Lecture 3: Principal-Agent Problem[1]

Luís Vasconcelos

October 2011

[1] Adapted from slides originally created by Andrea Galeotti.

Chapter 14 of MWG.

In previous lectures we studied situations in which asymmetric information existed at the time of contracting.

We now study the implication of contracting when parties anticipate that information asymmetries will develop in their relationship.

Anticipating the development of such information asymmetries, the parties will seek to design a contract that mitigates the possible difficulties caused.

We distinguish between hidden actions and hidden information.

An example of hidden actions (moral hazard) is when a firm (principal) hires a manager (agent) and the profit of the firm depends on the effort chosen by the manager, which is however not fully observable.

An example of hidden information is when the manager comes to possess superior information about the firm's opportunities.

Of course, the interest is when these asymmetric information creates incentives to the informed agent to behave in a way that has negative effect on the principal. The question is how the principal can create incentives for the agent to take certain actions even if such actions are not observable.

Hidden Actions

We consider the following framework. A firm must hire a manager for a one-time project. The firm offers a contract to the manager. If the manager accepts the contract, he must choose an effort $e \in \{e_L, e_H\}$, where $e_H > e_L$.

The profit of the project is π which is distributed in $[\underline{\pi}, \bar{\pi}]$.

The key point is that the profit of the project is stochastically related to e. In particular, the distribution function is $F(\pi|e)$, with the property that $F(\pi|e_H)$ first order stochastic dominates $F(\pi|e_L)$. That is $F(\pi|e_H) \leq F(\pi|e_L)$, for all $\pi \in [\underline{\pi}, \bar{\pi}]$. f is the density function.

The assumption on F implies that the expected profit of the project when the manager chooses high effort is higher than the expected profit of the project when the manager chooses low effort, i.e.,

$$\int_{\underline{\pi}}^{\bar{\pi}} f(\pi|e_H)\pi d\pi \geq \int_{\underline{\pi}}^{\bar{\pi}} f(\pi|e_L)\pi d\pi$$

The profit of the firm is $\pi - w$, where w is the wage. The utility of the manager for wage w and effort e is

$$u(w, e) = v(w) - g(e)$$

where v is increasing and concave in w, g is increasing in e. This implies that the manager likes income, he is risk averse over income lotteries and he dislikes effort.

The utility of the firm is the profit of the project net of the wage. So the firm is risk neutral.

The case of observable effort. We now ask what is the optimal contract that the firm can offer when effort is observable.

If effort is observable the contract of the firm will specify an effort that the manager should exert and the wage that the firm will pay given the observed profit of the project, i.e., $(e, w(\pi))$, where $e \in \{e_L, e_H\}$ and $w : [\underline{\pi}, \bar{\pi}] \rightarrow R^+$.

We also consider that the manager can always get u by simply rejecting the contract.

The problem of the firm is a simple maximization problem: find the effort level and the associated wage profile such that the expected profit of the firm is maximized subject to the constraint that the manager will accept the contract, e.g., participation constraint. Formally,

$$\max_{e \in \{e_L, e_H\}, w(\pi)} \int [\pi - w(\pi)] f(\pi|e) d\pi$$

subject to:

$$\int v(w(\pi)) f(\pi|e) d\pi - g(e) \geq \bar{u}$$

We can break the problem in two steps. In the first step, we fix e and we find $w(\pi)$ which solves the maximization problem. In the second step, we use the solution of the first step and find the optimal e.

First Step. Once e is fixed, we can rewrite the problem as follows

$$\min_{w(\pi)} \int w(\pi) f(\pi|e) d\pi$$

subject to the above participation constraint.

Next, note that in an optimal contract the constraint must be binding. Otherwise, the firm can decrease wage for some realization of π, the manager will still be willing to accept the contract and the firm will be better off.

Given that the constraint binds, we solve the problem using Lagrange method. We obtain that the optimal wage function satisfies:

$$\frac{1}{v'(w(\pi))} = \gamma$$

where γ is the Lagrangian multiplier. [Note that it is enough to solve the problem for each realization of π. The condition above is simply the first order condition].

The condition above allows us to differentiate between two cases.

