Ecole Polytechnique
ECO 560 Economics of Contracts

Raicho Bojilov
25 January 2013

Lecture 3: Adverse Selection, Extensions; Review of Concepts in Game Theory

1 General Model: Review

Consider a Principal who faces a continuum of types $\theta \in \left[\underline{\theta}, \overline{\theta}\right]$ with density $f(\theta)$. The buyer's utility of type θ from declaring type θ' is $U(\theta, \theta') = \theta q(\theta') - t(\theta')$. This utility function satisfies the so called single crossing property

$$\frac{\partial^2 U}{\partial q \partial \theta} \cdot \frac{\partial q}{\partial \theta} > 0$$

Alternative notation for this derivative that we will use for brevity is U_{12}. The single-crossing property is must be satisfied when we use a more general formulation of the Agent's utility. Let $q(\theta)$ be the quality purchased by θ and $t(\theta)$ be the transfer made from the Agent θ to the Principal. The Principal's cost is C(q), where C(q) is strictly convex, twice differentiable, C'(0) = 0 and C'(∞) = ∞ The Principal's profits from serving a customer θ are $t(\theta) - C(q(\theta))$.

The problem of the Principal is:

$$\max_{t(\theta), q(\theta)} \int_{\underline{\theta}}^{\overline{\theta}} (t(\theta) - C(q(\theta))) f(\theta) d\theta$$

s.t. for all θ, θ'

$$\theta q(\theta) - t(\theta) \geq \theta q(\theta') - t(\theta') \quad (IC)$$

$$\theta q(\theta) - t(\theta) \geq 0 \quad (IR)$$

We have shown that this problem is equivalent to:

$$\max_{t(\theta), q(\theta)} \int_{\underline{\theta}}^{\overline{\theta}} (t(\theta) - C(q(\theta))) f(\theta) d\theta$$

s.t.

$$\theta q(\underline{\theta}) - t(\underline{\theta}) \geq 0 \quad (IR(\underline{\theta}))$$

$$U(\theta, \theta) = \underline{\theta} q(\underline{\theta}) - t(\underline{\theta}) + \int_{\underline{\theta}}^{\theta} q(x) dx \quad (IC)$$

$$q(\theta) \text{ increasing}$$

The objective function can be restated as:

$$\int_{\underline{\theta}}^{\overline{\theta}} \left[\underbrace{\left(\theta - \frac{1 - F(\theta)}{f(\theta)} \right)}_{virtual\ type} q(\theta) - C(q(\theta)) \right] f(\theta) d\theta$$

which leads to very simple first-order conditions:

$$\left(\theta - \frac{1 - F(\theta)}{f(\theta)} \right) = C'(q(\theta))$$

1.1 Properties of the solution:

- The high type gets an efficient allocation of the commodity.
- All other types get subefficient amounts of the commodity.
- All types, except for the lowest type, get positive information rents, which increases with the type.
- The lowest type gets a zero surplus.
- In general little can be said about the shape of $T(q)$ schedule.

1.2 Ironing and Bunching

Since C is convex, $q(\theta)$ is increasing if and only if $\left(\theta - \frac{1-F(\theta)}{f(\theta)}\right)$ is increasing, which is the case if the hazard rate $\frac{f(\theta)}{1-F(\theta)}$ is increasing. The hazard rate is increasing for many probability distributions: normal, uniform, Pareto, logistic, exponential, etc. When the condition that q(θ) is increasing is violated, it becomes necessary to "iron" the schedule $q(\theta)$ by assigning the same contract to a set of individuals also known as bunching. Sometimes bunching provides the basis of empirical tests of the theory. For a rigorous treatment of the subjects, please see pp. 88-93 in Bolton and Dewatripont (2005).

2 Optimal Deterministic Mechanism

In the analysis above, we have considered only deterministic contracts based on one's type. However, under the assumptions of the model stated above, the optimal contract is deterministic. is strictly concave because $C(q(\theta))$ is strictly convex

Proposition 1: *Under the assumptions of the model, the solution of the Principal's problem is deterministic and unique.*

Proof. Suppose in contradiction that there is a stochastic solution $(Q(\theta), T(\theta))$ defined by a probability distribution g over $(q(\theta), t(\theta))$and let $(E(Q(\theta)), E(T(\theta)))$ be the deterministic contract obtained by taking expectations of the random terms in the optimal solution.The objective function

$$\left(\theta - \frac{1-F(\theta)}{f(\theta)}\right) q(\theta) - C(q(\theta))$$

is strictly concave because C($q(\theta)$) is strictly convex. By Jensen's inequality,

$$\int_{\underline{\theta}}^{\overline{\theta}} E_g \left[\left(\theta - \frac{1-F(\theta)}{f(\theta)}\right) q(\theta) - C(q(\theta))\right] f(\theta)\, d\theta$$

$$\leq \int_{\underline{\theta}}^{\overline{\theta}} \left[\left(\theta - \frac{1-F(\theta)}{f(\theta)}\right) E(Q(\theta)) - C(E(Q(\theta)))\right] f(\theta)\, d\theta$$

which is a contradiction with the assumption that $(Q(\theta), T(\theta))$ was optimal. ■

