European Journal of Operational Research 70 (1993) 83–103
North-Holland

Theory and Methodology

The principal/agent paradigm: Its relevance to various functional fields

Ann van Ackere
London Business School, Sussex Place – Regent's Park, London NW1 4SA, UK

Received June 1990; revised November 1991

Abstract: The main characteristics of the basic principal/agent model are discussed. Several applications of this model to the fields of accounting, industrial organization, finance and marketing are presented. The potential for application to management science problems is highlighted, and a batch-size problem is analyzed to illustrate the concepts involved. The paper concludes with a discussion of the major criticisms of the agency model.

Keywords: Principal/agent theory; Management science; Batch size

1. Introduction

Agency theory has been around for over twenty years and has been applied extensively in a variety of contexts including accounting, finance, industrial organization and marketing. Baiman (1982, p. 155) gives a concise definition of what that theory is about: "Agency theory research focuses on the optimal contractual relationships among members of a firm, where each member is assumed to be motivated solely by self-interest". As Ross (1973) points out, most contractual arrangements contain an element of agency. The literature on moral hazard is also closely related.

The basic model considers the interaction between two individuals, one of which is called the principal, and the other the agent. In the simplest situation the principal would like to delegate a task to the agent. He wants to determine what (minimal) reward he should offer the agent to ensure that
(1) the agent accepts to perform the task, and
(2) the agent performs the task in a satisfactory way.

The former constraint is usually referred to as the *Individual Rationality Constraint* (do not accept to perform a task if this makes you worse off). The latter constraint is commonly labelled *Incentive Compatibility Constraint* (motivate the agent to act in the interest of the principal, even when the principal cannot observe the action).

In the classical example of agency theory the principal is interpreted as being the shareholder or owner of a company and the agent its manager. Attention to the incentive problems ensuing from the agent managing the principal's capital goes back to Adam Smith (1776, p. 700), who already noted that "The directors of such [joint-stock] companies, however, being the managers rather of other people's money than of their own, it cannot well be expected, that they should watch

Correspondence to: A. van Ackere, London Business School, Sussex Place – Regent's Park, London NW1 4SA, UK

0377-2217/93/$06.00

over it with the same anxious vigilance with which the partners in a private copartnery frequently watch over their own... Negligence and profusion, therefore, must always prevail, more or less, in the management of the affairs of such a company".

Similar conflicts of interests arise in many situations. Still, these are often overlooked in traditional approaches. For instance, most studies of scheduling rules look at issues such as minimizing expected delays, maximizing throughput, etc, subject to some capacity constraint. Little attention is paid to whether or not the resulting schedule can be implemented. A medical practice may schedule patients close to each other, to avoid practitioners having to wait for the next patient. Patients will soon realize that if they arrive at the scheduled time, they are likely to face a considerable delay, and will adapt their behaviour accordingly, thereby defeating the purpose of the original schedule. Unless the incentives of all parties are taken into account, this type of behaviour is bound to occur.

Ross (1973) was one of the first to present a formal analysis of the principal/agent problem in its standard form. Spence and Zeckhauser (1971) had considered a related problem in the context of insurance contracts. They study the optimal form of the insurance contract, under the constraint that the insurance company (the principal) should break-even (rather than maximize his profits). Wilson (1968) discusses how a group of individuals can optimally reach a common decision under uncertainty, in a situation where this decision results in a pay-off to be shared among the members of a group. Stiglitz (1974) has discussed the appropriateness of the sharecropping system to solve incentive and risk-sharing problems. Most earlier work had overlooked the high-risk aspect of agriculture. Stiglitz formulates a simple competitive agricultural model, which allows him to focus on the risk-sharing and incentive properties of alternative contracts such as sharecropping, pure wage systems, pure rental systems, or a mixture of the latter two. He also incorporates the impact of the risk-aversion of workers and landlords and points out that when the level of effort is not observable, share-cropping has positive incentive effects.

Empirical work relating to the principal/agent paradigm is often limited to a detailed study of one or two sites, as the opportunities to test a variety of reward systems in comparable environments is generally limited. See for instance Blanchard, Chow and Noreen (1986). Alternatively, one might consider the use of experimental studies, as for instance in Waller and Bishop (1990). They conduct two experiments to test the efficiency of various payment schemes in coping with the problems of misrepresentation of private information by agents, and of consumption versus investment of allocated resources.

As mentioned previously, the basic principal/agent model and many of its extensions have been applied in a variety of functional fields. Surprisingly, very few applications to the field of Management Science (MS) appear in the literature, although many problems addressed by MS could benefit from this approach.

The remainder of this paper is structured as follows. In section 2, I present the basic principal/agent model more formally and discuss its main characteristics. In Section 3, I briefly discuss some of the many extensions of the basic model. Section 4 contains a selection of applications of the basic model and its extensions to the fields of accounting, industrial organization, marketing and finance. Section 5 discusses potential applications of agency theory to the field of MS. Section 6 contains some concluding remarks.

2. The basic principal / agent model

In the basic model, the principal (P) wants to hire an agent (A) to perform a job (e.g. preparing a bid for a contract). The outcome of the job (x; e.g. whether or not the contract is won) depends on two elements:

(1) the level of effort exerted by the agent (e; e.g. time-input) which is a disutility for the agent and cannot be observed by the principal, and

(2) a random factor (r; e.g. the bid of the competitor) which the agent learns after selecting his level of effort and which the principal does not observe.

Both the principal and the agent observe the outcome. The principal is assumed to be risk-neutral.

The principal's aim is to design a payment scheme (I), also called an incentive scheme, that maximizes his utility (U_P), while ensuring that

(1) the agent accepts to perform the job, and
(2) the agent selects the effort-level the principal favours.

The agent can refuse to perform the job. If he rejects the principal's contract, he obtains a certain utility level, usually called his reservation level ($\bar{U}_A$; for instance, unemployment benefit). Consequently, he will only accept the contract if his expected utility (denoted by U_A) exceeds this reservation level.

Letting E_r denote the expected value with respect to the random element r, the principal's problem can be formulated as follows:

$$\max_{I(x),e} E_r\{U_P(I(x), x(e, r))\}$$

subject to

$$E_r\{U_A(I(x(e, r)), e)\} \geq \bar{U}_A, \qquad (*)$$

$$E_r\{U_A(I(x(e, r)), e)\} \geq E_r\{U_A(I(x(e', r)), e')\} \quad \text{for all } e' \neq e. \qquad (**)$$

Constraint $(*)$ is the *individual rationality constraint*, constraint $(**)$ the *incentive compatibility constraint*. Note that when selecting a payment scheme ($I(x)$), the principal implicitly determines the agent's optimal effort level. This implies that the effort level e is, in fact, a decision variable of the principal. It is therefore usually included in the choice variables as above, in case there is a tie in the agent's preferences.

The following example illustrates the concepts introduced so far. The principal wants the agent to prepare a bid for a contract, for which there is only one other competitor. To keep things simple, assume that the agent can either exert a high level of effort ($e = 1$) or a low level of effort ($e = 0$). Assume also that the bid of the competitor is either competitive ($r = 1$) or overpriced ($r = 0$), and let the outcome be $x = 1$ if the contract is obtained and $x = 0$ if the contract is lost. Assume that the probabilities of winning the contract, given the agent's effort level e and the quality of the opponent's bid r, are as shown in Table 1.

