9 Adverse Selection

This is designed for one 75-minute lecture using *Games and Information*. Probably I have more material than I will end up covering.

This is just for sections 9.1 and 9.6.

October 7, 2006

Production Game VI: Adverse Selection

Players

The principal and the agent.

The Order of Play

(0) Nature chooses the agent's ability a, observed by the agent but not by the principal, according to distribution $F(a)$.

(1) The principal offers the agent one or more wage contracts $w_1(q), w_2(q), \ldots$

(2) The agent accepts one contract or rejects them all.

(3) Nature chooses a value for the state of the world, θ, according to distribution $G(\theta)$. Output is then $q = q(a, \theta)$.

Payoffs

If the agent rejects all contracts, then $\pi_{agent} = \overline{U}(a)$, which might or might not vary with his type, a; and $\pi_{principal} = 0$.
Otherwise, $\pi_{agent} = U(w, a)$ and $\pi_{principal} = V(q - w)$.

Under adverse selection, it is not the worker's effort, but his ability, that is noncontractible.

No uncertainty—- Either high or low output might be observed in equilibrium, unlike under moral hazard.

Offering multiple contracts can be an improvement over offering a single contract– perhaps a flat-wage contract for low-ability agents and an incentive contract for high-ability agents.

Production Game VIa puts specific functional forms into the game to illustrate how to find an equilibrium.

Production Game VIa: Adverse Selection with Particular Parameters

Players

The principal and the agent.

The Order of Play

(0) Nature chooses the agent's ability a, unobserved by the principal, according to distribution $F(a)$, which puts probability 0.9 on low ability, $a = 0$, and probability 0.1 on high ability, $a = 10$.

(1) The principal offers the agent one or more wage contracts $W_1 = (w_1(q = 0), w_1(q = 10)), W_2 = (w_2(q = 0), w_2(q = 10)) \ldots$

(2) The agent accepts one contract or rejects them all.

(3) Nature chooses a value for the state of the world, θ, according to distribution $G(\theta)$, which puts equal weight on 0 and 10. Output is then $q = Min(a + \theta, 10)$.

Payoffs

If the agent rejects all contracts, then depending on his type his reservation payoff is either $\overline{U}_{Low} = 3$ or $\overline{U}_{High} = 2$ and the principal's payoff is $\pi_{principal} = 0$.
Otherwise, $U_{agent} = w$ and $V_{principal} = q - w$.

A separating equilibrium is

Principal: Offer $W_1 = \{w_1(q=0) = 3, w_1(q=10) = 3\}$,
$W_2 = \{w_2(q=0) = 0, w_2(q=10) = 3\}$

Low agent: Accept W_1

High agent: Accept W_2

As usual, this is a weak equilibrium. Both Low and High agents are indifferent about whether they accept or reject their contract. The equilibrium indifference of the agents arises from the open-set problem; if the principal were to specify a wage of 3.01 for W_2, for example, the high- ability agent would no longer be indifferent about accepting it instead of W_1.

In hidden-action models, the principal tries to construct a contract which will induce the agent to take the single appropriate action. In hidden-knowledge models, the principal tries to make different actions attractive to different types of agent, so the agent's choice depends on the hidden information.

(1) **Incentive compatibility** (the agent picks the desired contract and actions).

(2) **Participation** (the agent prefers the contract to his reservation utility).

In a model with hidden knowledge, the incentive compatibility constraint is customarily called the **self-selection constraint**.

There can be one IC constraint and one Part. constraint for each type of agent.

Here, what action does the principal desire from each type of agent?

The agents do not choose effort, but they do choose whether or not to work for the principal, and which contract to accept.

The low-ability agent's expected output is 0.5(0) + 0.5(10)= 5, compared to a reservation payoff of 3, so the principal will want to hire him if $Ew \leq 5$.

