5 Principal-Agent Problem

(Dates Covered: February 19 - 28, 2002)

Suppose a restaurant owner (the principal) hires a waiter (the agent) to run the restaurant in his or her absence. If the waiter does not work hard (or "shirks"), fewer customers will come and the owner will lose potential revenue. If the waiter does work hard, more customers will come and the owner will absorb greater revenue than otherwise. However, this additional revenue is as a direct result of the waiter's effort who, depending on the terms of employment, could potentially reap no benefits from the increased effort. Thus, the restauranteur will want to design a contract and offer it to the agent, in the hopes that it will provide incentive to the waiter to work hard, such that both parties will mutually benefit. The preceding is a specific example of the **principal-agent problem**, which can generically be described in the following manner:

The principal designs the terms of the contract and offers it to the agent. The agent decides whether or not to accept the contract. Should the agent accept the contract, he or she decides on the level of effort that will be exerted. The firm's revenue is observed, and the principal pays the agent based on the terms agreed upon in the contract.

<u>Example 1</u>

Let's first look at the problem from the agent's perspective. We will define e, the level of effort exerted by the waiter, as

$$e = \begin{cases} 0 & \text{if } shirks \\ 2 & \text{if } workshard \end{cases}$$

In addition, we define the agent's utility as

$$U = \begin{cases} w - e & \text{if } devotes\text{e} \\ 10 & \text{if } rejectscontract \end{cases}$$

where w is the wage offered in the contract. Thus, the agent will either "take it or leave it." From the principal's perspective, we define $R(e)$ as the revenue for the restaurant as a function of the waiter's effort:

$$R(e) = \begin{cases} H & \text{if e= 2} \\ L & \text{if e= 0} \end{cases}$$

The potential profit for the principal is

$$\pi = R(e) - w$$

Of course, the principal will design the contract with the objective of maximizing π. We want to define w^H as the wage rate the principal promises to pay when the effort is high and w^L as the wage rate that the agent will receive if the effort is low.

We will finish this example during the next lecture.

(End of 2-19 Lecture. Start of 2-21 Lecture.)

Now we will finish the example we started during the previous lecture.

The principal will design a contract offer two wage rates: one, w^H, when the agent exerts high effort and another, w^L, when the agent exerts low effort. First, the principal will design the **participation contraint**

$$w^H - 2 \geq 10,$$

to ensure that the agent participates (agrees to accept the contract). This constraint is designed such that the difference between the high wage rate and maximal effort is greater than or equal to any other alternative that the waiter has. In short, even if the agent has to exert maximal effort, it would be in his or her best interest to accept the principal's contract offer.

In addition to wanting the agent to accept the contract, the principal also wants the agent to agree to work hard. Recall our assumption from the previous lecture that profits generated when the agent works hard (π_H) are greater than profits generated when the agent "shirks" (π_L) since customers are more likely to be repeat visitors when there is high customer service. Thus, the principal will also construct an **incentive constraint**

$$w^H - 2 \geq w^L,$$

which gives the agent the incentive to work hard. This constraint is designed such that the difference between the high wage rate and the maximal effort is greater than or equal to the low wage rate. In short, the high wage rate is set such that the agent is better off by accepting it **and** working hard as opposed to accepting the low wage rate and shirking.

The desired combined effect of these two contraints is to produce a contract such that the agent will both accept the principal's offer and agree to work hard. In this example, the participation and incentive constraints generate a high wage rate w^H of 12 and a low wage rate w^L of 10. The corresponding profit functions are

$$\pi^H = H - w^H = H - 12$$
$$\pi^L = L - w^L = L - 10,$$

where H and L are the revenues associated with the agent working hard and "shirking" respectively.