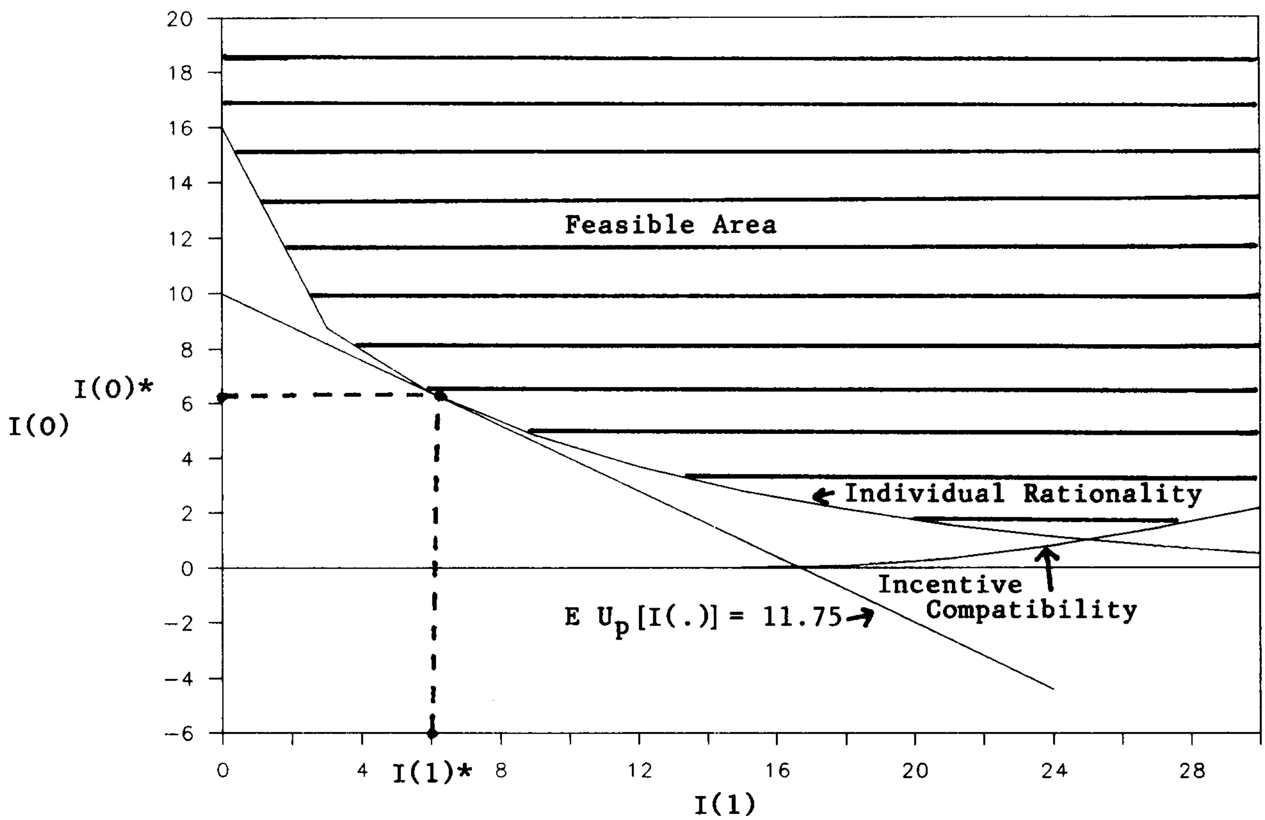

Figure 1a. $e = 1$

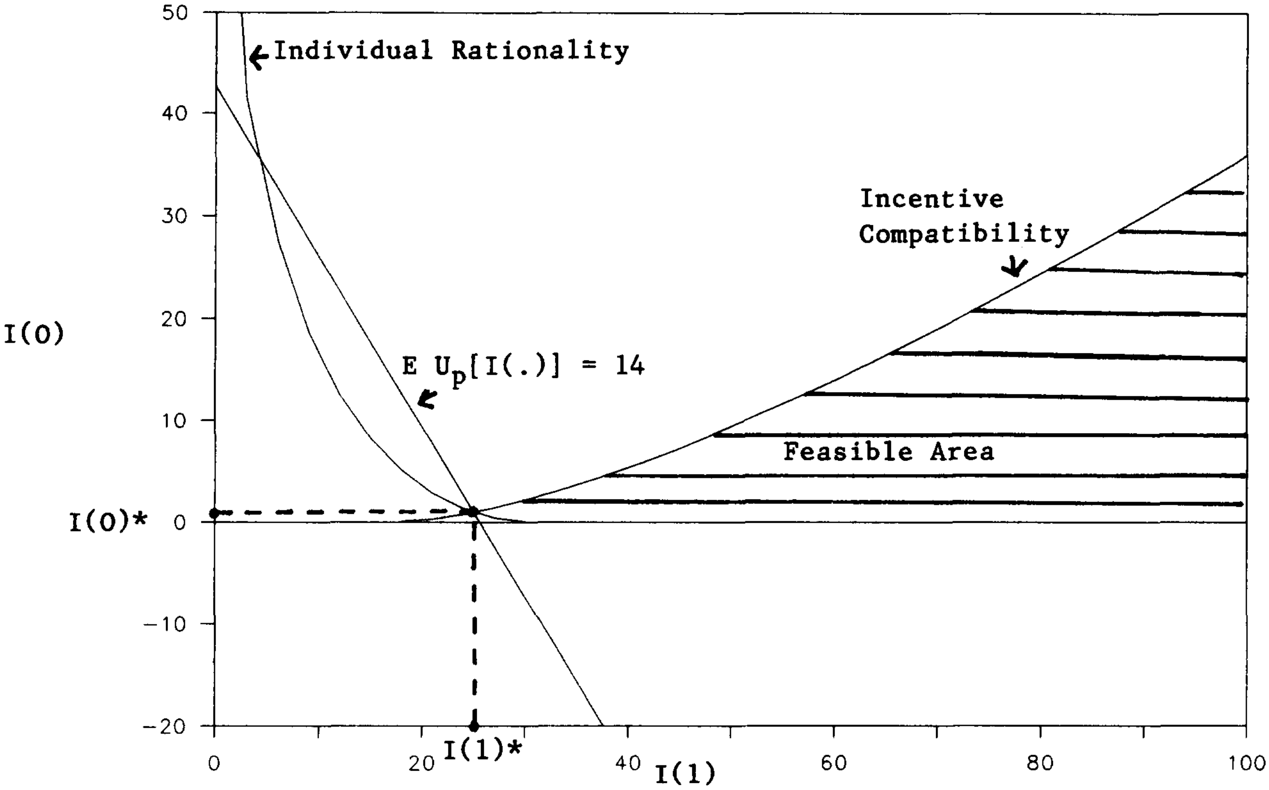

Figure 1b. $e = 0$

Assume further that the agent and the competitor can obtain no information about each other's bid before having to submit their own. Based on prior experience, the principal estimates that the competitor is equally likely to submit a competitive bid or an overpriced bid. The principal's profit (excluding any payments to the agent) is 48 if the bid is won and 0 if the bid is lost. This implies that $U_P(x) = 48x - I(x)$. Let us assume that the agent's utility equals $\sqrt{I(x)} - e$, and that his reservation utility equals 2.5. For this example the principal's problem can be written as follows:

$$\max_{I(x),e} E_r\{48x(e, r) - I(x(e, r))\}$$

subject to

$$E_r\left\{\sqrt{I(x(e, r))} - e\right\} \geq 2.5,$$

$$E_r\left\{\sqrt{I(x(e, r))} - e\right\} \geq E_r\left\{\sqrt{I(x(e', r))} - e'\right\}$$

for $e' \neq e$.

The principal has two options: induce a high level of effort ($e = 1$) or induce a low level of effort ($e = 0$). If he decides to induce a high level of effort, his problem becomes

$$\max_{I(x)} (48 - I(1)) * \tfrac{5}{8} + (-I(0)) * \tfrac{3}{8}$$

subject to

$$\left(\sqrt{I(1)} - 1\right) * \tfrac{5}{8} + \left(\sqrt{I(0)} - 1\right) * \tfrac{3}{8} \geq 2.5,$$

$$\left(\sqrt{I(1)} - 1\right) * \tfrac{5}{8} + \left(\sqrt{I(0)} - 1\right) * \tfrac{3}{8} \geq \sqrt{I(1)} * \tfrac{3}{8} + \sqrt{I(0)} * \tfrac{5}{8}.$$

Table 1

r	e	Probability that	
		$x = 0$	$x = 1$
0	0	0.5	0.5
0	1	0.25	0.75
1	0	0.75	0.25
1	1	0.5	0.5

Letting λ and μ denote the Lagrange multipliers of the first and second constraint, some algebra yields the following result: $I(1) = 25$, $I(0) = 1$, $\lambda = 7$, $\mu = 7.5$, $E_r(U_P) = 14$, $E_r\{U_A(e = 1)\} = 2.5 = E_r\{U_A(e = 0)\}$; i.e., the agent is indifferent between the two levels of effort. This is illustrated in Figure 1(a). I assume that, when indifferent between two actions and their associated rewards, the agent selects the action favoured by the principal.

If the principal decides to induce a low level of effort, his problem becomes

$$\max_{I(x)} (48 - I(1)) * \tfrac{3}{8} + (-I(0)) * \tfrac{5}{8}$$

subject to

$$\sqrt{I(1)} * \tfrac{3}{8} + \sqrt{I(0)} * \tfrac{5}{8} \geq 2.5,$$

$$\sqrt{I(1)} * \tfrac{3}{8} + \sqrt{I(0)} * \tfrac{5}{8} \geq \left(\sqrt{I(1)} - 1\right) * \tfrac{5}{8} + \left(\sqrt{I(0)} - 1\right) * \tfrac{3}{8}.$$

Letting again λ and μ denote the lagrange multipliers, some algebra yields $I(1) = I(0) = 6.25$, $\lambda = 5$, $\mu = 0$, $E_r\{U_P\} = 11.75$, $E_r\{U_A(e = 0)) = 2.5 > E_r\{U_A(e = 1)) = 1.5$. This situation is illustrated in Figure 1(b). Note that in this case the agent's reward does not depend on the outcome (i.e., no risk-sharing takes place), while if the principal decides to induce the higher level of effort, the agent has the same expected utility, but his actual reward could be 1 or 25, depending on the outcome.

It follows that in this example the principal is better off offering an incentive contract $I(1) = 25$, $I(0) = 1$, to induce a high effort level ($e = 1$).

If the value to the principal of winning the bid was only 32 (instead of 48), resolving this problem would yield that the principal can either offer a contract $I(1) = 25$, $I(0) = 1$ to induce an effort level $e = 1$ and get an expected utility of 4, or offer a contract $I(1) = I(0) = 6.25$ to induce an effort level $e = 0$ and get an expected utility of 5.75. In this case it is not worth offering an incentive contract, inducing the agent to expend a high level of effort: the increase in expected utility due to the higher probability of obtaining the contract is more than offset by the required increase in the agent's reward.

It is worth comparing these results to a situation without incentive problems; i.e., either the principal observes the agent's level of effort, or he is able to trust him. In this case he could offer a contract $I(1) = I(0) = 12.25$ if $e = 1$, nothing if $e = 0$. The agent would be better off choosing $e = 1$, as this guarantees him a utility of $\sqrt{12.25} - 1 = 2.5$. The principal would have an expected utility of $30 - 12.25 = 17.75$ in the first example, and $20 - 12.25 = 7.75$ in the second example. In both cases he is strictly better off, while the agent has the same expected utility but no longer carries any of the risk.

The model presented so far is known as the basic moral hazard model (see for instance Demski and Feltham, 1978). It is useful at this point to mention the adverse selection model, which can be interpreted as a special case of moral hazard. In the adverse selection model, the principal again has a task that he wants to delegate to an agent. But now there are several candidates to perform this task, each with a different level of ability. The principal is unable to observe the ability of the agents.

In the moral hazard model, the principal was unable to observe the agent's level of effort. He therefore designed a contract that would induce the agent to select an appropriate level of effort. In the adverse selection model, rather than attempting to induce the appropriate level of effort from a specific agent, the principal tries to select an agent with the appropriate level of skill. He will therefore write one (or more) contracts such that only the agents with the desired ability accept a specific contract.

One of the first papers to address this issue of asymmetric information about ability (the agent knows his ability, but the principal does not) or, more generally, about quality, is Akerlof's model of 'The Market for Lemons' (Akerlof, 1970), where he analyses the problems of asymmetric information about quality in the context of the second hand car market.

