Economics Letters 74 (2002) 291–300

economics letters

www.elsevier.com/locate/econbase

Informed principal, moral hazard, and the value of a more informative technology

Hector Chade*, Randy Silvers

Department of Economics, Arizona State University, Main Campus, PO Box 873806, Tempe, AZ 85287-3806, USA

Received 4 April 2001; accepted 16 May 2001

Abstract

We analyze a principal-agent model with moral hazard in which the principal has private information about the technology. We characterize Perfect Bayesian Equilibria of the contracting game that possess the following properties: (i) a principal with a more informative technology ends up earning less profits than a principal with a less informative one does; (ii) compared to the complete information case, the actions implemented by the privately informed principal can be distorted; (iii) the agent can end up being better off when the principal has private information.

Keywords: Moral hazard; Principal-agent; Informed principal

JEL classification: D82

1. Introduction

A standard paradigm in contract theory is the principal-agent model, where a risk neutral principal hires a risk averse agent to perform a certain task, but she cannot observe the action taken by the agent. The contract is based on an observable stochastic outcome that depends on the agent's action. The basic problem for the principal in this situation is to design a compensation scheme that maximizes her expected profits (Holmström, 1979; Grossman and Hart, 1983).

Although there are extensions of the principal-agent model where the agent has (or will come to have) private information (Holmström, 1979; Myerson, 1982; Sobel, 1993), so far most of the literature on moral hazard assumes that the principal does not have private information at the contracting date. However, it is easy to come up with examples where this assumption is too restrictive. For instance, a firm hiring a manager may have more information about the market

*Corresponding author. Tel.: +1-480-965-4714; fax: +1-480-965-0748.
E-mail address: hector.chade@asu.edu (H. Chade).

0165-1765/02/$ – see front matter
PII: S0165-1765(01)00557-2

conditions than the manager, whose effort is unobservable; in procurement contracts, the government can be privately informed about some features of the project that are unbeknown to the firm, whose actions can affect the cost of the project; in insurance, a company may have better information about the accident probabilities than the client, whose level of care is unobservable.

In this note, we analyze a principal-agent model with moral hazard in which the principal has private information about the technology, i.e. the stochastic relationship between the action and the observable outcome, and the contract that she offers may signal this information to the agent. Following the terminology introduced by Maskin and Tirole (1992) (see also Myerson, 1983), this is a principal-agent model with an *informed principal*. We characterize Perfect Bayesian Equilibria of the game that possess the following properties that do not arise in its complete information counterpart: first, a principal with a more informative technology ends up earning less profits than a principal with a less informative one does; second, compared to the complete information case, the actions implemented by the informed principal can be distorted, and the distortion can even be in an upward direction (i.e., a higher action is implemented under incomplete information); finally, in equilibrium the agent can end up being better off when the principal has private information.[1]

To be sure, our work is related to the large literature on the principal-agent model with moral hazard; our results show that allowing the principal to have private information can have nontrivial consequences. The paper is also related to an important paper by Maskin and Tirole (1992), where an informed principal contracts with an agent (who can also be privately informed) and the contract proposal may reveal part of the principal's private information. Unlike their work, we allow the agent to take an unobservable action after accepting the contract but in most of our analysis we do not allow the principal to use mechanisms as general as the ones considered by Maskin and Tirole (1992). We view our results as a first step towards a generalization of their framework to the case in which the agent takes an unobservable action. Also, our work is related to the literature that analyzes the value of a more informative technology in the principal-agent model (Grossman and Hart, 1983; Kim, 1995; Jewitt, 1997); our results imply that this value can be negative when the principal is privately informed. Another related paper is Sobel (1993), who shows in a two outcome model that when the agent can have access to private information, a risk neutral principal prefers informed to uninformed agents; the separating equilibria that we characterize have the property that, from an ex-ante point of view (i.e. before she knows her type), the principal prefers the agent to know the type of principal he faces. Finally, our paper is related to Inderst (2001), who also analyzes an agency model with an informed principal but under the assumption that both parties are risk neutral; he shows that private information induces the more profitable principal to provide less than first-best incentives (i.e., a 'flatter' contract) in the least-cost separating equilibrium.

