the fixed incentive. For the very simple version of the static model described here, it is a natural interpretation to see α as the fraction of x that should reward the agent (the share of the agent in the jointly observed result). For $r_A \rightarrow 0$ or $\sigma^2 \rightarrow 0$, we get $\alpha \rightarrow 1$. In addition, $\alpha \rightarrow 0$ follows from $\sigma^2 \rightarrow \infty$. An incentive equal to unity leads to a situation equivalent to first-best. At $c'(0) = 0$, for instance, with an incentive equal to zero, no effort is made. The bounds on values of the optimal incentives that follow from our discussion of (12), are then $0 < \alpha < 1$. More importantly, the strict version of these bounds is inappropriate since — for there to be a problem — it is assumed that r_A , σ^2 and c'' are non-zero and finite. However, tighter bounds could be placed upon α as a result of the specification of the agent’s cost function, and because the minimum expected net acceptable to the principal may have an effect. At $c'' = 0$, the agent is infinitely responsive to incentives, and at $c'' \rightarrow \infty$, we get $\alpha \rightarrow 0$. Finally, after having recognised that c'' is not necessarily a constant, it is too simplistic to say that α is computed by (12). In general, the optimal α is the value of the incentive parameter that satisfies $\alpha = 1/(1 + r_A \sigma^2 c''(a(\alpha)))$.

From the re-arranged version of the first-order condition in (10) and the insertion of $\alpha = c'(a)$ and $da/d\alpha = 1/c''$, we get the equilibrium condition

$$(1 - \alpha)/c'' = r_A \alpha \sigma^2. \tag{13}$$

Lastly, this alternative to (11) expressed in the incentive parameter α is investigated. Again, the left-hand side is the marginal increase in the expected net $a - c(a)$, and the right-hand side is the marginal risk premium. For constant reactions to incentives (c'' constant), this version of the equilibrium condition offers an interesting opportunity to illustrate the trade-off in the optimal contracting problem. This straightforward and simple illustration is given in Fig. 4.




Fig. 4. The determination of an optimal incentive parameter α .

5. Some observations

The classical principal-agent model, particularly starting from Holmström (1979) has been a very useful tool for applied research, for instance in accounting. The diffusion of applied research based on the optimal linear result, however, seems to have been slower.⁶ Given the merits of the linearity result, and the continuous search for new points of view that could produce further contributions, this is surprising. Some speculations are offered here. (i) The linear model is special, and seemingly less general and more restrictive. This may have led to certain reservation about basing applied work on the model.⁷ (ii) More specifically, the lack of wealth effects, in combination with the infinite fixed support of normal distributions, may have been seen as problematic. (iii) The classical principal-agent model and the Holmström–Milgrom model differ in the way information is understood. The simple signal-to-noise interpretation of the Holmström–Milgrom model could have been regarded erroneously as less fundamental. (iv) Linearity is contrary to the insights of the classical model, and with some justification, linearity on an ad hoc basis is seen as little short of a deadly sin, especially with respect to the investigation of problems of information. This (relevant) point of view may have been carried over to linearity that is a priori optimal. (v) Less crucially, the sunk costs of the expertise established could have been to the disadvantage of the Holmström–Milgrom model. In addition, there could be a snob effect arising from the mathematical simplicity of this model. Naturally, there are forces due to tractability that pull in the opposite direction.

Nevertheless, a significant amount of applied research has been carried out in the format of the linearity result of Holmström and Milgrom (1987). This research is surveyed fairly comprehensively in Lundesgaard (1996), which also presents some original applied work. The investigation of multitask problems is probably the most important offspring of the work on linear reward schemes. In a multitask problem there will be more than a single type of outcome, and the agent usually provides more than a single type of effort. In Holmström and Milgrom's original paper this problem was not investigated in any great detail (although the 1987 paper is based on a formulation of a multitask problem of some generality). Since then, particularly in Holmström and Milgrom (1991), the multitask problem has been investigated in a more specialised manner. The focus is on the effects of interaction among tasks due to the agent's effort costs. Feltham and Xie (1994) offer a more general formulation of the multitask

⁶ Circulation of the linearity result began in 1984, when it was presented in a working paper from Northwestern University. In 1985, it was included as Cowles Discussion Paper No. 742 in the working paper series of the Cowles Commission.