In the remainder of this section I will focus on the basic moral hazard model. Let us first consider the case where the principal is able to observe the agent's level of effort. In this case the incentive scheme can depend on the level of effort. Consequently, it is quite easy for the principal to write a forcing contract: if the agent selects the desired effort-level he receives a payment insuring a utility $\bar{U}_A$, otherwise he receives less than $\bar{U}_A$. Consequently, the agent will accept the contract and select the appropriate effort-level. Whenever the effort-level is not observable, the incentive scheme cannot depend on effort and a forcing contract is not implementable.

There may be situations where effort is observable, and still forcing contracts cannot be selected due to laws and regulations. For instance, a union agreement may prohibit wages below $\tilde{U}_A$. Also, they may violate the Bill of Rights in the (amended) United States Constitution. Under such circumstances, the forcing contract described above would be unlawful.

Next let us consider the case where the agent is risk-neutral. By this, I mean that the agent is indifferent between receiving a fixed payment (e.g., 5) or an uncertain payment with the same expected value (e.g. probability 0.5 of 0 and probability 0.5 of 10). In this case it can be shown that it is optimal to write a contract such that the agent receives all the benefits from the job (e.g. the sales revenue) less a fixed amount. The principal selects this amount to be as large as possible, while still ensuring that the agent accepts the contract. This incentive scheme induces the agent to select the effort-level that maximizes total revenue and the principal does as well as when effort is not observable. This situation is referred to as the first-best solution. Note that under this scheme the agent carries all the risk: it is basically a franchise agreement. There may again be problems with the implementation of such a contract. For instance, under this payment scheme the agent has no guaranteed minimum wage, which in many countries is unlawful.

In most applications, though, it is more plausible to assume that the agent is risk averse (i.e., he prefers a fixed payment to an uncertain payment with the same expected value) and the principal risk neutral. In this case, from a risk-sharing perspective, it would be optimal to design the incentive contract to be a lump sum, independent of the outcome, causing the risk-neutral principal to carry all the risk while asking the agent to select the optimal level of effort. But this is not implementable as the agent would have no incentive at all to exert any effort. Consequently it is necessary to trade off risk sharing and incentive compatibility, to obtain what is known as a second-best solution.

These concepts can be illustrated using the numerical examples discussed previously. For instance, in the first example the agent receives an expected payment equal to $\frac{5}{8}I(1) + \frac{3}{8}I(0) = 16$, but the utility he derives from this payment equals $\frac{5}{8}\sqrt{I(1)} + \frac{3}{8}\sqrt{I(0)} = 3.5$. A lump sum payment of $(3.5)^2 = 12.25$ would yield the same utility. This latter amount is referred to as the agent's certainty equivalent: he is indifferent between receiving 12.25 for sure and receiving 25 with probability $\frac{5}{8}$ and 1 with probability $\frac{3}{8}$. The difference between the expected payment and the certainty equivalent is known as the risk-premium, 3.75 in this case. It is the additional expected payment the agent requires to compensate for the uncertainty. Note that in this example the agent does carry a considerable amount of risk: if the contract is lost despite a high level of effort, his utility equals $\sqrt{I(0)} - 1 = 1 - 1 = 0$! This is strictly less than his reservation level of 2.5.

In the basic principal/agent model the second best solution exhibits the following properties:
(1) If lump sum payments are possible, the individual rationality constraint is binding (why pay the agent more than he asks for?)
(2) If the decision and outcome space are continuous, then the agent's marginal disutility of effort equals his marginal expected utility of income (this is a classical result of micro-economics).

In the next section I briefly present a series of extensions of the basic model. Section 4 then looks at the application of the basic model and some of its extensions to several functional fields. Two points deserve emphasis at this stage:
(1) The list of examples given is illustrative in nature, and not at all exhaustive.
(2) The focus is on the modelling aspect (what type of questions have been addressed by agency theory) rather than on detailed analysis of the models presented.

3. Extensions of the basic principal / agent model

In this section, I briefly discuss some extensions of the basic principal/agent model: communication, multiple agents, moral hazard and adverse selection, monitoring and multiperiod models.

In many situations the agent is better informed than the principal regarding the uncertain events affecting the outcome. The basic model does not allow the agent to pass on this information to the principal (i.e., no communication). Christensen (1981) analyses an agency problem in which the agent receives information after accepting a contract, but before selecting an action. He can send a message to the principal after receiving this information. Christensen shows that providing more information to the agent does not imply that the agency (agent plus principal) is better off, because the agent could use this information to exploit the principal.

Agency models typically look at the relationship between one agent and one principal. Many real world situations involve the presence of multiple agents, one example being a sales-force consisting of many salesmen. The presence of multiple agents may give rise to phenomena such as collusion, peer pressure and free rider problems, which do not occur in a single agent environment. On the other hand, if several agents are subject to correlated random events, risk-sharing may be improved by linking one agent's payment to the other agent's performance.

Antle (1982) analyses a multi-agent model in the context of auditing. The principal represents the owner. There are two agents: the manager and the auditor. The principal hires the auditor to induce truthful reporting by the agent (e.g. financial reporting). The auditor is modelled as an expected utility maximizer. Consequences from this assumption include the following:
(1) a multi-agent model;
(2) the need to address the subgame between the manager and the auditor;
(3) the need to address the auditor's incentive problems; and
(4) the potential for gains from risk-sharing with the auditor.

The work by Demski and Sappington (1989) relates both to communication and multiple agents. They consider a situation where a principal designs an incentive contract, a worker supplies labour input, and a supervisor can produce information about the production environment which may influence the worker's performance. They show that when the supervisor can acquire information unavailable to the worker, it is optimal for the principal to use the worker's performance to evaluate the supervisor.

In Section 2, I briefly mentioned adverse selection models, where the principal wants to select one out of many potential agents. McAfee and McMillan (1987) analyze a model in which they combine both adverse selection (several candidate agents of varying abilities, known only to the agent), and moral hazard (level of effort not observable). They assume that both the principal and the agents are risk neutral. They show that the principal's optimal contract insures that
(a) the agent with the highest ability will be hired;
(b) the higher an agent's ability, the higher his optimal effort level and the higher his share of the risk.

This result is interesting, as it shows that when both adverse selection and moral hazard problems are present, it is not optimal to have a risk-neutral agent carry all the risk, because extracting information rents becomes an issue.

In the models discussed so far, it was assumed either that the agent's action was not observable, or that it was observable. In many practical situations, the agent's action is not observable as such, but the principal can invest in monitoring, which will yield him (imperfect) information about the agent's action. This information reduces the agency costs, but this gain has to be compared to the cost of monitoring.

An interesting paper by Dyl (1988) looks at the impact of shareholders monitoring the manager's compensation. His hypothesis is that, in closely held companies (i.e., companies whose shares are controlled by a small number of individuals), major stockholders will monitor the manager's behaviour more closely. This will lead to a reduction of agency costs reflected by comparatively lower compensation for the Managing Directors of such companies. Dyl gives empirical evidence that supports this hypothesis. The idea underlying this observation is that major stockholders gain more from a certain amount of monitoring than small shareholders. Consequently, the presence of major shareholders will increase the optimal level of monitoring.

All models presented so far assumed that the principal and the agent interact only once. In many practical situations, this interaction will occur repeatedly. The main difference between one-period and multi-period models is the potential for learning and for reputation building. Consequently, the issues involved are modified considerably, as an agent's behaviour in one period can influence his rewards in future periods. It turns out that in multi-period models, assumptions about who knows what play a crucial role. The slightest modification can lead to completely different outcomes. This raises the issue of the potential advantages of a multi-period contract over a series of one period contracts.

For instance, Antle and Fellingham (1990) considers a situation where a principal hires a manager to implement a project. The manager knows the project's rate of return, while the principal only knows the distribution. This information asymmetry causes inefficiencies, which result in organizational slack and rationing of resources. Antle shows that when the same situation is repeated over time, multi-period contracts can reduce these inefficiencies.

Fudenberg, Hölmstrom and Milgrom (1990) look at sufficient conditions for the implementation of efficient long-term contracts as a sequence of short-term contracts, and conclude that the agent's information advantage plays a key role in determining the value of long-term contracts. In their model, long-term contracts only serve to avoid recontracting under asymmetric information: if at all times of recontracting the principal and the agent share the same beliefs about the pay-off relevant future, long-term contracts provide no gains.

Malcolmson and Spinnewyn (1988) have shown that the presence of moral hazard is not a sufficient condition to insure that a long term contract outperforms a sequence of short term contracts in terms of improving the trade-off between incentives and risk sharing. For models where the interaction between the agent and the principal is repeated infinitely often, Radner (1985) has shown that, provided the discount rate is sufficiently low, the first best solution can be approximated by a sequence of one-period contracts, eliminating the need to consider long term contracts. When the discount rate is higher, it is more difficult to penalise the agent in future periods for bad behaviour now. Consequently, multi-period contracts can yield an improvement over a sequence of one period contracts.

Hölmstrom and Milgrom (1987) look at the problem of providing appropriate incentives over time to an agent who has constant absolute risk aversion (i.e., $-U_A''(\cdot)/U_A'(\cdot)$, where $'$ and $''$ respectively denote first and second derivatives, is constant). They consider both a discrete model, where a single period multi-nomial model (the agent selects the mean of a multi-nomial distribution) is repeated a finite number of times, and a continuous time version, which can be interpreted as a limit case of the discrete model, where the length of a period goes to zero. For the discrete model they show that if the agent observes his performance before going on to the next period, the optimal scheme depends linearly on aggregates of the outcomes of past periods, even though the principal observes the outcomes of each period separately. A similar result holds for the continuous case. Assuming that the agent's action consists of selecting the instantaneous drift of a Brownian motion, it is shown that one can compute the optimal incentive scheme as if the agent were choosing the mean of a normal distribution only once, and the principal was restricted a priori to using a linear rule. The two key conclusions from this work are:

(1) There is no need to use all the available information (aggregate performance over time is sufficient).

(2) Optimal rules in a rich environment must be sufficiently robust to work well in a broad range of circumstances, and therefore should not be complicated functions of the outcome (in this particular model they turn out to be linear).