The high-ability agent's expected output is 0.5(10) + 0.5(10)= 10, compared to a reservation payoff of 2, so the principal will want to hire the high-ability agent if $Ew \leq 5$.

The participation constraints are, if we let $\pi_i(W_j)$ denote the expected payoff an agent of type i gets from contract j,

$$\begin{aligned} \pi_L(W_1) &\geq \overline{U}_{Low}; \;\; 0.5w_1(0) + 0.5w_1(10) \geq 3 \\ \pi_H(W_2) &\geq \overline{U}_{High}; \;\; 0.5w_2(10) + 0.5w_2(10) \geq 2. \end{aligned} \tag{1}$$

Clearly the contracts in our conjectured equilibrium, $W_1 = (3,3)$ and $W_2 = (0,3)$, satisfy the participation constraints. In the equilibrium, the low- and the high-output wages both matter to the low-ability agent, but only the high-output wage matters to the high-ability agent. Both agents, however, end up earning a wage of 3 in each state of the world, the only difference being that contract W_2 would be a very risky contract for the low-ability agent despite being riskless for the high-ability agent. principal would like to make W_1 risk-free, with the same wage in each state of the world.

In our separating equilibrium, the participation constraint is binding for the "bad" type but not for the "good" type.

This is typical of adverse selection models (if there are more than two types it is the participation constraint of the worst type that is binding, and no other).

The participation constraint is binding for the "bad" type but not for the "good" type.

The principal makes the bad type's contract unattractive for two reasons. First, if he pays less, he keeps more.

Second, when the bad type's contract is less attractive, the good type can be more cheaply lured away to a different contract. The principal can never extract all the gains from trade from the good type unless he gives up on making either of his contracts acceptable to the bad type.

Another typical feature of this equilibrium is that the low-ability agent's contract not only drives him down to his participation constraint, but is riskless.

The self-selection constraints are

$$\pi_L(W_1) \geq \pi_L(W_2); \quad 0.5w_1(0) + 0.5w_1(10) \geq 0.5w_2(0) + 0.5w_2(10)$$

$$\pi_H(W_2) \geq \pi_H(W_1); \quad 0.5w_2(10) + 0.5w_2(10) \geq 0.5w_1(10) + 0.5w_1(10) \tag{2}$$

The first inequality in (2) says that the contract W_2 has to have a low enough expected return for the low-ability agent to deter him from accepting it. The second inequality says that the wage contract W_1 must be less attractive than W_2 to the high-ability agent. The conjectured equilibrium contracts $W_1 = (3, 3)$ and $W_2 = (0, 3)$ do this, as can be seen by substituting their values into the constraints:

$$\begin{aligned} \pi_L(W_1) \geq \pi_L(W_2); &\quad 0.5(3) + 0.5(3) \geq 0.5(0) + 0.5(3) \\ \pi_H(W_2) \geq \pi_H(W_1); &\quad 0.5(3) + 0.5(3) \geq 0.5(3) + 0.5(3) \end{aligned} \tag{3}$$

The self-selection constraint is binding for the good type but not for the bad type.

This, too, is typical of adverse selection models.

The principal wants the good type to reveal his type by choosing the appropriate to the good type as the bad type's contract. It does not have to be *more* attractive though (here notice the open-set problem), so the principal will minimize his salary expenditures and choose two contracts equally attractive to the good type. In so doing, however, the principal will have chosen a contract for the good type that is strictly worse for the bad type, who cannot achieve so high an output so easily.

All that remains to check is whether the principal could increase his payoff.

He cannot, because he makes a profit from either contract, and having driven the low- ability agent down to his reservation payoff and the high-ability agent down to the minimum payoff needed to achieve separation, he cannot further reduce their pay.

Competition and Pooling

Although it is true, however, that the participation constraints must be satisfied for agents who accept the contracts, it is not always the case that they accept different contracts in equilibrium.

If they do not, they do not need to satisfy self-selection constraints.