⁷ The seminal paper of Holmström and Milgrom (1987) on dynamic incentives represents an important start to a complex discussion of the underpinnings of analyses based on optimal linear incentives. Hellwig and Schmidt (1998) point out that elements in the argument of the initial discussion are rather sketchy, and it is their ambition to provide formal proofs of a more satisfactory kind. The problem in the original version to which attention is drawn refers to the underpinnings of the multitask version of the problem. It is believed that the problem focused by Hellwig and Schmidt may have contributed to the caution displayed by some researchers.

problem. In light of the importance of this problem, a presentation of the multitask problem based on their specification is included as an appendix to the present paper.

Appendix A. A generalisation to the multitask case

In this more general version of the Holmström–Milgrom model, the representation

$$\begin{aligned}
 x_1 &= \mu_{11}a_1 + \mu_{12}a_2 + \dots + \mu_{1n}a_n + \varepsilon_1, \\
 x_2 &= \mu_{21}a_1 + \mu_{22}a_2 + \dots + \mu_{2n}a_n + \varepsilon_2, \\
 &\vdots \\
 x_m &= \mu_{m1}a_1 + \mu_{m2}a_2 + \dots + \mu_{mn}a_n + \varepsilon_m
 \end{aligned}$$

is used.⁸ Error terms are jointly and normally distributed. In describing how the problem investigated differs from that of Holmström and Milgrom (1991), Feltham and Xie (1994) point to that “[m]uch of [their, i.e. H&M] analysis focuses on the alternative forms of [cost functions] and they simplify the impact of the performance measures by assuming that each measure is influenced by a single task (although there need not be a measure for each task [...].” Thus, in contrast, Feltham and Xie’s focus is shifted from the structure of costs of effort to the structure of information generating elements. Above, parameters μ_{ij} could be interpreted as zero or one inclusion parameters. For instance, for the multitask problem investigated by Holmström and Milgrom (1991) we have⁹

$$\mu = \begin{bmatrix} 1 & 0 & \cdot & 0 \\ 0 & 1 & \cdot & 0 \\ \cdot & \cdot & \cdot & \cdot \\ \cdot & \cdot & \cdot & \cdot \\ 0 & 0 & \cdot & 1 \end{bmatrix}$$

⁸ The expected values of the observations made are linear in the elements of effort provided (more compactly we have $x = \mu a + \varepsilon$). Further, it is noted that Holmström and Milgrom (1990) is based on the somewhat more general $x = f(a) + \varepsilon$. The corresponding formulation in classical principal-agent theory is $x = f(a, \theta)$ with θ being random.

⁹ Alternatively, this representation is seen as a normalisation of a representation in which the diagonal of μ consists of the non-zero impact elements $\mu_{ij}(i = j)$ that are allowed to differ from unity. Normalisation is performed by “pushing impact elements back into the cost function” (Holmström and Milgrom). Moreover, the case in which $\mu = [\mu_{ij}] = [1]$, is obviously less interesting. This case leads to linear dependence in expected values in the elements of vector x , and what is observed is a series of identical aggregates with different noise terms.

More generally, the parameters μ_{ij} are interpreted as impact on the observations made, per unit of effort provided. On the basis of the observations, and of the linear reward scheme, the agent is compensated.¹⁰ The agent's problem is

$$\begin{aligned} \text{Max}_{a_1, a_2, \dots, a_n} \sum_{i=1}^m \alpha_i \cdot (\mu_{i1} a_1 + \mu_{i2} a_2 + \dots + \mu_{in} a_n) \\ + \beta - c(a_1, a_2, \dots, a_n) - R_A(\alpha, \Sigma), \end{aligned}$$

in which $R_A(\alpha, \Sigma)$ is the risk premium implied by the variance in the problem as represented by Σ , and by the vector of incentives α . The first-order conditions are

$$\begin{aligned} \alpha_1 \mu_{11} + \alpha_2 \mu_{21} + \dots + \alpha_m \mu_{m1} - c'_1 &= 0, \\ \alpha_1 \mu_{12} + \alpha_2 \mu_{22} + \dots + \alpha_m \mu_{m2} - c'_2 &= 0, \\ &\vdots \\ \alpha_1 \mu_{1n} + \alpha_2 \mu_{2n} + \dots + \alpha_m \mu_{mn} - c'_n &= 0. \end{aligned}$$