The preceding principal-agent problem also illustrates a concept called moral hazard. *A* **moral hazard** *occurs when the agent takes a decision or action that affects his or her utility as well as the principal's, the principal only observes the "outcome" (an imperfect signal of the action taken), and the agent (given a spontaneous choice) does not necessarily choose Pareto optimality.* Since the action is unobservable, the principal cannot force the agent to choose an action that is Pareto optimal. He or she can only influence the choice of an action by conditioning the agent's utility to the outcome, which is the only observable variable. In our example, the principal does not observe whether or not the agent is working hard or "shirking" (action). The only thing that the principal observes is the revenue generated as a result of the agent's effort. Thus, it is possible that the agent will accept the high wage, agree to work hard, but "shirk." The principal will not know this until he or she observes the revenue figures and calculates profit (the outcome). Hence, there is a "moral" hazard associated with these sorts of contracts. (Moral hazards essentially occur in all economic relationships.)

Some potential complications to the principal-agent problem include asymettric information, risk differences, and revenue uncertainty. We will now look at an example that has the latter complication.

Example 2

Suppose we have the same parameters as in Example 1, only the revenues H and L are uncertain – as indicated by the following "states of nature:"

$$R(2) = \begin{cases} H & \text{with probability } 0.8 \\ L & \text{with probability } 0.2 \end{cases}$$

and

$$R(0) = \begin{cases} H & \text{with probability } 0.4 \\ L & \text{with probability } 0.6 \end{cases}$$

Here, the agent affects the probability of each realization of $R(e)$ depending on the level of effort chosen. As a result, the utility function of the agent must be modified, and is now

$$U = \begin{cases} E[w] - e & \text{if devotes e} \\ 10 & \text{if rejects contract} \end{cases}$$

The wage is now an expected value $E[w]$ because of the uncertainties associated with the revenue function. Based on the revenue functions,

$$E[w] = \begin{cases} 0.8w^H + 0.2w^L & \text{when } e = 2 \\ 0.4w^H + 0.6w^L & \text{when } e = 0 \end{cases}$$

Thus, the new participation constraint is

$$0.8w^H + 0.2w^L - 2 \geq 10,$$

and the new incentive constraint is

$$0.8w^H + 0.2w^L - 2 \geq 0.4w^H + 0.6w^L.$$

Rewriting the participation constraint in terms of w^L yields

$$w^L = 60 - 4w^H,$$

and rewriting the incentive constraint in terms of w^L yields

$$w^L = w^H - 5.$$

Solving this as a system of equations yields $w^L = 8$, and then, using back substitution, we obtain $w^H = 13$. The uncertainty results in the agent incurring more risk which, in turn, generates a higher w^H value (Example 1 $w^H = 12$).

For homework, model this as an extensive form game, and solve SPE.

5.1 Risk Aversion

Suppose there are two consumers i and j. *Consumer i is more* **risk averse** *than consumer j if, when consumer j prefers a fixed sum of money over a lottery, then consumer i also prefers the fixed amount.* This is explained using the following example.

Example 3

Suppose that the principal and agent think about risk differently. One way to model this is the use of "subjective probability" that measures the probabilities that each player assigns to the realization of the states of nature. Suppose the principal has the same revenue functions as above:

$$R^P(2) = \begin{cases} H & \text{with probability } 0.8 \\ L & \text{with probability } 0.2 \end{cases}$$

and

$$R^P(0) = \begin{cases} H & \text{with probability } 0.4 \\ L & \text{with probability } 0.6 \end{cases}$$

The agent, however, has a different belief about how much revenue will be generated as a result of his or her effort. The agent believes that the principal's revenue functions will be as follows:

$$R^A(2) = \begin{cases} H & \text{with probability } 0.7 \\ L & \text{with probability } 0.3 \end{cases}$$

and

$$R^A(0) = R^P(0).$$

In this case, the agent is more risk averse since he or she is more skeptical about the realization of the "high" state of nature. In fact, it is typically assumed that the principal is "risk neutral" (i.e. less risk averse) and the agent is risk averse.