Strictly risk averse manager. If the manager is strictly risk averse, then

$w(\pi)$ must be constant in π. Hence, the optimal wage, given e, is

$$\int v(w^*)f(\pi|e)d\pi = \bar{u} + g(e)$$

that is

$$v(w^*) = \bar{u} + g(e).$$

Therefore $w^* = v^{-1}(\bar{u} + g(e))$. Note that the optimal contract insures the manager from income risk, i.e., the wage is constant regardless of the realization of profit.

We can now solve for the optimal e.

Here, note that if the manager chooses e_H then $w^*(e_H) = v^{-1}(\bar{u} + g(e_H))$, whereas under e_L then $w^*(e_L) = v^{-1}(\bar{u} + g(e_L))$, and since $e_H > e_L$, g is increasing in e and v is increasing in his argument, it follows that $w^*(e_H) > w^*(e_L)$.

So, the manager faces the following trade-off. To ask for high effort, getting high expected profit but paying high wage, or asking for low effort getting low expected profit but paying a low wage. Formally, the firm chooses e_H if and only if

$$\int \pi f(\pi|e_H)d\pi - v^{-1}(\bar{u} + g(e_H)) \geq \int \pi f(\pi|e_L)d\pi - v^{-1}(\bar{u} + g(e_L))$$

Note that if the *manager is risk neutral*, then v' is constant and therefore there are many optimal contracts. Among these, the one that we have derived is still an optimal contract.

But what if effort is not observable?

The optimal contract under observable effort has two properties: one, it specifies an efficient effort and two, it insures fully the manager from income risk.

When the effort is not observable these two goals are in conflict: the only way to let the manager to choose optimal effort is to relate the wage to the realized outcome of the project. This property does not allow for fully insurance.

To see that it is exactly this issue to create inefficiencies we start to look at the manager who is risk neutral and therefore insurance against risk is not an issue. We will see that in this case the fact that effort is not observable, does not create inefficiencies.

A Risk Neutral Manager. Suppose that $v(w) = w$. Note that under observable effort and risk neutrality we have that: $w^* = \bar{u} + g(e^*)$ and $e^* = \arg\max_{e \in \{e_L, e_H\}} = E[\pi|e] - \bar{u} - g(e)$.

We now show that there exists a contract which achieves the same outcome of the optimal contract under observable effort. Hence, that contract must be optimal.

Consider the contract $w(\pi) = \pi - \alpha$. That is, the firm asks the manager to pay α and the manager gets the π which is realized. This is as if the manager becomes owner of the firm.

Suppose that the manager accepts this contract. Once he gets the contract he will choose an effort to maximize his utility, i.e.,

$$\begin{aligned}\max_{e\in\{e_L,e_H\}} \quad & \int w(\pi)f(\pi|e)d\pi - g(e) \\ = \quad & \int \pi f(\pi|e)d\pi + \alpha - g(e) \\ = \quad & E[\pi|e] + \alpha - g(e)\end{aligned}$$

Hence, the optimal effort of the manager is the same as the effort induced under the optimal contract when effort is observable.

So, this contract induces the optimal level of effort, once the contract has been signed by the manager.

Next, given this, note that the manager will accept the contract if and only if α satisfies

$$E[\pi|e^*] - \alpha - g(e^*) \geq \bar{u}$$

which implies that the firm will set a α^* which makes the above constraint to bind, i.e.,

$$\alpha^* = E[\pi|e^*] - g(e^*) - \bar{u}$$

and α^* is the payoff of the firm, which is the same as the payoff she obtains under observable effort.

A risk averse manager. In this case, a manager will have incentive to set high effort only if he is compensated for taking high risk. To find for the optimal contract we can proceed as before.

First, suppose the optimal contract requires to implement effort e. What is the the optimal incentive scheme that the firm can design to implement e. This amounts in finding the wage schedule $w(\pi)$ which minimizes the cost of the firm, under two constraints: one, the manager accepts to undertake the project at that effort (participation constraint) and two, the manager has indeed the incentive to exert effort e once he signed up for the project (incentive constraint).