Other work attempts to capture the dynamic aspects by subdividing the period of interests into several sub-periods, where the agent takes an action and observes the relevant outcome in each sub-period. For instance, Lambert (1984) explains income smoothing in this context. He assumes that the agent exerts some level of effort and observes the resulting income for the first part of the year. He shows that it is rational for both the principal and the agent to design an incentive structure where, if income is above (below) expectations in the first sub-period, it is optimal for the agent to exert a low (high) level of effort in the second part of the year, which leads to income being smoothed towards the ex-ante expected value for the year.

4. Illustrations from various functional fields

In this section I present a selection of applications of agency theory to the fields of Accounting, Industrial Organisation, Marketing and Finance.

The aim is to illustrate the impact the principal/agent paradigm has had in these areas, rather than to provide an exhaustive list of all applications. It is worth pointing out that some applications could fit in more than one field. For instance, I have incorporated sales-force compensation into the marketing section, but one could classify this subject as part of labour relations, and include it under the heading of industrial organization. Similarly, the financial structure of a firm could be included under the heading of industrial organization, rather than finance.

4.1. Accounting

It is in the accounting context that the principal/agent paradigm was first formalized. A very broad spectrum of applications has been developed, and the principal/agent approach has become an integral part of the standard accounting textbook (see for instance Kaplan and Atkinson, 1989). I will only discuss two of them: budgetary control systems and variable cost allocation. The former one was selected on the basis that it is one of the first to be developed, and the latter one because cost allocation is a much discussed topic these days.

Demski and Feltham (1987) define a budget-based contract as having the following properties:
(1) The compensation depends on some observable output-measure.
(2) The budgeted outcome (or standard) partitions the set of feasible outcomes into desirable and undesirable outcomes.
(3) Two compensation functions are defined, one for the set of favourable outcomes and one for the set of unfavourable outcomes.

Two models are considered: the pure moral hazard model, where the agent selects a level of effort and the pure adverse selection model, where skills vary across workers. Demski and Feltham show, among other things, that necessary conditions for a budget based contract to be Pareto superior to other forms of contract are: (1) the agent is risk averse, (2) the state (random element) is costly (or impossible) to observe and (3) effort (moral hazard model) or skills (adverse selection model) are costly (or impossible) to observe.

Gjesdal (1988) considers a model with a risk-neutral agent, where moral hazard and adverse selection are both present. He focuses on piecewise linear compensation functions (i.e., linear functions with a kink at the target outcome). He shows that for a broad range of settings, this simple functional form is effective, and that when efficiency is not obtained, any reasonable function is likely to fail.

As mentioned earlier, empirical work concerning the principal/agent model is rather limited. An exception is Blanchard, Chow and Noreen (1986), who analyze the hospital budgeting system in the state of Washington. In their model, the principal represents the regulatory organization and the hospital manager is the agent. The model shows how the existing budgeting procedure, imposed by the regulatory organization, induces the managers to bias some of the data (estimated variable costs and volume) which they are required to report. They give empirical evidence which supports the conclusions from the model. This study provides a nice illustration of the need to take the agent's incentive constraint into account when determining a budgetary control system, or more generally any control or reward system. Section 5.1 discusses related issues in the context of centralized versus decentralized planning.

Magee (1988) considers a model where the principal (owner) provides a resource (new materials, capital equipment,...) to an agent (worker). The agent selects an effort level (not observable to the principal) and the level of usage of resources (observable to the principal). He shows that if the agent has private information about the value of the resources provided by the principal, then (1) the optimal compensation depends on resource usage; and (2) it is possible that compensation increases as resource usage increases.

The first part of this result simply states that allocation of variable costs is desirable. The second result is more surprising: higher consumption of a resource may lead to higher compensation. This contradicts the traditional cost allocation literature, where higher usage leads to more costs being allocated and possibly penalties. This phenomenon tends to occur when agent effort and resource usage are substitutes. Magee provides a numerical example which illustrates this possibility.

4.2. Industrial organization

This section will be kept very brief, as good surveys of various aspects of the use of the principal/agent paradigm can be found in for instance Hölmstrom and Tirole (1989), Baron (1989) or Tirole (1989). The area of industrial organization provides many examples of moral hazard and adverse selection models. As already mentioned, Akerlof (1970) shows how asymmetric information may lead to a breakdown of the market. Hölmstrom and Tirole survey recent work in industrial organization with an emphasis on moral hazard and incentive problems. The topics discussed include capital structure (which we discuss briefly in Section 4.4), disciplining of the product, labour and capital markets, and the role of supervision.

Baron (1989) provides a survey of recent work in the area of regulation. A regulatory relationship can be interpreted, among others, as a situation where the regulator delegates the pricing decision to the firm, specifying only a menu of pricing mechanisms from which the firm can choose. The firm has private information about its 'type' (production functions, factor prices, managerial ability, etc). The regulator faces the problem of choosing the menu of pricing mechanisms so as to induce maximization of total surplus, while satisfying the firm's incentive compatibility and individual rationality constraints. This situation is similar to that of an employer offering a menu of contracts to potential salesmen, who are better informed about their own ability, as discussed for instance in Lal and Staelin (1987). We discuss this issue in more detail in Section 4.3.

Firms are sometimes observed to be taking seemingly inefficient actions: i.e., they do not maximize profits. As explained in Tirole (1989), this seemingly sub-optimal behaviour may be the consequence of constraints imposed by the inability to observe the actions of all the agents. Also, managers have the power to exploit the information advantage they have over shareholders. This ability is to some extent limited by market constraints: sub-optimal decisions could lead to take-over bids, or even bankruptcy, which are harmful to the manager.

The manager also faces the dilemma of furthering his own career at the expense of the firm: decisions maximizing short term profits may provide him with a larger bonus and attractive job offers, at the expense of the future profitability of the firm. This problem will be discussed further when discussing work by Hölmstrom and Ricart i Costa (1986) in Section 4.4.

A related issue is that of labour contracts. Epstein (1985) notes that 60–65% of labour contracts in the USA are 'at will'; i.e., both parties can terminate the contract if they so desire. He argues that this type of contracts provides a high level of flexibility in adjusting the imperfections of the original contract, as neither party has an incentive to stick rigidly to the original contract. Two other advantages are that neither party has an incentive to use litigation to settle disagreements, and the loss of severing the relationship are larger to the party which misbehaves. But this type of contracts is not applicable in a unionized environment. A sufficient level of stability must be achieved, which cannot be guaranteed by 'at will' contracts. On the other hand, if firing is impossible, some agents can free-ride on the other agents' efforts. Epstein suggests that the 'for-cause' contracts (the contract can only be terminated if there is a good cause to do so) are an attempt at balancing these two elements, but argues that this type of contracts increases agency costs. The issue of rewarding a team of agents is closely related: if the effort of individual members cannot be observed, how can one avoid the free rider-problem?

Another area where agency problems arise is that of limited partnerships. Wolfson (1985) studies these in the context of the oil and gas industry, and tests several hypotheses empirically. Typically, the limited partner (the principal) provides the capital, while the general partner (the agent) makes the decisions. Agency problems arise because

(1) The agent has much more information than the principal.

(2) Cost allocations and tax effects imply that decisions beneficial to the agent may be sub-optimal for the principal.

Therefore limited partnerships are expected to occur less commonly in instances where these problems are more severe (developmental drilling). Agents can decrease these problems by building a reputation for success at the cost of lower short term profits. Doing so should allow them to obtain investment funds on more favourable terms. Wolfson shows that both these hypotheses are supported by empirical data.

4.3. Marketing

The relationship between a sales-manager and a salesmen is a natural candidate for a principal-agent approach. I discuss work that addresses three topics:
(1) What are the characteristics of an optimal compensation plan?
(2) When should a menu of compensation plans be offered?
(3) When should pricing be delegated to the salesmen?

Basu et al. (1985) model the sales-force compensation problem as a principal/agent relationship. The principal is the sales manager and the agent is the salesman. The principal's goal is to select a compensation plan $I(\cdot)$ for his sales-force, so as to maximize the expected value of revenue minus cost. Revenue is determined by the level of sales (x; the outcome). Costs consist of (variable) production costs and the compensation paid to the sales-force. In selecting a compensation plan the principal is constrained by the usual two conditions: the agent must be willing to accept the suggested contract and it must be optimal for the agent to select the effort-level desired by the principal.

Given a compensation scheme, the agent selects his effort-level (e in this case is time input) so as to maximize the expected utility of the compensation he receives, minus the cost of effort. He will only accept the contract if this makes him better off than rejecting it. The actual sales level depends on two elements:
(1) the agent's time-input, and
(2) random factors not observable to the principal.

Basu et al. show that, under a set of assumptions similar to the traditional set of assumptions in this literature (i.e., separability of utility in income and effort, decreasing marginal utility of income, increasing marginal disutility of effort, etc.), the characteristics of the optimal payment scheme depend on the agent's risk tolerance, where risk tolerance is defined as the inverse of Pratt's measure of absolute risk aversion (i.e., $T(\cdot) = -U_A'(\cdot)/U_A''(\cdot)$, where $'$ and $''$ respectively denote first and second derivatives).

Lal and Staelin (1986) address the following question: Why would a principal want to offer a menu of compensation plans to potential agents? They consider a situation where the sales manager would like to hire a subset of a pool of agents. Each agent is either a high performer or a low performer. Agents know their type, but the principal does not. He only has a prior over the number of high (low) performers in the pool of applicants. This set-up is similar to the work by McAfee and McMillan (1987) mentioned earlier, who show that when the principal must select one out of many potential agents whose abilities are unknown to him, the optimal contract insures that the highest ability agent is selected.

The main results can be summarized as follows:
(a) Based on his prior assessment over the number of high (low) performers in the pool of applicants and their reservation levels, the principal decides whether he wants to hire either only high performers or only low performers or both groups of agents.
(b) If the principal decides to hire agents from one group only, he offers a single contract which only agents from the desired group will accept.
(c) If the principal decides to hire agents from both groups, he offers a pair of contracts such that high performers will select one contract and low performers the other.

In other words, based on his decision of whom to hire, the principal designs contracts to induce self selection by the agents. Lal and Staelin also mention that the same result holds if the agents are characterized by either high or low risk-tolerance rather than performance.