Thus, in this case, the participation constraint will be

$$0.7w^H + 0.3w^L - 2 \geq 10 \rightarrow w^H = \tfrac{12-0.3w^L}{0.7},$$

and the incentive constraint will be

$$0.7w^H + 0.3w^L - 2 \geq 0.4w^H + 0.6w^L \rightarrow w^H = \tfrac{2}{0.3} + w^L.$$

The objective of the principal is to minimize expected wage if $H - L$ is sufficiently large:

$$\min_{w^H, w^L} E^0[w] = 0.8w^H + 0.2w^L,$$

which yields the optimum values $w^H = 14$ and $w^L = \frac{22}{3}$. The principal's expected wage will exceed the agent's reservation utility plus effort:

$$E^0[w] = 0.8w^H + 0.2w^L = 12.66 > 10 + 12 = 12.$$

Plotting all of this yields the following graphic:

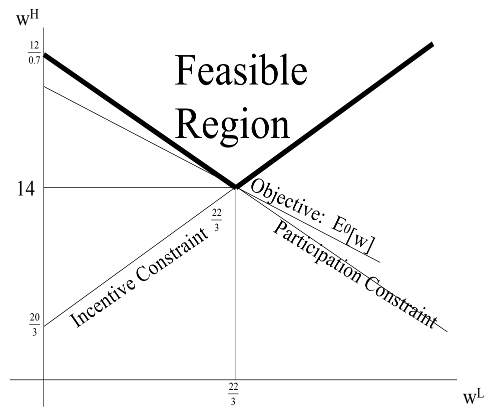

Figure 21: Risk Averse Principal-Agent Example

The following list of items is a general summary of a principal-agent model:

1. The agent chooses between n possible actions: $a_1, \ldots, a_n$.

2. These actions produce m possible outcomes: $x_1, \ldots, x_m$. (The outcome is a signal that gives the principal information on the action that the agent chooses.)

3. There is a stochastic relationship between actions and outcomes (called a "technology") such that when the agent chooses action a_i, the principal observes outcomes x_j with probability p_{ij}.

4. If the principal observes outcome x_j, he or she will pay the agent a wage w_j and keep the remaining $(x_j - w_j)$.

5. Assuming the agent's utility, $U(w) - a$, is separable, U is increasing and concave.

6. Assuming the principal is risk neutral, his or her utility is $x - w$.

First, we look at the agent's problem. When the principal offers a contract w_j, the agent chooses his or her best action by solving:

$$\max_{i=1,\ldots,n}[\sum_{j=1}^{m} P_{ij}U(w_j) - a_i].$$

If the agent chooses a_i, this gives $n-1$ incentive compatibility constraints (IC_k)

$$\sum_{j=1}^{m} P_{ij}U(w_j) - a_i \geq \sum_{j=1}^{m} P_{kj}U(w_j) - a_k \ \forall k \neq i.$$

Suppose $\underline{U}$ is the utility for the agent if he or she does not accept the contract. Then, the participation constraint (known, in this case, as an individual rationality constraint (IR)) is

$$\sum_{j=1}^{m} P_{ij}U(w_j) - a_i \geq \underline{U}.$$

Now, we look at the principal's problem. The principal wishes to choose a contract $(w_1, \ldots, w_m)$ that maximizes the expected utility, while taking into account the dependence of the contract on the agent's decisions. Thus, the objective function is

$$\max_{w_1,\ldots,w_m,i} \sum_{j=1}^{m} P_{ij}(x_j - w_j),$$

subject to $IC_k(\lambda)$ and $IR(\mu)$ from the agent's problem.

(End of 2-21 Lecture. Start of 2-26 Lecture.)