Formally, given effort e, we need to solve the problem:

$$\min_{w(\pi)} \int w(\pi) f(\pi|e) d\pi$$

subject to: (i) participation constraint and (ii) incentive constraint

$$\int v(w(\pi)) f(\pi|e) d\pi - g(e) \geq \bar{u}$$

$$e = \arg \max_{e' \in \{e_L, e_H\}} \int v(w(\pi)) f(\pi|e') d\pi - g(e')$$

Implementing low effort. Suppose the firm wants to implement $e = e_L$. In this case the firm can offer a constant wage, like in case of observable effort, $w^*(e_L) = v^{-1}(\bar{u} + g(e_L))$, and since wage is independent of π, the manager will indeed choose effort level e_L.

Implementing high effort. Suppose the firm wants to implement $e = e_H$. We can rewrite the incentive constraints as

$$\int v(w(\pi))f(\pi|e_H)d\pi - g(e_H) \geq \int v(w(\pi))f(\pi|e_L)d\pi - g(e_L)$$

Writing the minimization problem, calling γ and μ the two multipliers, the KT-conditions are that for every $\pi \in [\underline{\pi}, \bar{\pi}]$ the following must hold:

$$-f(\pi|e_H) + \gamma v'(w(\pi))f(\pi|e_H) + \mu[f(\pi|e_H) - f(\pi|e_L)]v'(w(\pi)) = 0$$

or

$$\frac{1}{v'(w(\pi))} = \gamma + \mu \left[1 - \frac{f(\pi|e_L)}{f(\pi|e_H)}\right]$$

Lemma. It is easy to establish that $\gamma > 0$ and $\mu > 0$.

First, suppose $\mu = 0$. Then the wage schedule must be constant. But this implies that the manager will choose a low effort, which contradicts that high effort can be implemented.

Second, suppose $\gamma = 0$. First order stochastic dominance of F implies that there exists value of π such that $f(\pi|e_L)/f(\pi|e_H) > 1$. Since $\mu \geq 0$ and $\gamma = 0$, it follows that for that values of π it must be the case that $v'(w(\pi)) < 0$, which is impossible.

Since $\gamma > 0$ and $\mu > 0$, it follows that both the participation and incentive constraint must bind at the optimal contract which wants to implement e_H.

Remark. Take the fixed wage $\hat{w}$: such that $\gamma = 1/v'(\hat{w})$. Note that for all π which are more likely under high effort than under low effort ($f(\pi|e_L)/f(\pi|e_H) < 1$), the firm pays higher wage than the fixed wage, $w(\pi) > \hat{w}$. The reverse holds for the other realizations.

This is all driven by incentive effects: by structuring compensation in this way, it provides the manager an incentives to provide high effort instead of low effort.

Remark. Note that $w(\pi)$ is increasing in π if and only if

$$\frac{f(\pi|e_L)}{f(\pi|e_H)}$$

is decreasing in π. This is the monotone likelihood ration property: as π increases the likelihood of getting profit π under e_H relative to the likelihood of getting π under e_L must increase.

To see this note that if $w(\pi)$ is increasing in π then $v'(w(\pi))$ is decreasing in π and therefore $1/v'(w(\pi))$ is increasing in π, which requires that

$$\gamma + \mu \left[1 - \frac{f(\pi|e_L)}{f(\pi|e_H)} \right]$$

is increasing in π, which requires MLRP.

Remark MLRP is not implied by FOSD relation. Hence, the optimal contract to implement high effort must not have necessarily that wage increases in π.

Remark. It is easy to check that the expected wage induced by the optimal wage schedule for implementing effort e_H is higher than the fix wage under observable effort associated to implement high effort.

Remark. The remark above implies that, for a firm, it is more costly to implement high effort under unobservable effort, while the cost of implementing low effort is not altered by non observable effort. This implies that nonobservability of effort can lead to inefficient level of effort.

Hidden Information. Suppose now that effort is observable. However, once the contract has been signed, the manager realizes his ability to undertake the project and this defines his disutility of efforts. This is private information.

Consider that the manager can choose effort e. Given e the profit of the project is $\pi(e)$ which is increasing and concave in effort, and $\pi(0) = 0$.

Once the contract has been signed, the manager observes whether he is suited for the project, θ_H, or not, θ_L. The probability that $\theta = \theta_H$ is λ.

Under e and θ the disutility of effort is $g(e, \theta)$ which is: increasing and convex in e, it is decreasing in θ, the cross partial derivative is negative and $g(0, \theta) = 0$.

The utility of the manager is

$$u(w, e, \theta) = v(w - g(e, \theta)),$$

where v is increasing and strictly concave.