Lal (1986) analyzes conditions under which it is optimal for the sales manager to delegate the pricing decision to the sales-force. In this situation the principal can select a compensation plan which depends on both the sales level and the selected price. The agent selects both an effort-level and a pricing policy. The actual price charged by the agent depends on the information revealed by his effort. In other words, the agent is able to expend effort to gain information about his market and/or customers, and then select a price based on this information.

The main results can be summarized as follows:
(a) Delegating the pricing decision does not make the principal worse off.
(b) Information asymmetry is a necessary (but not sufficient) condition for price delegation to be strictly better; i.e., the agent must be able to obtain more information on the sales environment than the principal.

The intuition behind the second part of the result is that the agent is able to exploit his information to charge a different price depending on the prevailing situation. Lal also notes that in practice, price delegation is observed in environments where the sales manager is significantly removed from the scene of action; i.e., where there is indeed a considerable likelihood that the agent is better informed than the principal.

Coughlan and Sen (1986) give an overview of other work in the area of sales-force compensation. They state a number of hypotheses derived from theoretical work and look at empirical evidence which either confirms or contradicts these hypotheses. As stated in their introduction, their results are more suggestive than conclusive, but represent a first attempt at testing these hypotheses.

Other aspects of sales force compensation may also benefit from this approach. For instance, how does one induce a salesman to allocate his time optimally between various products, or between following up on existing customers and looking for new customers?

4.4. Finance

The principal/agent approach has been used extensively in the financial modelling literature. In most work, the principal represents the shareholders, investors or owners of the company, while the agent represents the manager. Kaplan (1984) has expressed concern about this approach, on the grounds that characterizing managers as being effort averse is not very realistic. This point will be picked up in the concluding section.

The traditional analysis of the capital structure of the firm assumes an environment of full information. It has been shown that, under this assumption, the financial structure of the firm is irrelevant, which makes one wonder why we observe such a variety of complex financial structures. Ross (1977) takes an alternative approach. He assumes that managers have inside information about the value of their firm; i.e., there is asymmetric information. He suggests a managerial compensation scheme which has the property that the financial structure (amount of debt) selected by the manager reveals his inside information. Note that in this approach there is no 'effort' involved: the agent's only action consists of choosing a financial structure.

This compensation scheme is structured as follows:
(a) A bankruptcy cost is imposed on the manager (there is no need for the firm to actually incur this cost in case of bankruptcy).
(b) The manager receives a payment which is a function of the value of the firm.

Based on this scheme, the manager will select the smallest amount of debt (minimize the probability of bankruptcy) that differentiates his firm from the less valuable ones (increase his payment). A consequence of this scheme is that the debt-equity ratio rises as the value of a firm increases.

Haugen and Senbet (1981) consider a situation where the manager owns the firm; i.e., he retains complete control, but issues debt to outsiders. In their model the agent is the manager-owner and the principal represents the bondholders. They argue that the agent has the possibility (and the incentive) to increase his wealth at the expense of the bondholders by
(a) overconsuming perks (slack management), and
(b) pursuing riskier investment strategies than bondholders or a sole owner like.

They suggest that it is possible to reduce, and in some cases eliminate, these problems by requiring the manager-owner to own a combination of call and put options. The scheme works as follows:
(a) The manager owns a call option to buy back the whole firm at a set price.
(b) The bondholders own a put option to sell the firm to the manager.

It is intuitively clear that each of these make the agent bear the full cost of perk consumption, thus reducing his consumption to the level of a sole owner.

Haugen and Senbet also argue that the combination of these call and put options may reduce the agent's incentive to select risky projects. The argument is that selecting more risky projects will increase the value of the agent's call option as well as the value of the principal's put option. By selecting appropriate exercise prices, these two effects cancel each other out as far as the agent's portfolio is concerned, thus eliminating the incentive to select riskier projects. One problem with this approach is that the appropriate exercise prices may not exist.

Lambert (1986) analyses the selection of risky projects, focusing on an executive's motivation to expend effort on obtaining information that will allow him to select the best project. He shows that the conflict of interests between the principal and the executive can lead to under or over investment in risky projects, and indicates how in a multi-period setting, the use of future-period incentives is required to motivate the executive to spend effort on both productive use and information gathering.

Thompson and Wright (1989) have interpreted the recent surge in management buyouts as a way to reduce the consumption of managerial perks. These buy-outs usually consist of a group of managers receiving a large proportion of the shares for a relatively small investment, the balance being financed by debt. This transaction results in a reunification of ownership and control, thereby reducing (or eliminating) the agency problem. Using a simple model, Thompson and Wright show how this type of buy-out can reduce the consumption of managerial perks, thereby raising the value of the firm.

Bhattacharya and Pfleiderer (1985) consider a situation where a principal desires to
(1) hire agents (portfolio-managers) from a pool of candidates with heterogenous abilities, and
(2) motivate the hired agents to act in his interest.

Note that this paper combines the problems of screening and delegating. The authors consider an environment of asymmetric information where agents have two elements of private information, one of which is received at the pre-contracting stage (e.g. forecasting ability) and one is received after contracting, but before selecting an action (e.g. about security return).

Both the action of the agent and the outcome are assumed to be observable. The authors show that despite this, it remains necessary to deviate from an optimal risksharing agreement, in order to induce the agent to reveal his information (screening) at the pre-contracting stage. In other words, the presence of asymmetric information before contracting prevents an optimal risk-sharing agreement.

Hölmstrom and Ricart i Costa (1986) suggest a different reason for the divergence of interests between the manager (agent) and the owner (principal). They describe the cause of the incentive problem as a different valuation of the returns of a project: the manager worries about 'human capital' returns (career concerns) while the owner is interested in financial returns. This divergence leads to a difference in risk preferences.

They present a two-period model which allows for learning about managerial talent, the agent's second period wage being a function of his performance in the first period. Each period the agent has the possibility to recommend a project to the manager, based on his private information about potential projects. This information will not be known to the principal until after an investment takes place.

This information structure implies that the agent can refrain from presenting a specific project, but cannot misrepresent his information. The main results can be summarized as follows:
(a) The second best contract consists of giving the agent a downward rigid wage (partial insurance), which can be interpreted as an option on the value of his human capital.
(b) If the manager has a sufficiently low risk aversion, the principal must ration capital (by an investment rule) to avoid over-investment.
(c) Under the second best contract less projects are undertaken than in the first best (full insurance) case: the effective cost of capital exceeds the market rate, because of the residual human capital risk carried by the option (see (a)).

5. Principal / agent theory and Management Science

Very few applications of agency theory to Management Science (MS) problems can be found in the literature, although many might benefit from this approach. In this section, I first discuss one of these applications: the choice of centralized versus decentralized production planning. It is worth noting that the work discussed in Section

5.1 was done before agency theory became an accepted paradigm. Also, the authors assume that both parties are risk neutral, but do not consider the possibility of letting the agent carry all the risk. In Section 5.2, I discuss potential applications of agency theory to MS.

5.1. Centralized versus decentralized production planning

In this section we look at several bonus schemes a central planner may select to induce plant managers (or managers of any type of resource) to produce the desired amounts of output. In principal/agent terminology, the principal is the central planner and the agent is the plant manager. It is generally assumed that the agent has more information about the capacity of the plant than the principal. In a classical reward system, the central planner sets a target for the plant, and the manager receives a bonus if this target is achieved or exceeded. Such a scheme is not without problems, especially if one assumes that the manager knows more about capacity than the central planner.

One category of problems relates to how the target level is determined. If the central planner uses the manager's estimate of capacity to set the target, then the manager has an incentive to underestimate capacity, as this will make his life easier. On the other hand, if the target is a function of past performance, the agent has a strong incentive to underproduce, so as to keep his target down in future periods. The principle of giving a bonus for exceeding the target is in itself problematic, especially if the target is not determined appropriately. The concept of inducing managers to overachieve contradicts the initial purpose of production planning.

Several incentive schemes, addressing one or more of these problems have been developed. Ijiri, Kinard and Putney (1968) provide a detailed discussion of the problem of setting an appropriate target level. They use simple budgeting examples to show how target level, actual level and performance interact, given a system of penalties. They suggest how this approach can be applied to responsibility accounting by having each responsibility centre submit a plan, which form the basis for central planning. Performance of each responsibility centre is evaluated using the usual accounting methods, as well as penalties for discrepancies between plans and actual figures. The choice of appropriate penalties is important if the plan of one responsibility centre affects performance at other centres.

Fan (1975) suggests an incentive scheme in the spirit of the one suggested by Ijiri, Kinard and Putney (1968), which induces the manager to reveal correctly the (expected) capacity and, given any target, to produce up to the (realized) capacity. Bonin (1976) extends Fan's result and shows that for any distribution of the capacity, it is optimal for the manager to reveal the median capacity (rather than the expected capacity, which is a consequence of the symmetry of the bonus scheme around the production target). Bonin (1976) also takes Fan's analysis a step further. He introduces the concept of 'planning from below'. Under this scenario, the planner selects an incentive scheme, and the agent selects his own target; i.e., there is no longer a need for the planner to obtain information from the manager. He suggests a bonus scheme which allows the planner to give the manager an incentive to reveal a specific fractile of the output distribution, and to produce up to the realized capacity.

Weitzman (1977) discusses the properties of an incentive scheme that was first introduced in the Soviet Union in 1971. The main characteristic of this plan is that the target results from a two-stage process. The planner first selects a tentative bonus and a tentative target. The manager is then allowed to revise the target. The actual bonus depends on these three quantities as well as on the realized production. The properties of this scheme are similar to those of the scheme suggested by Bonin.

This system does require the central planner to gain some information about capacity to be able to set a reasonable tentative target. The manager's opportunity to revise this target makes an accurate decision less crucial and also reduces the problem ensuing from the relationship between target and either past performance or information revealed by the manager.

It is interesting to relate this to Gjesdal's (1988) work mentioned earlier. The scheme described above can be interpreted as a piece-wise linear system, with a slight complexity introduced by the special way in which the target bonus is determined. The situation considered here is simpler as there is no adverse selection issue.