If all types of agents choose the same strategy in all states, the equilibrium is **pooling***. Otherwise, it is* **separating***.*

The distinction between pooling and separating is different from the distinction between equilibrium concepts.

A model might have multiple Nash equilibria, some pooling and some separating.

Moreover, a single equilibrium— even a pooling one— can include several contracts, but if it is pooling the agent always uses the same strategy, regardless of type.

If the agent's equilibrium strategy is mixed, the equilibrium is pooling if the agent always picks the same mixed strategy, even though the messages and efforts would differ across realizations of the game.

The possibility of a pooling equilibrium reveals one more step we need to take to establish that the proposed separating equilibrium in Production Game VIa is really an equilibrium:

Would the principal do better by offering a pooling contract instead, or a separating contract under which one type of agent does not participate?

All of my derivation above was to show that the agents would not deviate from the proposed equilibrium, but it might still be that the principal would deviate.

First, would the principal prefer pooling?

Then all that is necessary is that the contract as cheaply as possible induce both types of agent to participate.

Here, that would require that we make the contract barely acceptable to the type with the lowest ability and highest reservation payoff, the low-ability agent.

The contract $(3, 3)$ offered by itself would do that, but it would not increase profits over W_1 and W_2 in our equilibirum above.

Either pooling or separating would yield profits of $0.9(0.5(0-3)+0.5(10-3))+0.1(0.5(10-3)+0.5(10-3))=2.5$.

Second, would the principal prefer a separating contract that "gave up" on one type of agent?

The principal would not want to drive away the high-ability agent, of course, though he could do so by offering a high wage for $q = 0$ and a low wage for $q = 10$, because the high-ability agent has both greater output and a lower reservation payoff (if we had $\overline{U}_{High} = 11$ then the outcome would be different).

But if the principal did not have to offer a contract that gave the low-ability agent his reservation payoff of 3, he could be more stingy towards the high-ability agent.

If there were no low-ability agent, the principal would offer a contract such as $(0, 2)$ to the high-ability agent, driving him down to his reservation payoff and increasing the profits from hiring him.

Here, however, there are not enough high-ability agents for that to be a good strategy for the principal.

His payoff would decline to $0.9(0) + 0.1(0.5(10 - 2) + 0.5(10 - 2)) = 0.8$, a big decline from 2.5.

If 99% of the agents were high-ability, instead of 10%, things would have turned out differently.

Production Game VII: Adverse Selection and Moral Hazard COMBINED

Players

The principal and the agent.

The Order of Play

(0) Nature chooses the state of the world s, observed by the agent but not by the principal, according to distribution $F(s)$, where the state s is Good with probability 0.5 and Bad with probability 0.5.

(1) The principal offers the agent a wage contract $w(q)$.

(2) The agent accepts or rejects the contract.

(3) The agent chooses effort level e.

(4) Output is $q = q(e, s)$. where $q(e, good) = 3e$ and $q(e, bad) = e$.

Payoffs

If the agent rejects all contracts, then $\pi_{agent} = \overline{U} = 0$ and $\pi_{principal} = 0$.
Otherwise, $\pi_{agent} = U(e, w, s) = w - e^2$ and $\pi_{principal} = V(q - w) = q - w$.

Thus, there is no uncertainty, both principal and agent are risk neutral in money, and effort is increasingly costly.

In this model, the first-best effort depends on the state of the world. The two social surplus maximization problems are

$$\underset{e_g}{\overset{Maximize}{}} \quad 3e_g - e_g^2, \tag{4}$$

which is solved by the optimal effort $e_g = 1.5$ (and $q_g = 4.5$) in the good state, and

$$\underset{e_g}{\overset{Maximize}{}} \quad e_b - e_b^2, \tag{5}$$

which is solved by the optimal effort $e_b = 0.5$ (and $q_b = 0.5$)in the bad state.

The problem is that the principal does not know what level of effort and output are appropriate.