The next section describes the model. Section 3 contains the main results of the paper. Section 4 casts some light on the robustness of the results. Section 5 concludes. Due to space limitations, the proofs of the results have been omitted; they can be found in the (downloadable) working paper version Chade and Silvers (2001).

[1] In the complete information version of the model, a well-known result of Grossman and Hart (1983) (Proposition 13, p. 35) shows that the value of a more informative technology is always positive for the principal. Moreover, the individual rationality constraint always holds with equality in the additively separable case.

2. The model

The principal-agent relationship we analyze can be described as follows. A risk-neutral principal wants to hire an agent to perform a certain task, but she cannot monitor the agent's actions. The agent is risk-averse and his preferences over income and action are summarized by a vonNeumann–Morgenstern utility function $U(I, a) = V(I) - a$, where $I \in (\underline{I}, \infty)$ is his income and $a \in \{a_1, a_2\} \subset \mathbb{R}_+$ is his action, with $a_1 < a_2$, $V'(I) > 0$, $V''(I) < 0$, and $\lim_{I \to \underline{I}} V(I) = -\infty$. If he does not work for the principal, the agent can obtain a reservation utility $\overline{U}$.

The principal has private information about the technology that relates the action of the agent with his performance. There are two possible technologies: $\theta = (q, \pi)$, where $q = (q_f, q_s) \in \mathbb{R}^2_+$ denotes the vector of gross profit levels for the principal, with $q_f < q_s$ (f is a mnemonic for 'failure' and s for 'success'), $\pi(a) = (\pi_f(a), \pi_s(a))$ is the probability distribution over the outcomes given the action a, $\pi(a) \gg 0$ for every a, $\pi_s(a_2) > \pi_s(a_1)$ (monotone likelihood ratio property); and $\theta' = (q', \pi')$, where $q' = (q'_f, q'_s)$, $\pi'(a) = (\pi'_f(a), \pi'_s(a))$, and $\pi'(a) \gg 0$ for every a. The two technologies are related by a 2×2 stochastic matrix (the elements of each column sum to one) $R \gg 0$ as follows: (i) $\pi'(a) = R\pi(a)$ for each a (here π' and π are column vectors) and (ii) $q'R = q$. The elements of R are denoted by r_{ij}, and it is assumed that $r_{22} > r_{21}$; this ensures that $\pi'_s(a_2) > \pi'_s(a_1)$ and $q'_f < q'_s$. The agent's prior belief that the principal's type is θ is given by $p \in (0, 1)$.

Intuitively, the transformation from $\pi(a)$ to $\pi'(a)$ described in (i) corresponds to a decrease in informativeness in Blackwell's sense (Grossman and Hart, 1983, p. 35, GH hereafter); a principal whose type is θ will be referred to as the principal with the more informative technology. The assumption $q'R = q$ (also used by GH) ensures that the expected revenues of types θ and θ' if they want to implement action a are equal; formally, $B(a) = B'(a)$, where $B(a) = \pi_f(a)q_f + \pi_s(a)q_s$ and $B'(a) = \pi'_f(a)q'_f + \pi'_s(a)q'_s$. As we shall see below, this simplification allows us to focus on the differences in the cost incurred by both types of the principal when each implements a particular action a.

In this setting, a contract for a principal of type θ is a compensation scheme $I(\theta) = (I_f, I_s)$ that specifies the remuneration to be paid if 'failure' (f) or 'success' (s) is observed; similarly, a contract for θ' is given by $I(\theta') = (I'_f, I'_s)$.

Remark. Notice that we have implicitly assumed that contracts are written conditional on 'success' and 'failure' and not on the actual values of q and q'. The reason is that, given our simplifying assumption $q'R = q$, we want to avoid the trivial and uninteresting case in which all the information gets revealed once the agent observes a contract written in terms of salaries paid when the components of q or q' are realized. One way to justify this assumption formally is to assume that what is contractible is not q or q', but a signal or performance measure $x \in \{0, 1\}$ with 1 corresponding to 'success' and 0 to 'failure', whose conditional distribution depends on the type of the principal and the action chosen by the agent. An alternative is to assume that $\theta' = (q, \pi')$ (with q fixed) and deal with different expected revenues for the two types of the principal; we present an example in Section 4 that reveals that similar equilibrium properties emerge in this set up as well.