5.2. Potential applications of principal / agency theory to Management Science

In this section, I present three management problems which involve issues of task delegation and which may therefore benefit from a principal/ agent approach. The first two are only briefly discussed. The third one is described in more detail to illustrate what kind of insights can be expected from this approach.

The first example considers issues of production planning, and has already been discussed to some extent in the previous section. Elements one may want to consider include: should goals be set by one central authority (the principal) who tries to enforce them, or should this decision be delegated to the production unit (the agent); what type of payment scheme will induce achievement of a stated goal, thus avoiding excessive over- or underproduction; what is the impact of the payment scheme on the type of individuals who are willing to run a production unit (self selection). These concepts seem relevant in the present crisis of the EEC's agricultural policy. The incentive system created to insure an adequate minimum revenue has led to large amounts of overproduction and to abuse of the system. How could this system be revised to keep the positive aspect of a guaranteed minimum income, while avoiding the undesirable side-effects?

The issue of production planning could be extended to incorporate aspects of quality: it is not solely the amounts produced that matters, but also the quality of this output and the amount of rejected output. This latter element is especially important as it influences production requirements at the upstream units. If planned reject rates are exceeded, the overall production plan may be severely affected.

The second example concerns scheduling. Consider a situation where there is one facility, say a piece of specialized equipment in a lab, and several engineers (agents) desire to use this piece of equipment to carry out a test. They submit their requests for a specific amount of time (which may be random) to the lab manager (the principal), who schedules a starting time for each.

Given that times required to perform a test are subject to uncertainty, the manager does not know when exactly the equipment will be available. If he could perform the tests himself, his job would be easy. He could collect the items to be tested, carry out all the tests sequentially, and return the items with the test results. If this is not possible, he must select a schedule for the engineers. If he schedules the tests close to each other, engineers may have to wait (the equipment is not available when they arrive). If this occurs frequently, they may take on the habit of arriving late.

On the other hand, if the manager schedules too much buffer time between tests, the equipment will be idle for a non-negligible part of the day, implying low utilization and possibly long waiting lists to obtain access to the equipment. Consequently, when selecting a schedule, the lab manager must trade off the engineers' waiting time against idleness of the equipment, and also ensure that the engineers will arrive at their scheduled starting times. Although situated in a very different environment, the structure of this problem is very similar to the medical practice example mentioned in the introduction.

Using a principal/ agent approach would allow one to incorporate some of the following aspects: asymmetric information (the engineers know their disutility of having to wait, they have a better idea of how long the test will take, etc); self selection (a specific scheduling rule may induce some engineers to look for another lab, e.g. engineers with a high disutility for waiting in the case of a tight scheduling rule); delegation (should the lab-manager select a schedule and attempt to implement it, or should he limit his intervention to determining the order of tests (e.g., first applied, first served) and let the engineers decide when to arrive; moral hazard (did an engineer come late on purpose to avoid waiting or because he was unexpectedly held up by a phone call?).

The key issue of this example is that in the principal/ agent context (or more generally, in the game theoretic context), scheduling can be analyzed as a multi-person problem, rather than as a situation where one individual selects a schedule and the others obey blindly. A similar problem is analyzed in van Ackere (1990) in the context of hospital operating rooms. There is a single operating room, in which several operations, each performed by a different surgeon, need to be scheduled sequentially. The duration of the various surgical procedures is subject to uncertainty and surgeons may arrive late, either on purpose or due to events beyond their control. The surgeons' behaviour is analyzed, and an optimal scheduling rule which takes this behaviour into account is derived.

The third example considers the selection of batch-size in an environment with high set-up costs. The purpose is to produce enough good units to satisfy a specific order. The reject rate depends on the worker's level of effort, and on a number of random elements (e.g. quality of the raw materials, precision of the set-up, etc.). In such a situation, should the batch-size be determined by the supervisor (principal) or should he delegate this decision to the worker (agent) who knows what level of effort he will select and who may have better information about the random elements? Under each of these alternatives, what is the optimal reward system? How can the principal induce the optimal batchsize/effort level combination? Does an appropriate reward system allow the supervisor to obtain the 'right' type of workers for the job (self-selection)?

In the remainder of this section I present a simple model that enables us to illustrate some of these points. Consider a manager who has received a special order for a limited amount of a product. Examples include shirts with a special logo for a sports club, textile of a very specific colouring, and packaging materials for a special deal. Any excess items must be disposed off, while a shortage results in penalties. The manager must select an optimal run-length (batch-size), knowing that a second run is infeasible, either due to excessive set-up costs, time constraints or technical infeasibility (e.g. obtaining a second batch of identical colouring).

For mathematical simplicity I assume that the order is of size 1, and infinitely divisible. The batch-size is denoted by B. The yield depends on the worker's level of effort, denoted by e, where $e \in [e^-, e^+]$, $e^- \geq 0$, $e^+ \leq 1$, and on a random component r, $0 \leq r \leq 1$, with density $f(\cdot)$ and cumulative distribution $F(\cdot)$. For a given batch-size, level of effort and random factor, the number of good units is assumed to equal erB. Units up to the ordered quantity can be sold at a price P, while any excess units are valueless. If insufficient good units are available, a shortage cost of S per unit of shortage is incurred. This cost can be interpreted either as a penalty for non-delivery or as the cost of reworking a defective unit.

Let us assume that a worker can produce one unit per hour, independently of his level of effort, and that he is paid a wage w. Variable material costs equal c per unit. First it is assumed that there are no incentive compatibility problems and no individual rationality issues. Next the individual rationality issue is included, and finally incentive compatibility problems are considered, assuming that the worker is risk averse. A numerical example is included to illustrate how accounting for individual rationality and incentive compatibility affects the optimal decision.

Given this framework, the manager's profit function for a given batch-size, effort level and random shock equals

$$\Pi(B, e, r) = P * \min(1, erB) - S * \max(0, 1 - erB) - (c + w) * B. \quad (1)$$

The first term refers to sales revenue: good units up to the order quantity of 1 can be sold at price P. The second term represents the shortage cost: If $erB < 1$, a shortage cost equal to $S(1 - erB)$ is incurred. The last term consists of the variable production and labour costs. Taking the expected value with respect to the random factor r yields an expected profit of

$$\Pi(B, e) = (P + S) \int_0^{1/(eB)} erBf(r)\, \mathrm{d}r + (1 - F(1/eB))P - SF(1/eB) - (c + w)B. \quad (2)$$

To keep the problem interesting, the following assumption is made:

$$E(R)e^- \geq (c + w)/(P + S).$$

This assumption insures that for any effort level e, it is optimal to select a batch-size B such that $eB \geq 1$; i.e., under ideal circumstances ($r = 1$) the manager would be able to satisfy the order. Optimizing (2) with respect to the batch-size, for a given level of effort, yields

$$\int_0^{1/(eB^*(e))} erf(r)\, \mathrm{d}r = (c + w)/(P + S). \quad (3)$$

This expression shows that the optimal batch size depends both on the effort level and on the distribution of the random factor. Intuitively, one would expect the batch-size $B^*(e)$ to be decreasing in the level of effort e. This turns out to be the case for most distributions of r, but not for all. From (3) it follows that

$$\frac{\mathrm{d}B^*(e)}{\mathrm{d}e}=\frac{eB^3\int_0^{1/(eB)} rf(r)\,\mathrm{d}r}{f(1/(eB))}-\frac{B}{e}. \qquad (4)$$

If the random factor r is uniform on [0, 1], (4) yields

$$\frac{\mathrm{d}B^*(e)}{\mathrm{d}e}=\frac{-B^*(e)}{2e}<0,$$

which implies that the optimal batch-size is indeed decreasing in the level of effort e. If the random component r has density $f(r)=2(1-r)$, i.e., a high probability of a low yield, then

$$\frac{\mathrm{d}B^*(e)}{\mathrm{d}e}=\frac{4-3eB^*(e)}{6e^2(1-eB^*(e))},$$

which is negative for $eB^*(e)<\frac{3}{4}$, and positive otherwise. This phenomenon can be explained intuitively as follows: when the yield is very low, additional good units are likely to be needed and are therefore valuable. At the same time, an increase in the level of effort makes these units less costly to produce. Consequently, the desire for more good units leads to a larger optimal batch-size, despite the increase in yield.

As this phenomenon is the exception rather than the rule, I will focus on the case where $B^*(e)$ decreases in e. In the remainder of this section, I will assume more specifically that r is distributed uniformly on [0, 1]. In this case

$$\Pi(B,e)=P-\frac{P+S}{2eB}-(c+w)B, \qquad (5)$$

which yields

$$B^*(e)=\sqrt{\frac{P+S}{2(c+w)e}} \quad \text{and}$$

$$\Pi(B^*(e),e)=P-\sqrt{2(c+w)(P+S)/e}. \qquad (6)$$

Given that the expected profit function is increasing in e, and that the manager can costlessly induce any level of effort, he will select $e=e^+$.