He does not want to require high output in both states, because if he does, he will have to pay too high a salary to the agent to compensate for the difficulty of attaining that output in the bad state.

Rather, he must solve the following problem:

$$\overset{Maximize}{q_g, q_b, w_g, w_b} \quad [0.5(q_g - w_g) + 0.5(q_b - w_b)], \tag{6}$$

where the agent has a choice between two forcing contracts, (q_g, w_g) and (q_b, w_b), and the contracts must induce participation and self selection.

The self-selection constraints are based on efforts of $e = q/3$ for the good state and $e = q$ for the bad state.

In the good state, the agent must choose the good-state contract,so

$$\pi_{agent}(q_g, w_g|good) = w_g - \left(\frac{q_g}{3}\right)^2 \geq \pi_{agent}(q_b, w_b|good) = w_b - \left(\frac{q_b}{3}\right) \tag{7}$$

and in the bad state he must choose the bad-state contract,so

$$\pi_{agent}(q_b, w_b|bad) = w_b - q_b^2 \geq \pi_{agent}(q_g, w_g|bad) = w_g - q_g^2. \tag{8}$$

The participation constraints are

$$\pi_{agent}(q_g, w_g|good) = w_g - \left(\frac{q_g}{3}\right)^2 \geq 0 \tag{9}$$

and

$$\pi_{agent}(q_b, w_b|bad) = w_b - q_b^2 \geq 0. \tag{10}$$

The bad state's participation constraint will be binding, since in the bad state the agent will not be tempted by the good-state contract's higher output and wage. Thus, we can conclude from constraint (10) that

$$w_b = q_b^2. \tag{11}$$

The good state's participation constraint will not be binding, since there the agent will be left with an informational rent– the principal must leave the agent some surplus to induce him to reveal the good state.

The good state's self-selection constraint will be binding, since in the good state the agent *will* be tempted to take the easier contract appropriate for the bad state. Thus, we can conclude from constraint (7) that

$$\begin{aligned} w_g &= \left(\frac{q_g}{3}\right)^2 + w_b - \left(\frac{q_b}{3}\right)^2 \\ &= \left(\frac{q_g}{3}\right)^2 + q_b^2 - \left(\frac{q_b}{3}\right)^2, \end{aligned} \tag{12}$$

where the second step substitutes for w_b from equation (11). The bad state's self-selection constraint will not be binding, since the agent would then not be tempted to produce a large amount for a large wage.

Now let's return to the principal's maximization problem. Having found expressions for w_b and w_g we can rewrite (6) as

$$\underset{q_g,\, q_b}{Maximize} \quad [0.5\left(q_g - \left(\frac{q_g}{3}\right)^2 - q_b^2 + \left(\frac{q_b}{3}\right)^2\right) + 0.5(q_b - q_b^2)] \tag{13}$$

with no constraints. The first-order conditions are

$$0.5\left(1 - \frac{2q_g}{9}\right) = 0, \tag{14}$$

so $q_g = 4.5$, and

$$0.5\left(-2q_b + \frac{2q_b}{9}\right) + 0.5(1 - 2q_b) = 0, \tag{15}$$

so $q_b \approx .26$. We can then find the wages that satisfy the constraints, which are $w_g \approx 2.32$ and $w_b \approx 0.07$.

Thus, in the second-best world of information asymmetry, the effort in the good state remains the first-best effort, but second-best effort in the bad state is lower than first-best. This results from the principal's need to keep the bad- state contract from being too attractive in the good state. Bad-state output and compensation must be suppressed. Good-state output, on the other hand, should be left at the first-best level, since the agent will not be tempted by that contract in the bad state.

Also, observe that in the good state the agent earns an informational rent. As explained earlier, this is because the good-state agent could always earn a positive payoff by pretending the state was bad and taking that contract, so any contract that separates out the good-state agent (while leaving some contract acceptable to the bad-state agent) must also have a positive payoff.