If we fix a_i, the Lagrangian of this maximization problem is

$$L(w, \lambda, \mu) = \sum_{j=1}^{m} P_{ij}(x_j - w_j) + \sum_{k=1;k\neq i} \lambda_k(\sum_{j=1}^{m} P_{ij}U(w_j) - a_i - \sum_{j=1}^{m} P_{kj}U(w_j) + a_k) + \mu(\sum_{j=1}^{m} P_{ij}U(w_j) - a_i - \underline{U}).$$

Taking the first order condition with respect to w_j, we get

$$\frac{\partial L}{\partial w_j} = -P_{ij} + \sum_{k=1;k\neq i}^{n} \lambda_k(P_{ij}U'(w_j) - P_{kj}U'(w_j)) + \mu P_{ij}U'(w_j) \equiv 0$$
$$\rightarrow \frac{1}{U'(w_j)} = \mu + \sum_{k=1;k\neq i}^{n} \lambda_k(1 - \frac{P_{kj}}{P_{ij}}),$$

which makes w_j as a function of j depend on the ratio $\frac{P_{kj}}{P_{ij}}$. (In equilibrium, at least one of the λ_k values must be positive; otherwise, we could neglect the incentive constraints, and the moral hazard problem would vanish.) Thus, if the agent has only two possible actions, as in the previous examples, our expression simplifies to

$$\frac{1}{U'(w_j)} = \mu + \lambda[1 - \frac{P_{1j}}{P_{2j}}].$$

(Note that, in statistics, an expression of the form $\frac{P_{aj}}{P_{bj}}$ is known as a **likelihood ratio,** *which measures the likelihood of observing x_j given that the agent chooses a to the likelihood of observing x_j given that the agent chooses b.* Thus, the higher the value of the ratio, the more likely the agent will choose a – and vice-versa.)

Recall that we said $U(w)$ is increasing and concave. This indicates that the ratio $\frac{1}{U'(w_j)}$ is increasing, and that w_j will tend to be higher as the ratio increases.

This is illustrated by the following graphic:

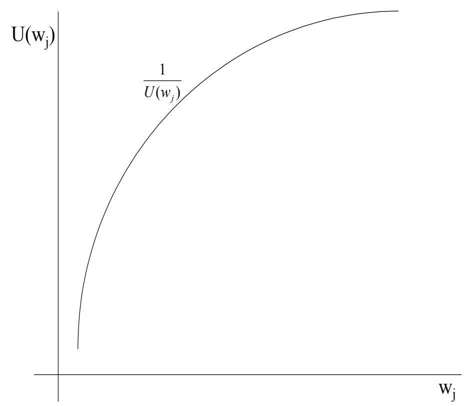

Figure 22: As likelihood ratio deceases, left-hand side increases in w_j

Since this is a monotone increasing function of x_j, the optimal incentive scheme is a linear function of the likelihood ratio. However, if the number of actions $n > 2$, the argument will not necessarily hold since wage depends on a weighted average of likelihood ratios. (It should be noted, however, that the intuition *will* hold in most cases.) In order to make the argument more concrete, we need to assume that the cumulative distribution function (cdf) is convex on actions $\{a_1, \ldots, a_n\}$. That is, for $i < j < k$ and $\lambda \in [0, 1]$,

$$a_j = a_i\lambda + (1 - \lambda)a_k.$$

This condition gives

$$P_{jl} \leq \lambda P_{il} + (1 - \lambda)P_{kl} \ \forall l = 1, \ldots, m.$$

An example of how to construct an incentive scheme for two actions and two outcomes is as follows:

$$\begin{cases} w_1 = w \\ w_2 = w + s(x_2 - x_1)(s \leq 1) \end{cases}$$

Here, the agent receives a base wage w and a bonus proportional to the increase in surplus when he or she suceeded in satisfying the principal's demands.

Using the Lagrangian statistics mentioned earlier, we can determine how the principal's optimizaed values change with respect to the agent's effort in the two action, two outcome case.