Implied in the first-order conditions are the reaction functions $a_i = a_i(\alpha_1, \alpha_2, \dots, \alpha_m)$.¹¹ Thus, for the vector of marginal costs in the first-order conditions above, we get $c'_i(a_1(\alpha_1, \alpha_2, \dots, \alpha_m), a_2(\alpha_1, \alpha_2, \dots, \alpha_m), \dots, a_n(\alpha_1, \alpha_2, \dots, \alpha_m))$. We have n equilibrium conditions. Implicit reactions to incentives can then be found from these conditions. For each of m incentives, there are n equations, such that

$$\begin{aligned} \mu_{11} &= c''_{11} \frac{\partial a_1}{\partial \alpha_1} + c''_{12} \frac{\partial a_2}{\partial \alpha_1} + \dots + c''_{1n} \frac{\partial a_n}{\partial \alpha_1}, \\ &\vdots \\ \mu_{1n} &= c''_{n1} \frac{\partial a_1}{\partial \alpha_1} + c''_{n2} \frac{\partial a_2}{\partial \alpha_1} + \dots + c''_{nn} \frac{\partial a_n}{\partial \alpha_1}, \end{aligned}$$

¹⁰ Feltham and Xie are careful to point out that rewards in their model are linear on an ad hoc basis. That is to say, they do not appeal to the arguments for optimal linearity found in Holmström and Milgrom (1987). Strictly speaking, in using the tractability of static versions of the Holmström-Milgrom model, optimal linearity should be checked by going back to an underlying dynamic version. Furthermore, in investigating specifications that are more complex and possibly different in some aspects of their underlying dynamics compared with what was investigated by Holmström and Milgrom (1987), the problem of optimality in a dynamic version becomes more critical. Thus, in basing their analysis of a complex static problem on the ad hoc argument, Feltham and Xie (1994) show a necessary degree of caution on this point. It is observed that Hellwig and Schmidt (1998) appear to offer some help in this connection.

¹¹ In the case analysed by Feltham and Xie $c''_{ij} = 0$ (for $i \neq j$), due to the additive character of their cost function. They also have constant marginal reactions to incentives ($c''_{ii} = \text{constant}$). In other words, the case of Feltham and Xie is void of some of the effects that are particularly interesting in Holmström and Milgrom (1991).

and

$$\begin{aligned} \mu_{n1} &= c''_{11} \frac{\partial a_1}{\partial \alpha_n} + c''_{12} \frac{\partial a_2}{\partial \alpha_n} + \dots + c''_{1n} \frac{\partial a_n}{\partial \alpha_n}, \\ &\vdots \\ \mu_{mn} &= c''_{n1} \frac{\partial a_1}{\partial \alpha_n} + c''_{n2} \frac{\partial a_2}{\partial \alpha_n} + \dots + c''_{nn} \frac{\partial a_n}{\partial \alpha_n}. \end{aligned}$$

and

$$\begin{aligned} \mu_{m1} &= c''_{11} \frac{\partial a_1}{\partial \alpha_m} + c''_{12} \frac{\partial a_2}{\partial \alpha_m} + \dots + c''_{1n} \frac{\partial a_n}{\partial \alpha_m}, \\ &\vdots \\ \mu_{mn} &= c''_{n1} \frac{\partial a_1}{\partial \alpha_m} + c''_{n2} \frac{\partial a_2}{\partial \alpha_m} + \dots + c''_{nn} \frac{\partial a_n}{\partial \alpha_m}. \end{aligned}$$