Next let us assume that the worker's utility, when receiving a wage w and selecting a level of effort e, equals $w-e^2$. A quadratic term is used for the cost of effort to capture decreasing marginal returns of effort: increasing the fraction of good units from $0.1r$ to $0.2r$ requires less additional effort than increasing it from say $0.6r$ to $0.7r$. If the agent's reservation utility equals $\tilde{U}_A$, the manager must pay him $w=\tilde{U}_A+e^2$ when

Table 2
Numerical illustration of the batchsize example. Parameters: $P=16$, $S=4$, $0\le e\le 1$, $r=U[0,1]$

c	$\tilde{U}_A$	Risk neutrality					Risk aversion											
							Individual rationality					Incentive compatibility						
		e	B	$eB/2$	Wage	Profit	e	B	$eB/2$	Wage	Profit	w_d	w_g	e	B	$eB/2$	Wage	Profit
0	0.2	0.45	7.48	1.7	0.4	10.0	0.26	23.34	3.01	0.07	12.68	0.03	0.53	0.14	34.92	2.39	0.07	11.57
0	0.5	0.71	3.76	1.3	1.0	8.5	0.41	7.42	1.52	0.44	9.40	0.19	2.11	0.25	10.21	1.30	0.43	7.72
0	0.9	0.95	2.42	1.1	1.8	7.3	0.55	3.56	0.98	1.44	5.75	0.57	5.17	0.38	4.63	0.88	1.44	3.62
0.5	0.2	0.84	2.92	1.2	0.9	7.8	0.60	4.51	1.36	0.32	8.64	0.01	1.99	0.33	6.09	1.01	0.34	5.94
0.5	0.5	1	2.24	1.1	1.5	7.1	0.60	3.66	1.10	0.74	6.91	0.13	3.38	0.37	4.91	0.90	0.73	4.43
0.5	0.9	1	2.04	1.0	1.9	6.3	0.65	2.61	0.85	1.76	4.24	0.49	6.19	0.45	3.35	0.75	1.77	1.73
1	0.2	1	2.13	1.1	1.2	6.6	0.73	3.01	1.09	0.53	6.82	0.00	2.83	0.41	3.92	0.81	0.59	3.61
1	0.5	1	2.00	1.0	1.5	6.1	0.71	2.66	0.94	1.00	5.36	0.09	4.28	0.44	3.46	0.76	1.01	2.47
1	0.9	1	1.86	0.9	1.9	5.5	0.73	2.12	0.77	2.05	3.07	0.42	7.02	0.50	2.70	0.67	2.07	0.30

requiring a level of effort e. His problem becomes

$$\Pi(B, e) = (P+S)\int_0^{1/(eB)} erBf(r)\,\mathrm{d}r + (1 - F(1/eB))P - SF(1/(eB)) - (c + \tilde{U}_A + e^2)B, \quad (7)$$

which follows from (2) when replacing w by $\tilde{U}_A + e^2$. The first order conditions become

$$\int_0^{1/(eB)} erf(r)\,\mathrm{d}r = \frac{c + \tilde{U}_A + e^2}{P+S},$$
$$\int_0^{1/(eB)} rBf(r)\,\mathrm{d}r = \frac{2eB}{P+S}. \quad (8)$$

Again, $B^*(e)$ will be mostly decreasing in e, except in a few special cases. If r is distributed uniformly on [0, 1], (7) reduces to

$$\Pi(B, e) = P - \frac{P+S}{2eB} - (c + \tilde{U}_A + e^2)B. \quad (9)$$

This yields

$$B^*(e) = \sqrt{\frac{P+S}{2(c+\tilde{U}_A)^{3/2}}} \quad \text{and}$$

$$e^* = \sqrt{c + \tilde{U}_A} \quad \text{for } \sqrt{c + \tilde{U}_A} \in [e^-, e^+]. \quad (10)$$

The optimal level of effort e^* increases in both c and $\tilde{U}_A$, while the optimal batch size decreases in c and $\tilde{U}_A$. Expected profits equal

$$\Pi(B^*(e), e^*) = P - 2\sqrt{(P+S)\sqrt{c + \tilde{U}_A}}$$
$$\text{for } \sqrt{c + \tilde{U}_A} \in [e^-, e^+],$$

and decrease in c and $\tilde{U}_A$, as expected. A numerical example is given in Table 2.

Comparing this to the case where the principal did not account for the cost of the agent's effort shows that the optimal effort-level is lower for $c + \tilde{U}_A < (e^+)^2$, and unchanged otherwise.

So far the worker was assumed to be risk-neutral. Next, risk aversion is introduced by assuming that his utility equals the square root of any payments received, minus the disutility of effort: i.e., $U_A(w, e) = \sqrt{w} - e^2$. If there are no incentive compatibility problems, the manager must pay the worker $(\tilde{U}_A + e^2)^2$ when requiring a level of effort e. Assuming again that r is distributed uniformly on [0, 1], his problem becomes

$$\Pi(B, e) = P - \frac{P+S}{2eB} - \left(c + (\tilde{U}_A + e^2)^2\right)B, \quad (11)$$

which follows from (5) by replacing w by $(\tilde{U}_A + e^2)^2$. Solving this problem for B and e yields

$$B^*(e) = \sqrt{\frac{P+S}{2\left(c + (\tilde{U}_A + e^2)^2\right)e}} \quad \text{and}$$

$$e^* = \sqrt{\tfrac{1}{3}\left(\sqrt{4\tilde{U}_A^2 + 3c} - \tilde{U}_A\right)}. \quad (12)$$

Note that e^* is increasing in c, increasing in $\tilde{U}_A$ for $\tilde{U}_A > \frac{1}{2}\sqrt{c}$, but decreasing in $\tilde{U}_A$ for $\tilde{U}_A < \frac{1}{2}\sqrt{c}$. This is counter-intuitive, as one would expect e^* to increase in $\tilde{U}_A$ as defects become more expensive. Intuitively this can be explained as follows. When c is 'large', the desired level of effort is high. Therefore, an increase in $\tilde{U}_A$ will cause a considerable increase in the required wage $(\tilde{U}_A + e^2)^2$ to maintain the same level of effort. This creates an incentive to decrease the level of effort. When $\tilde{U}_A$ is small compared to c ($\tilde{U}_A < \frac{1}{2}\sqrt{c}$), this latter effect dominates.

This case is illustrated in Table 2. Note that the profits should not be compared to those of the risk neutral case, due to the difference in the utility function of the agent.

Next, let us incorporate incentive compatibility and assume that a worker attempts to maximize his hourly expected utility. The wage he receives depends on the observed outcome: he is paid w_g if the unit is good, and w_d if the unit is defective. Assuming r is uniform on [0, 1], the agent's problem thus becomes

$$\max_e \tfrac{1}{2}e\sqrt{w_g} + (1 - \tfrac{1}{2}e)\sqrt{w_d} - e^2.$$

This yields $e = \frac{1}{4}(\sqrt{w_g} - \sqrt{w_d})$. The principal's problem can then be formulated as

$$\max_{e,B,w_g,w_d} P - \frac{P+S}{2eB} - cB - \left\{\tfrac{1}{2}ew_g + (1 - \tfrac{1}{2}e)w_d\right\}B$$

$$\text{subject to} \quad \tfrac{1}{2}e\sqrt{w_g} + (1 - \tfrac{1}{2}e)\sqrt{w_d} - e^2 \geq \tilde{U}_A,$$
$$e = \tfrac{1}{4}\left(\sqrt{w_g} - \sqrt{w_d}\right).$$

Solving for the first order conditions and rearranging the terms yields the following equations:

$$e = \sqrt{\tilde{U}_A - \sqrt{w_d}}\,,$$

$$\sqrt{w_g} = \sqrt{w_d} + 4\sqrt{\tilde{U}_A - \sqrt{w_d}} = \sqrt{w_d} + 4e,$$

$$B = \sqrt{\frac{P+S}{48\left(\tilde{U}_A - \sqrt{w_d}\right)^2 - 8\left(\tilde{U}_A - \sqrt{w_d}\right)^{5/2}}}\,,$$

$$w_d = c + 2\tilde{U}_A\sqrt{w_d} + 4\left(\tilde{U}_A - \sqrt{w_d}\right)^2 - 16\left(\tilde{U}_A - \sqrt{w_d}\right)^{3/2}.$$

Table 2 illustrates these results for a numerical example. As expected, $w_d < w_g$, e^* increases in c and $\tilde{U}_A$, and B^* decreases in c and $\tilde{U}_A$. Both components of the wage increase in $\tilde{U}_A$, and w_g increases in c while w_d decreases. The optimal level of effort is considerably lower than in the case without the incentive compatibility constraint, the expected number of good units is lower, despite the considerably larger optimal batch-size, and expected profits are lower.

This simple example illustrates how incorporating the agent's ability to select the effort-level maximizing his utility can considerably affect the optimal policy: the optimal effort level may be reduced by as much as 50%, while the optimal batch-size may be increased by up to 50%. The drop in profits illustrates that it may be worthwhile investigating the possibility of monitoring the agent's level of effort, for instance by collecting information about factors influencing r.

This example illustrates that the optimal planning and control policy is context specific. In particular, it is not possible to separate the planning stage (what is the optimal batch size, what is the optimal level of effort) from the control stage (how does one induce the worker to exert that level of effort). Rather, the optimal plan will reflect the control costs associated with implementing the plan.

This example only scratches the surface of the issues involved. Other elements, such as information asymmetry (the agent has private information about r) could be incorporated. It was assumed that the agent maximizes his expected hourly utility. One might want to analyze what happens if he maximizes expected utility from the job: he may want to decrease the level of effort in order to increase the optimal batch-size, and therefore be employed a larger number of hours.

The purpose of these examples was to illustrate what type of results can be expected. The common point among these three examples is that a principal/agent approach (and more generally, the use of a multi-person approach) allows one to address some basic questions which are often overlooked by other methods of analysis, and see how sensitive optimal policies are to these issues.

6. Concluding remarks

Many criticisms have been directed towards agency theory. Some of these, but not all, can be answered by the extensions discussed in Section 3. I first discuss some of the more important ones, and end with some general comments about the principal/agent approach.

Kaplan (1984) mentions, among others, the following problems with agency theory:
(1) only simple models can be solved.
(2) Agency theory is a theory of contracting with production workers, not managers.
(3) The role of knowledge and innovation to create value in the firm is missing; i.e., agency theory is static.

The first of these criticisms is indeed a key issue: as soon as one goes beyond the basic set-up, agency models become quite complex to solve, and the complicated algebraic expressions hide the underlying intuition. Still, Grossman and Hart (1983) have shown that one can obtain considerable insight into the optimal structure of an incentive scheme without working out the full problem. They explain how the complexity of solving an agency problem can be reduced by splitting it up into two sub-problems if one is willing to assume that the agent's preferences over income lotteries are independent of the action he selects. In a first stage, one determines for each action an incentive scheme that minimizes the expected cost of inducing the agent to choose that action. This step can be formulated as a convex program. In the second stage the optimal action is determined. This stage is much harder to solve. Grossman and Hart show that solving for stage one only, allows to determine many characteristics of the optimal incentive contract. Gjesdal's (1988) work on the optimality of piece-wise linear incentive schemes is also relevant in this respect, and when turning to multi-period models, the work by Hölmstrom and Milgrom (1987) mentioned earlier is encouraging.