The Lagrangian first order conditions with respect to the agent's costs (effort) yield

$$\frac{\partial L}{\partial a_1} = \lambda,$$

and

$$\frac{\partial L}{\partial a_2} = -(\lambda + \mu),$$

based on the Envelope Theorem, which was outlined extensively in ISyE 6229.

This helps answer the question, "Which is better, the "carrot" or the "stick?". The "carrot" provides incentive to the agent to perform the action the principal wants by decreasing the cost (effort) of the chosen action a_2. The "stick" punishes the agent for not taking the action desired by the principal by increasing the cost of the alternative action a_1 by the same magnitude. As evidenced by the first order conditions, the "carrot" relaxes two constraints (IC and IR), while the stick only relaxes one (IC). Hence, a small decrease in the cost of the chosen action always increases the principal's utiity by a larger amount than an increase of the same magnitude in cost of the alternative action. In other words, *the "carrot" is always better than the "stick"*.

Suppose now we have a change in the probability distribution dP_{aj}. The effect on the principal's utility of such a change is

$$dL = -\lambda \sum_{j=1}^{m} dP_{aj} U(w_j).$$

Thus, when the IC constraint is binding ($\lambda > 0$), the interests of the principal and agent are diametrically opposed with respect to changes in the probability distribution. In other words, any action that helps the agent unambiguously hurts the principal.

5.2 Multiple Agents

We will conclude our discussion on the Principal-Agent problem with a look at the case of multiple agents. We will do so by looking at the following example.

Example 4

Consider an order picking operation in a warehouse with N pickers. The value V to the warehouse depends on the effort levels of the N pickers

$$V = \sum_{i=1}^{N} \sqrt{e_i}.$$

To simplify the analysis, assume that

$$\sum_{i=1}^{N} w_i = V,$$

where w_i is the wage for picker i. Furthermore, assume all pickers have the same utility

$$U_i = w_i - e_i \; i = 1, \ldots, N.$$

Suppose the warehouse manager equally divides the work. Then $w_i = \frac{V}{N}$. To determine NE, each worker takes the effort levels of his or her colleagues as given, and chooses his or her effort to maximize the utility. Formally, each worker chooses e_i to

$$\max_{e_i} U_i = \frac{\sum j \neq i \sqrt{e_j} + \sqrt{e_i}}{N} - e_i,$$

where $\frac{\sum j \neq i \sqrt{e_j} + \sqrt{e_i}}{N}$ is w_i. The corresponding first order condition is

$$\frac{\partial U}{\partial e_i} = \frac{e_i^{-\frac{1}{2}}}{2N} - 1 \equiv 0 \rightarrow e_i = \frac{1}{4N^2}.$$

As an addendum to the problem, if workers collude to maximize utility and each picker can observe the effort of others, the set $e_i = e$, and

$$\max_e U = w - e = \frac{V}{N} - e = \frac{N\sqrt{e}}{N} - e.$$

The first order conditions yield

$$V = \sum_{i=1}^{N} \sqrt{e_i} = N\sqrt{\frac{1}{4}} = \frac{N}{2},$$

and $e^* = \frac{1}{4}$.

From these results, we have learned that, if the "team" is a single worker, then that worker will provide an optimal level of effort. If there is more than one worker on a team, then each worker will devote less than the optimal level of effort. Thus, the larger the team is, the greater the incentive is to "shirk."

This will affect the welfare of both the principal and the agent. For the principal,

$$V^* - V^{NASH} = \frac{N}{2} - \frac{1}{2} = \frac{(N-1)}{2},$$

and the difference increases as N increases. For the agent (picker),

$$U_i = w_i - e_i = \frac{V^{NASH}}{N} - \frac{1}{4N^2} = \frac{1}{2N} - \frac{1}{4N^2},$$

and the first order condition, assuming N is continuous, is

$$\frac{\partial U_i}{\partial N} - \frac{-1}{2N^2} + \frac{1}{2N^3}(< 0 if N > 1),$$

and the agent's utility decreases in N.