In all there are $m \times n$ unknown marginal reactions and $m \times n$ equations. The system of equations consists of m blocks, such that block i is attributed with the structure $[\mu_{ij} = c''[\partial a_j / \partial \alpha_i]]$. The admissible set of actions taken by the agent is represented by \mathcal{A} , and for the vector \mathbf{a} of effort provided by the agent we have $\mathbf{a} \in \mathcal{A}$. In hidden action theory, however, there are no strong traditions for this sort of a priori restriction on the action set of the agent. Nevertheless, in many cases the destruction of output is not allowed for. The standard assumption, applied to the case investigated, is simply that of $\mathbf{a} \in \mathfrak{R}_+^n$. In the multitask problem investigated by Feltham and Xie, the problem of restrictions is less straightforward. In a multitask agency a desired set of actions is implemented by the vector of incentive parameters α of the contract. That is, the agency faces the question of possible limitations in the admissible set of actions, due to the two parties' access to informative and contractible observations. The set of observations obviously affects the vector of instruments α that is permitted and the set of actions \mathbf{a} for which implementation is a possibility. The access to independent and informative elements in the set of observations is important to the creation of incentive instruments that permit the implementation of desired actions. For the structure presented above, with reference to a simple set of arguments, attention is drawn to what is believed to be essential in Feltham and Xie's paper.

- First, consider the case in which $m = n$. For this case, the number of possible instruments in the form of incentives is equal to n . Further, this number is equal to the number of actions taken by the agent, so that

$$\begin{aligned} a_1 &= a_1(\alpha_1, \alpha_2, \dots, \alpha_n), \\ a_2 &= a_2(\alpha_1, \alpha_2, \dots, \alpha_n), \\ &\vdots \\ a_n &= a_n(\alpha_1, \alpha_2, \dots, \alpha_n). \end{aligned}$$

In block i , given that c''^{-1} exists, we have $[\partial a_j / \partial \alpha_i] = c''^{-1}[\mu_{ij}]$ so that $[\partial a_i / \partial \alpha_j] = c''^{-1} \mu$ for the matrix of all marginal reactions. We could of course also see the actions of the agent as ends attained by means of incentives contractually

implemented. That is, for a given set of targeted ends, the question concerns finding an appropriate set of means.¹² That is, more formally, the problem of finding the functions

$$\begin{aligned}\alpha_1 &= \alpha_1(a_1, a_2, \dots, a_n), \\ \alpha_2 &= \alpha_2(a_1, a_2, \dots, a_n), \\ &\vdots \\ \alpha_n &= \alpha_n(a_1, a_2, \dots, a_n).\end{aligned}$$

According to the Inverse Function Theorem, the marginal effects of implemented incentives, with respect to the levels of effort desired, are found in $[\partial\alpha_i/\partial a_j] = \mu c''$. Thus, if μ exists, an inverted “from-ends-to-means” form can be seen to exist for any $a \in \mathfrak{R}_+^n$. This inverted “from-ends-to-means” form is a sort of mirror image of the original “from-means-to-ends” reaction functions $a_1 = a_1(\alpha)$, $a_2 = a_2(\alpha)$, ..., $a_n = a_n(\alpha)$. In the multitask case focused by Holmström and Milgrom (1991) the same sort of correspondence is established, i.e. on a basis of the inverse function argument, $[\partial a_i/\partial \alpha_j] = [c''_{ij}]^{-1}$ is found after first establishing $[\partial \alpha_i/\partial a_j] = [c''_{ij}]$. Further, when $m = n$, and given a μ that is identical with a unity matrix, the original system of implied reactions to incentives is reduced to

$$\begin{aligned}1 &= c''_{11} \frac{\partial a_1}{\partial \alpha_1} + c''_{12} \frac{\partial a_2}{\partial \alpha_1} + \dots + c''_{1n} \frac{\partial a_n}{\partial \alpha_1}, \\ 1 &= c''_{21} \frac{\partial a_1}{\partial \alpha_2} + c''_{22} \frac{\partial a_2}{\partial \alpha_2} + \dots + c''_{2n} \frac{\partial a_n}{\partial \alpha_2}, \\ &\vdots \\ 1 &= c''_{n1} \frac{\partial a_1}{\partial \alpha_n} + c''_{n2} \frac{\partial a_2}{\partial \alpha_n} + \dots + c''_{nn} \frac{\partial a_n}{\partial \alpha_n}.\end{aligned}$$

Thus far, as regards the possibility of an implementation of any $a \in \mathfrak{R}_+^n$, the conclusion is that the specifications investigated by Holmström and Milgrom (1991) and Feltham and Xie (1994), lead to the same sort of insights. The case in which $m = n$ does not include any additional restrictions on the set of possible actions. In other words, we have to look to other aspects of the problem to find new and interesting dimensions in multitask theory.