3 Perfect Competition in Contracts

Observe that profits are positive for all types in both the two-type model and the model with a continuum of agents. For simplicity, here we consider only the two-type model. Recall that the Principal extracts all the surplus from the low type:

$$\pi_1 = \theta_1 q_1 - C(q_1)$$

and the profits on the high type are:

$$\pi_2 = \theta_2 (q_2 - q_1) + \theta_1 q_1 - C(q_2)$$

Therefore:

$$\begin{aligned} & \pi_1 - \pi_2 \\ = \; & \theta_2 (q_2 - q_1) - (C(q_2) - C(q_1)) \\ = \; & \int_{q_1}^{q_2} (\theta_2 - C'(q))\, dq > 0 \end{aligned}$$

Since $\theta_2 \geq C'(q)$ for all $q \in [0, \theta_2]$. Thus, the profits on the high type are greater than the profits on the low type. Now suppose that new Principals may enter freely in the market (i.e. there are no entry costs). The new Principals will be willing to enter as long as they make non-zero profits. For example, the new entrant may offer just slightly better conditions to θ_2 than the incumbent and, as a result, capture the high end of the market. The incumbent is, however, also likely to respond to this threat. The two will Bertrand compete driving their profits on the high type to 0. Thus, we observe that allowing for free entry of new Principals reduces the profits of the incumbent. In what follows, we investigate the magnitude of the profit reduction.

For the purposes of answering this question, we have to use an equilibrium concept. We follow the Rothschild -Stiglitz equilibrium defined below.

Definition. Let a profitable entrant be a contract that makes positive profits if all other existing contracts remain unchanged. Then, the Rothschild -Stiglitz equilibrium is defined as the set of contracts, such as there is no profitable entrant and no contract leads to losses.

That is, the Rothschild -Stiglitz equilibrium is a Nash equilibrium of a game played by the Principals, where the strategies of the players are contracts.

Proposition 2: The Irrelevance Result *With free entry of Principals, the set of Rothschild -Stiglitz equilibria under complete and asymmetric information are the same.*

Proof. First, consider the case of complete information. The Principals know the type of the agent and Bertrand compete, which drives the transfer down to the cost of production. Thus, the Principals offer a menu of contracts $(q^*(\theta), t^*(\theta))$, such that $(q^*(\theta), t^*(\theta))$ solves the profit maximization problem of the firm and that has the following properties:

- $t^*(\theta) = C(q^*(\theta))$;
- $q^*(\theta)$ is the solution of $\theta = C'(q^*(\theta))$.

By the convexity of the cost function, the equilibrium is unique.

Next, consider the case of asymmetric information. We will prove that menu of contracts $(q^*(\theta), t^*(\theta))$ is the unique Rothschild -Stiglitz equilibrium. Observe that $(q^*(\theta), t^*(\theta))$ allows type θ to extract all the surplus from the relation with the Principal. Consequently, θ will choose from the menu of contracts, the contract corresponding to her type. This applies for all types. Consequently, the menu of contracts $(q^*(\theta), t^*(\theta))$ is incentive compatible and, obviously, individually rational. Furthermore, since the profits of the incumbent principals are again zero, there are no entries. Thus, the menu of contracts $(q^*(\theta), t^*(\theta))$ constitutes a Rothschild -Stiglitz equilibrium. It is unique from the convexity of the costs.

4 Extensions

The model that we considered above extends with some modifications to the analysis of a number of environments:

- Credit rationing
- Optimal taxation
- Regulation
- Insurance
- Reforming the public sector

- Labor contracts, etc.

In this course, we will focus on the first five cases. Please see the notes for PC 3 for more details. If interested in the last topic and a general overview of the literature, please read Chapter 13 B of Mas-Colell, Whinston, and Green.

5 Bayesian Nash Equilibrium: Review

In this section, we review some game theoretic concepts that are necessary for the analysis of signalling. Namely, we review the definitions of Nash equilibrium, Bayesian equilibrium, and Perfect Bayesian equilibrium. We start with the definition of a strategic game.