The timing of the interaction is as follows: first the principal proposes a contract to the agent; after observing the contract, the agent updates his beliefs about the principal's type and then accepts or rejects the offer. If the contract is rejected, the game ends and the agent collects his reservation utility.

If he accepts, then he chooses an action that is unobservable to the principal; after that, both observe whether the result is a 'success' or a 'failure', and then the payoffs are distributed.

The solution concept we use is Perfect Bayesian Equilibrium (PBE); that is, we look for a profile of strategies and agent's beliefs such that, at any stage of the game, strategies are optimal given the beliefs, and these are derived using Bayes' rule whenever possible. Loosely speaking, we need to specify a vector $(I(\theta), I(\theta'))$ for the principal, and an acceptance rule, a choice of action, and beliefs for the agent for each possible contract, in such a way that the aforementioned optimality and consistency conditions are satisfied.

3. Main results

3.1. Preliminaries

Consider first the complete information case, i.e., the agent knows what type of principal he is facing. Then the following results are well-known from GH:

1. If the principal wants to implement a_1, then no matter what her type is, the optimal contract is to offer a constant remuneration

$$I^* = h(\overline{U} + a_1), \tag{1}$$

where $h = V^{-1}$.

2. If the principal's type is θ and she wants to implement a_2, then the optimal contract (I_f^*, I_s^*) is the unique solution to the following system of equations:

$$\begin{aligned} \pi_f(a_2)V(I_f) + \pi_s(a_2)V(I_s) &= \overline{U} + a_2 \\ (\pi_f(a_2) - \pi_f(a_1))V(I_f) + (\pi_s(a_2) - \pi_s(a_1))V(I_s) &= a_2 - a_1, \end{aligned}$$

where the first equation is the individual rationality constraint and the second is the incentive compatibility constraint of the agent. Therefore,

$$I_f^* = h\left(\overline{U} + a_2 - \frac{(a_2 - a_1)\pi_s(a_2)}{\pi_s(a_2) - \pi_s(a_1)}\right) \tag{2}$$

$$I_s^* = h\left(\overline{U} + a_2 + \frac{(a_2 - a_1)\pi_f(a_2)}{\pi_s(a_2) - \pi_s(a_1)}\right), \tag{3}$$

and the expected cost for the principal is $\pi_f(a_2)I_f^* + \pi_s(a_2)I_s^*$. Similarly, if the principal's type is θ', then

$$I_f^{*'} = h\left(\overline{U} + a_2 - \frac{(a_2 - a_1)\pi'_s(a_2)}{\pi'_s(a_2) - \pi'_s(a_1)}\right) \tag{4}$$

$$I_s^{*'} = h\left(\overline{U} + a_2 + \frac{(a_2 - a_1)\pi'_f(a_2)}{\pi'_s(a_2) - \pi'_s(a_1)}\right), \tag{5}$$

and her expected cost is $\pi'_f(a_2)I_f^{*'} + \pi'_s(a_2)I_s^{*'}$.

3. The action that principal θ implements depends on whether $B(a_2) - (\pi_f(a_2)I_f^* + \pi_s(a_2)I_s^*) \gtreqless B(a_1) - I^*$, and similarly for θ'.
4. The cost of implementing a_2 is greater for θ' than for θ (GH Proposition 13); that is, $\pi'_f(a_2)I_f^{*'} + \pi'_s(a_2)I_s^{*'} > \pi_f(a_2)I_f^* + \pi_s(a_2)I_s^*$. Since $B(a_2) = B'(a_2)$, it follows that the value of a more informative technology is *positive* for the principal if a_2 is the action that she implements; i.e., her profits are higher when the technology is θ. A straightforward implication of this result is that if the principal finds it optimal to implement a_1 when her technology is θ, then she will also implement a_1 when it is θ'.[2]
5. Since $B(a_2) - B(a_1) = B'(a_2) - B'(a_1) = (q_s - q_f)(\pi_s(a_2) - \pi_s(a_1))$, it follows that if (q_f, q_s) is such that

$$q_s - q_f \geq \frac{\pi'_f(a_2)I_f^{*'} + \pi'_s(a_2)I_s^{*'} - I^*}{\pi_s(a_2) - \pi_s(a_1)}, \tag{6}$$

then both θ and θ' will find it optimal to implement a_2.[3]

3.2. Equilibrium analysis

Let's turn now to the incomplete information case. Although a complete characterization of the set of equilibria is possible, we shall focus on PBE of the game that possess some interesting properties that do not arise under complete information.