- Second, consider the situation in which $m > n$. In this case the number of possible instruments is greater than the number of tasks to which the agent is supposed to

¹² When seen as expected values of observations, the approach is analogous with a well-known discussion started by Tinbergen (1952). A useful summary of this discussion can be found in the published volumes of Leif Johansen's (1977, 1978) lecture notes. In the second volume of these notes Johansen focuses on problems under uncertainty and the volume includes a discussion of a case resembling the one investigated by Feltham and Xie (1994).

direct his effort. Thus, the agency has access to more instruments than are strictly needed to implement a set of desired and admissible actions, i.e. there are alternatives when it comes to the implementation of actions. Consequently, in the implementation of a specific (targeted) allocation of effort, it is possible for the agency to take other concerns into account. In the problem investigated, the most immediate concern refers to the cost of exposing the agent to risk. In finally identifying a set of optimal incentives, this is of course again a problem about the trade-off between gross benefit and the costs of effort (including cost of exposure to risk). It can be seen that the $m > n$ case is a more general version of a classical problem of hidden action theory about the use and value of added information, i.e. the problem of the value of added information, when the information to which the agency already has access does not mean that restrictions are added to $a \in \mathfrak{R}_+^n$.¹³

- Third, consider the situation in which $m > n$. In this case there is a lack of instruments, and the effect of this is dramatic: the agency is unable to implement all $a \in \mathfrak{R}_+^n$, and we represent the implementable set of actions by the more restricted set $A^x \subset \mathfrak{R}_+^n$. From a practical point of view the case is both relevant and interesting, and it is the one focused by Feltham and Xie (1994).

In the effort levels a_1, a_2, \dots, a_n , the expected gross benefit at the principal's disposal is $\pi(a_1, a_2, \dots, a_n)$. In any of the three types of case above, we have $a_1(\alpha), a_2(\alpha), \dots, a_n(\alpha)$ for $A^x \subseteq \mathfrak{R}_+^n$ (the possibility that the full $a \in \mathfrak{R}_+^n$ is implementable is included). In the incentives, the formulation of the joint problem is

$$\text{MAX}_{\alpha} \pi(a_1(\alpha), a_2(\alpha), \dots, a_n(\alpha)) - c(a_1(\alpha), a_2(\alpha), \dots, a_n(\alpha)) - R_A(\alpha, \Sigma).$$

A detailed exploration of solutions to this very general formulation of the problem will not be made here. The kinds of solution that the third type of case leads to are particularly interesting, however, and some further comments will be made on their structure.

To be more specific about the consequences of the observations made, it can be said that the character of the information problem is determined by the structure of both μ^x and Σ^x (the vector of observations is represented by x). Starting from the observation of the x , the statistical inference problem is to establish an estimate of elements in the effort provided—a problem that seems to be related to the problem of identification in econometrics. The problem associated with the third type of case is clearly present in the case with at least a pair of tasks, and it arises because an information structure restricts the possible inferences that can be made to an aggregate of the efforts provided. Hence, in the relevant aggregate the allocation across tasks is left to the discretion of the agent. To put it more technically, the character of the problem is determined by the structures of μ^x and Σ^x (characteristics such as sparsity and linear independence are important in this connection).

¹³ The seminal contributions on this point are those of Holmström (1979) and Shavell (1979).

The costs to the agency due to this sort of problem are determined by the distance between the allocation chosen by the agent and what would otherwise have been implemented (taking the preferences of the principal into account). In other words this is a question of the preferred choice of the agent compared with the preferences of the principal, which is also a question of the degree of congruence (or non-congruence). An interesting aspect of Feltham and Xie's specification is that these authors are able to propose a precise measure in this respect. The effects of the information problem just described can naturally be more — or less — severe depending on the degree of congruence.