Definition: A strategic game $\left\{N, (S_i)_{i\in N}, (u_i)_{i\in N}\right\}$ consists of

- a finite set N of players, $1, ..., n$.
- a nonempty set S_i of (pure) strategies s_i for each player i.
- a payoff function u_i which maps each profile of pure strategies $s, s \in S = \times_{i\in N} S_i.$, to a real number.

The best response of a player given the strategies played by her opponents is the set of strategies that gives her the higher payoff, given the strategy of the other players. Formally,

$$B_i(s_{-i}) = \{s_i \in S_i | u_i(s_i, s_{-i}) > u_i(s_i', s_{-i}) \text{ for all } s_i' \in S_i\}$$

where s_{-i} is the set of strategies played by the opponents of i. Given these definitions, we are ready to introduce the solution concept known as Nash equilibrium.

Definition: A Nash equilibrium of a strategic game $\left\{N, (S_i)_{i\in N}, (u_i)_{i\in N}\right\}$ is a strategy profile s^*, such that for all i :

$$u_i\left(s_i^*, s_{-i}^*\right) > u_i\left(s_i', s_{-i}^*\right) \text{ for all } s_i' \in S_i$$

That is, there is no player i who would find it profitable to deviate from s_i^* and play something else given that everybody else plays s_{-i}^*. Nash equilibrium can also be stated in terms of best responses.

Definition: A Nash equilibrium of a strategic game $\left\{N, (S_i)_{i\in N}, (u_i)_{i\in N}\right\}$ is a strategy profile s^* such that for all i

$$s_i^* \in B_i\left(s_{-i}^*\right)$$

The classical strategic game involves the condition of common knowledge, which can be intuitively stated as follows. Each player knows who are the players, what are their strategies, and what are their payoff functions (i.e. knows the game). She also knows that the other players know the game. The other players know that she knows and they know that she knows that they know... ad infinimum. This condition is relaxed in the games of imperfect information. In these games, the players have a different information (sets). Nevertheless, a similar condition holds, which is outside of the scope of this course. Such games are known as Bayesian games and we will restrict our attention to games in which there is uncertainty about who the other players are, i.e. what are their payoff functions. We model this uncertainty by allowing for a set Ω of possible states of nature. In any play of the game, a particular state is realized, $\omega \in \Omega$. We model the information that a player knows with a function $\tau_i(\omega) = t_i$, $t_i \in T_i$. $\tau_i^{-1}(t_i)$ is a set in Ω. Thus, when the player observes t_i she deduces that the real states is in the set $\tau_i^{-1}(t_i)$.

We assume that the prior probability $p_i\left(\tau_i^{-1}(t_i)\right) > 0$ for all i and all t_i. Then she forms her posterior belief in a Bayesian way. Her posterior belief for each state $\omega \in \Omega$ is:

$$\frac{p(\omega)}{p_i\left(\tau_i^{-1}(t_i)\right)}$$

if $\omega \in \tau_i^{-1}(t_i)$ and zero otherwise.

Example: If $\tau_i(\omega) = \omega$, then the player knows exactly what the state of nature is. Alternatively, if p_i are independent across i and $\tau_i(\omega) = \omega_i$ then i does not learn anything about the state of the world upon observing t_i.

One particular issue that arises in such games is that player may find it optimal to mimic another player. For example, a low productivity worker may adopt some of the attributes of a high productivity worker in order to convince a potential employer that he should be paid a high wage. To model this feature of the environment, we allow for the payoffs of players to depend on both the strategy of the players and their identity, i.e. on (s, ω). We assume that the players are expected utility maximizers with a criterion function U_i.

Definition: A Bayesian game $\left\{N, (S_i)_{i\in N}, (u_i)_{i\in N}, (\Omega, T_i, \tau_i, p_i)_{i\in N}\right\}$consists of :

- a finite set N of players, $1, ..., n$.
- a finite set Ω of states ω.
- a nonempty set S_i of (pure) strategies s_i for each player i.
- a finite set T_i and a function $\tau_i : \Omega \to T_i$. T$_i$ and τ_i determine what i knows at the beginning of the game.
- a prior probability $p_i\left(\tau_i^{-1}(t_i)\right) > 0$ for all t_i.
- a payoff function u_i which maps each profile of pure strategies $s, s \in S = \times_{i \in N} S_i.$, to a real number.

The main difference here relative to what we have seen above in the context of Nash equilibrium is that each player's payoffs are affected by the information t_i. Thus, the outcome of the strategic interaction depends on (i, t_i) and we will define Bayesian Nash equilibrium for this pair. Intuitively, one can think of t_i as the type of player i. Let $U_{(i,t_i)}(s_i, s_{-i})$ denote the expected utility of (i, t_i) from strategy profile s.