The first question we ask is: can $((I_f^*, I_s^*), (I_f^{*'}, I_s^{*'}))$ be part of a separating PBE? The following result reveals that the answer is negative.[4]

Proposition 1. *The following inequalities cannot hold simultaneously:*

$$\begin{aligned} \pi_f(a_2)I_f^* + \pi_s(a_2)I_s^* &\leq \pi_f(a_2)I_f^{*'} + \pi_s(a_2)I_s^{*'} \\ \pi'_f(a_2)I_f^{*'} + \pi'_s(a_2)I_s^{*'} &\leq \pi'_f(a_2)I_f^* + \pi'_s(a_2)I_s^*. \end{aligned}$$

That is, given (I_f^*, I_s^*) and $(I_f^{*'}, I_s^{*'})$, it is always the case that at least one type of principal wants to mimic the other. In particular, if $\pi_s(a_2) > \pi'_s(a_2)$, then θ' wants to mimic θ.[5]

This proposition implies that, in any separating PBE where both types of the principal implement

---

[2] We shall see below that these assertions do not hold under incomplete information.

[3] Notice that the right side of (6) does not depend on q.

[4] We also get a negative answer in the case where (q_f, q_s) is such that, under complete information, θ finds it optimal to implement a_2 and θ' implements a_1. In particular, it is easy to show that if $\pi_s(a_2) > \pi'_s(a_2)$, then $((I_f^*, I_s^*), I^*)$ cannot be part of a separating PBE, for θ' will want to deviate and mimic θ.

[5] Given $\pi' = R\pi$, $\pi_s(a_2) > \pi'_s(a_2)$ is equivalent to $\pi_s(a_2) > r_{21}/(r_{21} + r_{12})$.

a_2, at least one of them will make less profits than in the complete information case. As the next proposition clearly illustrates, one can even construct PBE of this sort where type θ' ends up earning *higher* profits than θ does. The proof uses the following lemma:

Lemma 1. *Let* $\pi_s(a_2) > \pi'_s(a_2)$ *and assume* $p = 0$ (*i.e., the agent believes that the principal's type is* θ' *with probability one*). *Then the contract that minimizes the expected cost of implementing* a_2 *for* θ *is* $(I_f^{*'}, I_s^{*'})$, *defined by* (4) *and* (5).

Proposition 2. *Suppose* $\pi_s(a_2) > \pi'_s(a_2)$. *Then there exists a set of* (q_f, q_s) *such that the following profile of strategies and beliefs constitute a separating PBE*:

(*i*) *The principal offers* $(\hat{I}_f, \hat{I}_s)$ *if her type is* θ *and* $(I_f^{*'}, I_s^{*'})$ *if her type is* θ', *where* $(\hat{I}_f, \hat{I}_s)$ *is the unique solution to*

$$\pi'_f(a_2)\hat{I}_f + \pi'_s(a_2)\hat{I}_s = \pi'_f(a_2)I_f^{*'} + \pi'_s(a_2)I_s^{*'} \tag{7}$$

$$V(\hat{I}_s) = \frac{a_2 - a_1}{\pi_s(a_2) - \pi_s(a_1)} + V(\hat{I}_f), \tag{8}$$

and $(I_f^{*'}, I_s^{*'})$ *is given by* (4) *and* (5);

(*ii*) *The agent accepts* $(\hat{I}_f, \hat{I}_s)$ *and* $(I_f^{*'}, I_s^{*'})$, *and then chooses* a_2; *for any other contract, his acceptance decision and action choice are determined by the individual rationality and incentive compatibility constraints given his beliefs*;

(*iii*) *The agent's beliefs are as follows*: $\mu(\theta|(\hat{I}_f, \hat{I}_s)) = 1$, *and* $\mu(\theta|(I_f, I_s)) = 0$ *otherwise*.[6]