The second criticism has been partly addressed by reinterpreting the basic model. Rather than talking about 'effort', alternative interpretations look at career concerns, signalling, perks consumption, etc. Alternatively, one can interpret the principal/agent approach as a theory of 'life at the margin': assuming 'disutility of effort' to be U-shaped, we would be at a point of positively sloping personal cost. Still, this criticism is worth keeping in mind.

The third criticism is probably the most crucial one. In a rapidly changing world, the manager's role is to ensure that his firm is ready for tomorrow's requirements. To capture these elements a dynamic multi-period environment is desirable. But no theory can capture all aspects of a problem. Other fields have (and are) focusing on issues such as selecting an appropriate R&D strategy, choosing the right technology, or reacting to a competitor's action. Game theory has been used to address some of these issues. But there again, the problem of tractability arises as soon as attempts are made to make a model more realistic. This raises the question as to whether precise mathematical analysis is the appropriate tool to tackle a dynamic environment, but does not imply that agency theory should be put 'on the shelf' for problems involving dynamics. The insights derived from analyzing simple problems can be very useful when using other tools to handle the more complex framework. Again, the work by Hölmstrom and Milgrom (1987) is a first step in the process of capturing the multi-period aspect in a tractable way.

Demski and Kreps (1980) criticize one of the fundamental assumptions underlying agency theory: no recontracting is allowed after an action is taken (or the effort-level selected) but before the outcome is known, although in many instances this would make both parties better off. The key problem here is that, if this type of recontracting were allowed, then the original contract would no longer be an equilibrium, as the principal and the agent would expect recontracting to occur and adapt their behaviour accordingly. This observation emphasizes the need for further attention to issues of renegotiation and renegotiation-proof equilibria, as for instance in Fudenberg and Tirole (1990). They modify the standard principal/agent model in order to capture the notion of renegotiation. They show that any outcome that the principal can achieve by renegotiation after the agent selects his action, but before the outcome is observed, can also be achieved by a renegotiation-proof contract, and that the resulting equilibrium involves mixed strategies. The work by Fudenberg, Hölmstrom and Tirole (1990) referred to earlier, which looks at the conditions under which efficient long term contracts can be implemented as a sequence of short term contracts is also important in this regard.

Some aspects of agency theory appeared in several of the functional fields discussed above. Examples include the idea of delegation (portfolio management, pricing, resource usage), standards (accounting and POM (bonus schemes)), and self-selection (portfolio management and sales force selection). Concepts and results developed in one area can be of significant use in other fields, making joint work across various functional areas particularly desirable. This may avoid duplication of effort and allow the application of well-developed concepts to one's own field.

Despite its shortcomings, I remain convinced that agency theory has a lot to offer to a variety of functional fields, and to Management Science in particular. It allows to incorporate the multiperson aspect (incentives and risk attitude) too often overlooked in the traditional approaches. In-depth analysis of the extensions mentioned above should yield many interesting results and improve our understanding of contracting.

Acknowledgements

I would like to thank J.M. Harrison for suggesting I undertake this work, and Avi Mandelbaum and the participants at the Summer Workshop on Operations Management (MIT) for their useful comments.

References

Akerlof, G. (1970), "The market for 'Lemons': Quality and the market mechanism", *Quarterly Journal of Economics* 84, 488–500.

Antle, R. (1982), "The auditor as an economic agent", *Journal of Accounting Research* 20/2, 503–527.

Antle, R., and Fellingham, J. (1990), "Resource rationing and organizational slack in a two-period model", *Journal of Accounting Research* 28/1, 1–24.

Baiman, S. (1982), "Agency research in managerial accounting: A survey", *Journal of Accounting Literature* 1, 154–213.

Baron, D. (1989), "Design of regulatory mechanisms and institutions", in: R. Schmalensee and D. Willig (eds.), *Handbook of Industrial Organization*, North-Holland, Amsterdam, 1347–1447.

Basu, A., Lal, R., Srinivasan, V., and Staelin, R. (1985), "Salesforce compensation plans: An agency theoretic perspective", *Marketing Science* 4/4, 267–291.

Bhattacharya, S., and Pfleiderer, P. (1985), "Delegated portfolio management", *Journal of Economic Theory* 36/, 1–25.

Blanchard, G., Chow, C., and Noreen, E. (1986), "Information asymmetry, incentive schemes, and information biasing: The case of hospital budgeting under rate regulation", *The Accounting Review* LXI/1, 1–15.

Bonin, J. (1976), "On the design of managerial incentive structures in a decentralized planning environment", *The American Economic Review* 66/4, 682–687.

Christensen, J. (1981), "Communication in agencies", *The Bell Journal of Economics* 661–674.

Coughlan, A., and Sen, S. (1986), "Sales force compensation: Current research and managerial recommendations", WP G5, Yale School of Organisation and Management.

Demski, J., and Feltham, G. (1978), "Economic incentives in budgetary control systems", *The Accounting Review* LIII/2, 336–359.

Demski, J., and Kreps, D. (1980), "Models in managerial accounting", *Journal of Accounting Research* 20 Supplement 1980, 177–148.

Demski, J., and Sappington, D. (1989), "Hierarchical structure and responsibility accounting", *Journal of Accounting Research* 27/1, 40–58.

Dyl, E. (1988), "Corporate control and management compensation: Evidence on the Agency Problem", *Managerial and Decision Economics* 9/1, 21–25.

Epstein, R. (1985), "Agency costs, employment contracts, and labor unions," in: J.W. Pratt and R.J. Zeckhauser (eds.), *Principals and Agents*, HBS Press, 127–148.

Fan, L. (1975), "On the reward system", *The American Economic Review* 65/1, 226–229.

Fudenberg, D., Hölmstrom, B., and Tirole, J. (1990),"Short term contracts and long term relationships", *Journal of Economic Theory* 51/1, 1–31.

Fudenberg, D., and Tirole, J. (1990), "Moral hazard and renegotiation in agency contracts", *Econometrica* 58/6, 1279–1320.

Gjesdal, F. (1988), "Piecewise linear incentive schemes", *Scandinavian Journal of Economics* 90/3, 305–328.

Grossman, S., and Hart, O. (1983), "An analysis of the principal–agent problem", *Econometrica* 51/1, 7–45.

Haugen, R., and Senbet, L. (1981), "Solving the agency problems of external capital through options", *The Journal of Finance* XXXVI/3, 629–647.

Hölmstrom, B., and Milgrom, P. (1987), "Aggregation and linearity in the provision of intertemporal incentives", *Econometrica* 55 303–328.

Hölmstrom, B., and Ricart i Costa, J. (1986), "Managerial incentives and capital management", *The Quarterly Journal of Economics* 101, 835–860.

Hölmstrom, B., and Tirole, J. (1989), "The theory of the firm" in: R. Schmalensee and D. Willig (eds.), *Handbook of Industrial Organization*, North-Holland, Amsterdam, 61–133.

Ijiri, Y., Kinard, J., and Putney F. (1968), "An integrated evaluation system for budget forecasting and operating performance with a classified budgeting bibliography", *Journal of Accounting Research* 6/1, 21–28.

Kaplan, R. (1984), "The evolution of management accounting", *The Accounting Review* LIX/3, 390–418.

Kaplan, R., and Atkinson, A. (1989), *Advanced Management Accounting* second edition, Prentice-Hall, Englewood Cliffs, NJ.

Lal, R. (1986), "Delegating pricing responsibility to the salesforce", *Marketing Science* 5/2, 159–168.

Lal, R., and Staelin, R. (1986), "Salesforce compensation plans in environments with asymmetric information", *Marketing Science* 5/3, 179–198.

Lambert, R. (1984), "Income smoothing as rational equilibrium behavior", *The Accounting Review* October 604–618.

Lambert, R. (1986), "Executive effort and selection of risky projects", *Rand Journal of Economics* 17/1, 77–88.

Magee, R. (1988), "Variable cost allocation in a principal/agent setting", *The Accounting Review* LXIII/1, 42–54.

Malcolmson, J.M., and Spinnewyn, F. (1988), "The multiperiod principal-agent problem", *Review of Economic Studies* LV, 391–408.

McAfee, R.P., and McMillan, J. (1987), "Competition for agency contracts", *Rand Journal of Economics* 18/2, 296–307.

Radner, R. (1985), Repeated principal agent games with discounting", *Econometrica* 53, 1173–1198.

Ross, S. (1973), "Decision making under uncertainty", *American Economic Review* 63/2, 134–139.

Ross, S. (1977), "The determination of financial structure: The incentive-signalling approach", *The Bell Journal of Economics* 8/1, 23–40.

Smith, A. (1776), *The Wealth of Nations*, Cannan Edition, (Modern Library, New York, 1937) p. 700.

Spence, M., and Zeckhauser, R. (1971), "Insurance information and individual action", *American Economic Review* 61 (Supplement) 380–387.

Stiglitz, J. (1974), "Incentives and risk sharing in sharecropping", *Review of Economic Studies* 41, 219–255.

Thompson, R., and Wright, M. (1989), "Bonding, agency costs and management buyouts: A note", *Bulletin of Economic Research* 41/1, 69–75.

Tirole, J. (1989), *The Theory of Industrial Organization* MIT Press, Cambridge, MA.

van Ackere, A. (1990), "Conflicting interests in the timing of job", *Management Science* 36/8, 970–984.

Waller, and S., Bishop, R. (1990), "An experimental study of incentive pay schemes, communication and intra-firm allocation", *The Accounting Review* 65/4, 812–836.

Weitzman, M. (1977), "The new Soviet incentive model", *The Bell Journal of Economics* 7/1, 251–257.

Wilson, R. (1968), "The theory of syndicates", *Econometrica* 36/1, 119–132.

Wolfson, M. (1985), "Empirical evidence of incentive problems and their mitigation in oil and gas shelter programs", in: J.W. Pratt and R.J. Zeckhauser (eds.), *Principals and Agents*, HBS Press, 101–125.