Definition: A Bayesian Nash equilibrium of the game $\left\{N, (S_i)_{i \in N}, (u_i)_{i \in N}, (\Omega, T_i, \tau_i, p_i)_{i \in N}\right\}$ is a strategy profile s^*, such that for all i :

$$U_{(i,t_i)}\left(s_i^*, s_{-i}^*\right) \geq U_{(i,t_i)}\left(s_i, s_{-i}^*\right) \text{ for all } s_i \in S_i$$

where the expectation is taken with respect to the posterior belief of the player after observing t_i that assigns probability :

$$\frac{p(\omega)}{p_i\left(\tau_i^{-1}(t_i)\right)}$$

to (s^*, ω) and zero otherwise.

Example: Second-price sealed bid auction.

- the set of players N is $\{1, ..., n\}$.
- the set Ω is the set of valuations, say Θ.
- the set of strategies is R_+, i.e. the amounts that each player can bid.
- the set T_i is Θ.

- the function $\tau_i(\theta_1, ..., \theta_n) = \theta_i$, where the types are iid from a distribution.
- for this reason, the prior belief for a profile v of the players is $\prod_{j=1}^{n} \pi(\theta_j)$ for some distribution over the valuations Θ.
- as we have seen, this game has an equilibrium in which $s_i(i, \theta_i) = \theta_i$.

6 Extensive Games and Perfect Bayesian Equilibrium

When players interact sequentially rather than simultaneously, we say that they play an extensive game.

Definition: An extensive game with perfect information $\{N, H, P, u_i\}$has the following components:

- a set of players N.
- a set H of all possible sequences of strategies for the different stages of the game. An element of h is known as history. All players observe the played strategies.[1] The histories after which no one plays are known as terminal.
- a function P which assigns to each stage of the game (strictly speaking, each history), a player from N whose turn is to choose a strategy.
- a payoff function u_i which maps each profile of pure strategies $s, s \in S = \times_{i \in N} S_i.$, to a real number.

A Bayesian extensive game with observable actions is an extensive game in which the only uncertainty is about the identity (type) of the players. Its formal definition follows below.

Definition: A Bayesian game with observable actions $\left\{\Gamma, (\Theta_i)_{i \in N}, (\pi_i)_{i \in N}, (U_i)_{i \in N}\right\}$has the following components:

[1] The set H also needs to satisfy some technical properties. See Osborne and Rubinstein (1994), Chapter 6.

- $\Gamma = \{N, H, P\}$ is an extensive game with perfect information and simultaneous moves (that is, at each stage the players know the history of played strategies and move simultaneously.), where N is the set of players, H is the set of possible histories, and P assigns who plays when.

- Θ_i is the set of possible types that player i.

- π_i is a probability distribution over Θ_i, such that $\pi_i(\theta_i) > 0$ for all θ_i. π_i is known by all players and is independent across types.

- U_i is the expected utility of i and $U_i(.|\theta, h^t)$ is the expected payoff of i given a type profile θ and terminal history h.

Note that the introduction of imperfect information, in particular the uncertainty about the identity of the opponents, makes the solution of the game depend on the beliefs of the players about the other players. We encountered this feature for the first time in the case of Bayesian strategic games; extensive games make no exception.

Definition. Consider a Bayesian extensive game $\left\{\Gamma, (\Theta_i)_{i\in N}, (\pi_i)_{i\in N}, (U_i)_{i\in N}\right\}$ A perfect Bayesian equilibrium is a pair (σ^*, μ^*) with a generic element (s_i^*, μ_i^*), where s_i^* prescribes what player i does after each history and μ_i^* specifies the beliefs of player i after each history, which satisfy the following properties:

- *Sequential rationality:* For any nonterminal history h, every player (i, θ_i)

$$U_i\left((s_i^*, \mu_i^*)\left(s_{-i}^*, \mu_{-i}^*\right)|\theta, h\right) \geq U_i\left((s_i', \mu_i^*)\left(s_{-i}^*, \mu_{-i}^*\right)|\theta, h\right)$$

  for any $s_i' \in S_i$.

- *Correct initial belief:* The initial belief $\mu_i^0 = \pi_i$ for each player.

- *Action-determined beliefs:* If player i does not move after history h, then the beliefs of the other players about i do not change. If player i moves after history h, then the beliefs of the other players change only as a function of the history h and the strategy played by i.

- *Bayesian updating:* After history h and The beliefs about player i are updated in a Bayesian way:

$$\mu_i\left(\theta_i'|h,s\right)=\frac{\mu_i\left(h,s|\theta_i'\right)\pi_i\left(\theta_i'\right)}{\sum_{\widetilde{\theta}}\mu_i\left(h,s|\widetilde{\theta}_i\right)\pi_i\left(\widetilde{\theta}_i\right)}$$