It is easy to show that in this equilibrium θ earns less profits than θ' does. By construction, θ' is indifferent between $(I_f^{*'}, I_s^{*'})$ and $(\hat{I}_f, \hat{I}_s)$; that is, $B'(a_2) - (\pi'_f(a_2)\hat{I}_f + \pi'_s(a_2)\hat{I}_s) = B'(a_2) - (\pi'_f(a_2)I_f^{*'} + \pi'_s(a_2)I_s^{*'})$. Also, $\hat{I}_s > \hat{I}_f$ (which follows from the incentive compatibility constraint (8)) and $\pi_s(a_2) > \pi'_s(a_2)$ yield

$$\pi'_f(a_2)\hat{I}_f + \pi'_s(a_2)\hat{I}_s < \pi_f(a_2)\hat{I}_f + \pi_s(a_2)\hat{I}_s. \tag{9}$$

Since $B'(a_2) = B(a_2)$, (9) implies $B'(a_2) - (\pi'_f(a_2)I_f^{*'} + \pi'_s(a_2)I_s^{*'}) > B(a_2) - (\pi_f(a_2)\hat{I}_f + \pi_s(a_2)\hat{I}_s)$, and therefore θ' earns higher profits than θ does. The intuition is rather simple: given the agent's beliefs, if θ offers any alternative contract she will be believed to be θ', and any contract that satisfies individual rationality and incentive compatibility for those beliefs is less profitable than $(\hat{I}_f, \hat{I}_s)$ for θ (see Lemma 1). But then $\pi_s(a_2) > \pi'_s(a_2)$ makes the use of this contract more costly for her than for θ', who is in turn indifferent between this and her optimal contract under complete information.

The proof of Proposition 2 also shows that $(\hat{I}_f, \hat{I}_s)$ satisfies the agent's individual rationality constraint with strict inequality. Therefore, ignorance of the technology of the principal can make the agent *better off*.

Notice that (7) and (9) imply that the set of (q_f, q_s) that makes the profile of strategies and beliefs a PBE is contained within the set determined by (6); that is, the action implemented by the principal

[6] Eq. (8) is the agent's incentive compatibility constraint when he believes that the principal is of type θ with probability one.

cannot be distorted when compared to the complete information case, no matter what her type is. The following result shows that this need not be the case in a PBE.

Proposition 3. *Suppose* $\pi_s(a_2) > \pi'_s(a_2)$. *Then there exists a set of* (q_f, q_s) *such that the following profile of strategies and beliefs constitute a separating PBE*:

(*i*) *The principal offers* I^* *if her type is* θ *and* $(I_f^{*'}, I_s^{*'})$ *if her type is* θ', *where these contracts are defined by* (1), *and* (4) *and* (5), *respectively*:

(*ii*) *The agent accepts* I^* *and* $(I_f^{*'}, I_s^{*'})$, *and then chooses* a_1 *and* a_2, *respectively*; *for any other contract, his acceptance decision and action choice are determined by the individual rationality and incentive compatibility constraints given his beliefs*;

(*iii*) *The agent's beliefs are as follows*: $\mu(\theta|I^*) = 1$, *and* $\mu(\theta|(I_f, I_s)) = 0$ *otherwise*.

In this equilibrium, since $B(a_1) = B'(a_1)$ and θ' implements a_2, it follows that

$$B'(a_2) - (\pi'_f(a_2)I_f^{*'} + \pi'_s(a_2)I_s^{*'}) > B(a_1) - I^*,$$

and therefore θ' earns higher profits than θ does. Moreover, the action implemented by θ is evidently distorted *downward* when compared with the complete information case (see result 4 in Section 3.1).

A straightforward corollary of Propositions 2 and 3 is that in these PBE the principal, from an ex-ante point of view (i.e., before nature reveals her type), is *better off* when she faces an informed agent than when the agent does not observe nature's move.

Let $(\overline{\pi}_f(a), \overline{\pi}_s(a)) = (p\pi_f(a) + (1-p)\pi'_f(a), p\pi_s(a) + (1-p)\pi'_s(a))$ be the probability distribution of 'success' and 'failure' when the agent's belief is $0 < p < 1$; consider the contract $(\bar{I}_f, \bar{I}_s)$ given by

$$\bar{I}_f = h(\overline{U} + a_2 - \frac{(a_2 - a_1)\overline{\pi}_s(a_2)}{\overline{\pi}_s(a_2) - \overline{\pi}_s(a_1)}) \tag{10}$$

$$\bar{I}_s = h(\overline{U} + a_2 + \frac{(a_2 - a_1)\overline{\pi}_f(a_2)}{\overline{\pi}_s(a_2) - \overline{\pi}_s(a_1)}). \tag{11}$$

This is the contract that satisfies with equality the constraints of the agent when his belief is p. The next proposition reveals that the negative value of a more informative technology and the distortions in actions are also present in PBE of the pooling type.