Observable Information. Consider the problem when θ is observable. In this case the contract will be such that: the firm offer (w_i, e_i) when θ_i is observed, $i = H, L$. The maximization problem reads:

$$\max_{(w_L, e_L), (w_H, e_H)} \lambda[\pi(e_H) - w_H] + (1 - \lambda)[\pi(e_L) - w_L]$$

subject to participation constraint

$$\lambda v(w_H - g(e_H, \theta_H)) + (1 - \lambda) v(w_L - g(e_L, \theta_L)) \geq \bar{u}$$

It is clear that, at the solution, the participation constraint must bind. We can then solve the maximization problem and obtain that $[(w_H^*, e_H^*), (w_L^*, e_L^*)]$ satisfies:

$$v'(w_H^* - g(e_H^*, \theta_H)) = v'(w_L^* - g(e_L^*, \theta_L))$$

and

$$\pi'(e_i^*) = g_e(e_i^*, \theta_i)$$

The first condition implies that the marginal utility of income of the manager is equalized across states. The condition is satisfied if and only if

$$w_H^* - g(e_H^*, \theta_H) = w_L^* - g(e_L^*, \theta_L),$$

which implies that manager's utility is the same across state. In other words, the optimal contract provide fully insurance to the manager.

The second condition says that optimal level of effort in state θ_i is such that marginal benefit of effort equals its marginal disutility.

What does happen when information is hidden? In this case, if the manager is offered menu $[(w_H^*, e_H^*), (w_L^*, e_L^*)]$, when he has a high type, he may find it profitable to tell the firm he is of type θ_L.

To understand the optimal contract, we first need to identify the set of possible contracts that the owner can offer. This set is very large. For example, the owner can offer a compensation which is a function of the profit of the project and that leaves the effort choice to the manager. Or, the owner can restrict the set of efforts that the manager can choose.

Since the set of contract is so large, the problem of optimal contract seems very difficult to solve. We now show that the problem can be simplify substantially once we appeal to the *revelation principle*.

The idea is to consider a revelation mechanism: a contract that asks the manager to announce his ability and it associates a wage and an effort to be chosen for each possible announcement.

The revelation principle says that to find an optimal contract we can restrict on revelation mechanisms for which the manager always *reveals truthfully* his productivity. These are the *incentive compatible* revelation mechanisms.

Proposition. In search for an optimal contract, the owner can without loss of generality restrict himself to contracts of the following form:

I. After the state $\theta \in \{\theta_L, \theta_H\}$ is realized, the manager is required to announce which state has occurred.
II. The contract specifies an outcome $(w(\hat{\theta}), e(\hat{\theta}))$ for each possible announcement.
III. For every possible realized state, the manager finds it optimal to report

the state truthfully.

To see this consider the following contract: the wage depends on the realized profit of the project, $w(\pi)$, and the effort is chosen by the manager. Suppose that given $w(\pi)$, the effort chosen by the manager is e_H if he has ability θ_H and e_L if he has ability θ_L.

We now show that the owner can construct an incentive compatible revelation mechanism which leads to the same outcome. The contract is the following: if the manager reports θ_L then the contract specifies wage $w(\pi(e_L))$ and effort e_L, if the manager reports θ_H then the contract specifies wage $w(\pi(e_H))$ and effort e_H.

Clearly, this contract leads to the same outcome induced by the original contract. We now need to check that it is incentive compatible.

For this it must be the case that when the manager is, say, of low productivity he has an incentive to report the owner that he is of low productivity. Note that under the original contract, a θ_L manager could have achieved the outcome $(w(\pi(e_H)), e_H)$ by choosing effort e_H. However, by assumption, he preferred outcome $(w(\pi(e_L)), e_L)$ to $(w(\pi(e_H)), e_H)$. Hence, under the postulated revelation mechanism, he will have an incentive to report θ_L than θ_H.

Note that the revelation principle can be used in a variety of contexts. Its contribution is substantial as it allows to approach the derivation of optimal mechanisms in a rather simple way.

We now use the revelation principle to solve our problem. We consider the case in which the manager is infinite risk averse, i.e., the manager accepts the contract only if it assures his at least his outside option in each state.