A possible and less technical way of looking at the traditional analysis of added information, is that in a first round more information reduces the exposure to risk. In a second round the exposure to risk increases because an increase in the power of incentives is optimal. The net effect of this is an increase in the joint surplus. This is the problem described in the second type of case above (the agent is not given any discretion in his allocation of effort over tasks). The third type of case is different in this respect and thus contributes to new insight into the problem of the value of added information. Suppose that instead of the original vector of observations x , the more informational rich x^+ is observed (x^+ includes the original vector of observations x). This may well mean that the implementable set of actions is expanded, such that $A^x \subset A^{x^+} \subseteq \mathfrak{R}_+^n$. This may contribute to increased congruence, which means that the problem of the value of added information is made more complex and interesting.

Feltham and Xie (1994) was important in drawing attention to an interesting information problem that had been overlooked in the literature. Furthermore, the way in which their specification includes more traditional ones, is attractive. Holmström and Milgrom (1991) bring the problem of congruence into their discussion of immeasurability problems. However, their focus on interesting cost effects meant that the congruence problem was disregarded. As to the problem of added information, Feltham and Xie demonstrate the importance of information that is not only indirectly but also directly informative about the effort provided. This has meant a breakthrough in our understanding of the Controllability Principle in accounting (agents should be evaluated on a basis of observations that they are to some extent in control of themselves). It seems to be recognized that some correlation characteristics can render the contractual inclusion of a directly informative observation less favourable, or even (weakly) unprofitable.¹⁴ In a situation where this leads to greater congruence, a conclusion about the zero value of information could of course be altered. The problem of congruence obviously adds in a significant way to our understanding of the value of information. Further, in their investigation Feltham and Xie are able to bring incentive theory into line with more traditional ideas about planning, and there are interesting parallels here with the problem of identification in econometrics.

¹⁴ See, Lundesgaard (1996), Chapter VII for a formal discussion of this topic.

References

- Feltham, G. A., & Xie, J. (1994). Performance measure congruity and diversity in multi-task principal/agent relations. *Accounting review*, 69, 429–453.
- Gjesdal, F. (1976). Accounting in agencies. Manuscript at the Graduate School of Business, Stanford University.
- Hart, O., & Holmström, B. (1987). The theory of contracts. In T. F. Bewley, *Advances in Economics Theory*. Cambridge: Econometric Society Monographs, Cambridge University Press.
- Hellwig, M., & Schmidt, K. M. (1998). Discrete-time approximations of the Holmström-Milgrom Brownian-motion model for intertemporal incentive provision. Manuscript dated May 5 1998.
- Holmström, B. (1979). Moral hazard and observability. *Bell Journal of Economics*, 10, 74–91.
- Holmström, B., & Milgrom, P. (1987). Aggregation and linearity in the provision of intertemporal incentives. *Econometrica*, 55, 303–328.
- Holmström, B., & Milgrom, P. (1990). Regulating trade among agents. *Journal of Institutional and Theoretical Economics (Zeitschrift für die gesamte Staatswissenschaft)*, 146, 85–105.
- Holmström, B., & Milgrom, P. (1991). Multitask principal-agent analyses: Incentive contracts, asset ownership and job design. *Journal of Law, Economics & Organization*, 7 (Special Issue), 24–52.
- Holmström, B., & Milgrom, P. (1994). The firm as an incentive system. *American Economic Review*, 84, 972–991.
- Johansen, L. (1977). *Lectures on macroeconomic planning, Part 1 — General aspects*. Amsterdam: North-Holland.
- Johansen, L. (1978). *Lectures on macroeconomic planning, Part 2 — Centralization, decentralization, planning under uncertainty*. Amsterdam: North-Holland.
- Lundesgaard, J. (1996). *Information, budgeting and organization: Excursions into an optimal linear approach to the analysis of moral hazard*. Dissertation Series, Vol. 41. Economics Department, University of Oslo.
- Mirrlees, J. (1975). *The theory of moral hazard and unobservable behavior — Part I*: Nuffield College, Oxford.
- Mirrlees, J. (1976). The optimal structure of authority and incentives within an organization. *Bell Journal of Economics*, 7, 105–131.
- Shavell, S. (1979). Risk sharing and incentives in the principal and agent relationship. *Bell Journal of Economics*, 11, 55–73.
- Tinbergen, J. (1952). *On the theory of economic policy*. Amsterdam: North-Holland.