Proposition 4. *Suppose* $\pi_s(a_2) > \pi'_s(a_2)$. *Then there exists a set of* (q_f, q_s) *such that the following profile of strategies and beliefs constitute a pooling PBE*:

(*i*) *Both types of the principal offer* $(\bar{I}_f, \bar{I}_s)$;

(*ii*) *The agent accepts* $(\bar{I}_f, \bar{I}_s)$ *and then chooses* a_2; *for any other contract, his acceptance decision and action choice are determined by the individual rationality and incentive compatibility constraints given his beliefs*;

(*iii*) *The agent's beliefs are as follows*: $\mu(\theta|(\bar{I}_f, \bar{I}_s)) = p$, *and* $\mu(\theta|(I_f, I_s)) = 0$ *otherwise*.

It is immediate from $\pi_s(a_2) > \pi'_s(a_2)$ and $\bar{I}_s > \bar{I}_f$ that θ' makes *more* profits than θ does. Regarding the distortion in the actions implemented by the principals, the following result provides the answer:

Lemma 2. *If* $\pi_s(a_2) > \pi'_s(a_2)$, *then* $\pi_f(a_2)\bar{I}_f + \pi_s(a_2)\bar{I}_s > \pi_f(a_2)I_f^* + \pi_s(a_2)I_s^*$ *and* $\pi'_f(a_2)\bar{I}_f + \pi'_s(a_2)\bar{I}_s < \pi'_f(a_2)I_f^{*'} + \pi_s(a_2)I_s^{*'}$.

That is, $(\bar{I}_f, \bar{I}_s)$ is less costly than $(I_f^{*'}, I_s^{*'})$ for θ' and it is more costly than (I_f^*, I_s^*) for θ. It then follows that, although the action implemented by θ cannot be distorted, there exists a (q_f, q_s) that satisfies Proposition 4 and such that θ' would implement a_1 under complete information but a_2 under private information. Thus, compared to the complete information case, the action implemented by θ' can be distorted in an *upward* direction.

4. Discussion

The equilibria described in the previous section illustrate some interesting features that emerge when the contract proposal can be used to signal the principal's private information to the agent. For example, the value of a more informative technology is *not* necessarily positive for the principal when she is privately informed about it. As we mentioned before, GH showed that this cannot arise under complete information. Also, the equilibrium analysis shows that the actions implemented by each type of principal can be *different* than the corresponding ones without private information; the kinds of distortions illustrated by Propositions 3 and 4 have no counterpart in the complete information case. Finally, in equilibrium the agent's individual rationality constraint can hold with *strict* inequality, a feature that cannot arise under complete information and additively separable preferences.

One may wonder whether these features are robust to variations in the model or to stronger equilibrium requirements like the Cho–Kreps intuitive criterion. We now shed some light on these issues.

(i) *Equilibrium Refinements:* A standard refinement used in applications is the intuitive criterion of Cho and Kreps (1987). The following proposition shows that the separating equilibria derived in the previous section satisfy the requirements imposed by this refinement.

Proposition 5. *The PBE of Propositions* 2 *and* 3 *pass the Cho–Kreps intuitive criterion, but the PBE of Proposition* 4 *fails to pass this refinement.*

(ii) *Constant q:* The results were derived under the assumption that $q'R = q$ in order to focus on the cost of implementing the actions with the two technologies, without worrying about differences in expected revenues. To avoid a triviality, we also assumed that contracts were written conditional on 'success' and 'failure' and not on the actual values in q and q'. The following example illustrates that if we drop these assumptions and keep q constant while $\pi'(a) = R\pi(a)$, then it is easy to construct PBE where θ earns less than θ' does, a result that is driven mainly by the differences in expected revenues.