By restricting to incentive compatible revelation mechanism, the problem of optimal contract is:

$$\max_{w_H, e_H \geq 0, w_L, e_L \geq 0} \lambda[\pi(e_H) - w_H] + (1-\lambda)[\pi(e_L) - w_L]$$

subject to participation constraints, i.e.,

$$\begin{aligned} w_L - g(e_L, \theta_L) &\geq v^{-1}(\bar{u}) \quad (1) \\ w_H - g(e_H, \theta_H) &\geq v^{-1}(\bar{u}) \quad (2) \end{aligned}$$

and subject to incentive compatibility

$$\begin{aligned} w_H - g(e_H, \theta_H) &\geq w_L - g(e_L, \theta_H) \quad (3) \\ w_L - g(e_L, \theta_L) &\geq w_H - g(e_H, \theta_L) \quad (4) \end{aligned}$$

Under constraint 1 a low productivity manager prefers (w_L, e_L) to his outside option; under constraint 2 a high productivity manager prefers (w_H, e_H) to his outside option.

Under constraint 3 a high productivity worker reveals truthfully, given the specified contract; under constraint 4 a low productivity worker reveals truthfully, given the specified contract.

The problem can be solved using standard KT methods. Here we provide a series of lemmas to derive properties of optimal contract.

Lemma 1. If constraints 1 and 3 hold, then constraint 2 holds. Hence, we can drop the participation constraint of a high productivity worker.

Indeed

$$w_H - g(e_H, \theta_H) \geq w_L - g(e_L, \theta_H) \geq w_L - g(e_L, \theta_L) \geq v^{-1}(\bar{u})$$

where the first inequality is the incentive compatible constraint for the high type, the second inequality follows by assumption of g and the third inequality is the participation constraints for the low productivity manager.

Lemma 2. In the solution to our problem the participation constraint of the low productivity manager must be binding, i.e., $w_L - g(e_L, \theta_L) = v^{-1}(\bar{u})$.

Suppose it is not

$$w_L - g(e_L, \theta_L) > v^{-1}(\bar{u})$$

Then, the owner can offer the same contract but with a slightly lower wage, i.e., $w'_L = w_L - \epsilon$ and $w'_H = w_H - \epsilon$. We can find ϵ sufficiently small so that the participation constraint for the low productivity guy still holds, and not that the incentive constraints are not altered. So, the new contract increases the expected profit of the owner, a contradiction to the hypothesis that the contract was optimal.

Lemma In any optimal contract the manager's effort level in state θ_L, e_L is not more than the level that would arise if θ were observable, e^*_L. The manager's effort level in state θ_H, e_H, equals the level that would arise if θ were observable, e_H^*.

The first part follows by noticing that if $e_L > e_L^*$, then the manager can offer the first best point (w_L^*, e_L^*), without altering the utility of the manager.

The second part uses the fact that $e_L \leq e_L^*$ to show that the best contract that the owner can offer which satisfies constraints has effort equals to e_H^* followed by announcement of a high productivity manager.

Lemma In any optimal contract the manager's effort level in state θ_L, e_L is strictly less than the level that would arise if θ were observable, e_L^*.

It can be checked that the optimal level of low effort is such that:

$$[\pi'(e_L] - g_e(e_L, \theta_L)] + \frac{\lambda}{1-\lambda}[g_e(e_L, \theta_H) - g_e(e_L, \theta_L)]$$

Note that the first term is zero at $e_L = e_L^*$ and strictly positive at $e_L < e_L^*$, whereas the second term is always strictly negative.

Heuristically, in order to incentive the high productivity manager to exerts high effort he has to distort the allocation given to low productivity workers. This is done by decreasing effort in the low state.

That is, in determining the low effort level the owner weights the marginal loss in profit in the low state against the marginal gain in the high state.

In particular, the more likelihood is the high state relative to the low state (the higher is $\lambda/(1-\lambda)$), the more the owner is willing to distort the low state (the lower is e_L relative to e_L^*).

Summarizing:

Proposition. In the hidden information principal-agent model with infinite risk averse manager the optimal contract sets the level of effort of high productivity manger to the first best level. The effort level of the low productivity manager is distorted downward relative to the first best. The manager is not fully insured: it receives higher utility in the high state, and his outside option in the low state. In expectation the manager gets his outside option. The expected utility of the owner is lower than under complete information.