Example 1. Let $U(I, a) = \ln I - a$, $\{a_1, a_2\} = \{0, 1\}$, $(q_f, q_s) = (10, 40)$, $\overline{U} = 1$, $\pi(0) = (0.7, 0.3)$, $\pi(1) = (0.3, 0.7)$, $\pi'(0) = (0.4, 0.6)$, and $\pi'(1) = (0.2, 0.8)$. Then it is easy to show that the following profile of strategies and beliefs constitute a PBE: (i) the principal offers $(I_f^*, I_s^*) = (1.28, 15.64)$ if her type is θ and $I^* = 2.72$ if her type is θ'; (ii) the agent accepts (I_f^*, I_s^*) and I^* and chooses a_2 and a_1, respectively; for any other contract, he chooses in an optimal way given his out-of-equilibrium beliefs; (iii) his beliefs are: $\mu(\theta|I^*) = 0$ and $\mu(\theta|(I_f, I_s)) = 1$ for any alternative contract. Since $B(1) - (0.3 \times 1.28 + 0.7 \times 15.7) = 19.67$ and $B'(0) - 2.72 = 25.28$, it follows that θ' earns higher profits than θ does. That is, even though θ implements a_2 with her complete information contract and θ' implements a_1, θ ends up making less profits than θ' does.

(iii) *More General Contracts:* An important limitation of our model is that we restrict attention to a particular mechanism used by the principal, i.e., she offers one contract if she is of type θ and one if her type is θ'. In the spirit of Maskin and Tirole (1992), one could envision a more general model as follows: first each type of principal offers a *menu* of two contracts; the agent observes the menu offered, updates his beliefs, and then accepts or rejects; if he accepts, then the principal announces her type and thus her contract offer; after the announcement, the agent updates his beliefs and chooses an action, nature reveals the outcome, and the contract corresponding to the type announced is executed. The following result casts some light on this important extension, although a complete characterization of the equilibria of this model remains to be done:

Proposition 6. *There exists a PBE of the extended model with the same properties as the PBE described in Proposition* 2.

5. Conclusion

In summary, we have analyzed a principal-agent model with moral hazard where the principal is privately informed about the technology, and have obtained results that do not arise under complete information. To make the analysis tractable, we have restricted attention to a simple model with two actions, two outcomes, and two types. Although the intuition of the results does not seem to depend on this restriction, our proofs exploit some of the key aspects of this 'binary' model. Extending the results beyond this simple framework seems like an interesting albeit messy robustness exercise. A more important limitation of our analysis is the fact that, except for Proposition 6, we do not allow the principal to use more general mechanisms as in Maskin and Tirole (1992). The incorporation of moral hazard in their general characterization is, in our opinion, an important and challenging problem that has not been hitherto analyzed. Our results suggest that this extension can generate nontrivial insights about the principal-agent relationship.

Acknowledgements

We would like to thank Alejandro Manelli and Ed Schlee for their helpful comments and suggestions. The paper also benefited from comments by seminar participants at the Brown Bag Lunch Workshops at Arizona State University.

References

Chade, H., Silvers, R., 2001. Informed principal, moral hazard, and the value of a more informative technology. Working Paper, Department of Economics, Arizona State University (http://www.cob.asu.edu/pubs/ecn/pub.cfm).

Cho, I., Kreps, D., 1987. Signalling games and stable equilibria. Quarterly Journal of Economics 102, 179–221.
Grossman, S., Hart, O., 1983. An analysis of the principal-agent problem. Econometrica 51, 7–45.
Holmström, B., 1979. Moral hazard and observability. Bell Journal of Economics 10, 74–91.
Inderst, R., 2001. Incentive schemes as a signaling device. Journal of Economic Behavior and Organization 44, 455–465.
Jewitt, I., 1997. Information and principal agent problems. Mimeo, Department of Economics, University of Bristol.
Kim, S., 1995. Efficiency of an information system in an agency model. Econometrica 63, 89–102.
Maskin, E., Tirole, J., 1992. The principal-agent relationship with an informed principal, II: common values. Econometrica 60, 1–42.
Myerson, R., 1982. Optimal coordination mechanisms in generalized principal-agent problems. Journal of Mathematical Economics 10, 67–81.
Myerson, R., 1983. Mechanism design by an informed principal. Econometrica 51, 1767–1798.
Sobel, J., 1993. Information control in the principal-agent problem. International Economic Review 34, 259–269.