6.972: Game Theory	April 14, 2005

Lecture 18

Lecturer: Asu Ozdaglar

1 Agenda

In this lecture, we discuss:

- Repeated games with imperfect monitoring and trigger-price strategies.
- Static Games with incomplete information and Bayesian-Nash Equilibrium.

2 Imperfect Monitoring

2.1 Model

Recalling from last class, in a repeated game with imperfect monitoring the players can't observe each other's actions, but only a public outcome that is correlated to their actions.

Each action profile induces a probability distribution $\pi(y, a)$ of the public outcome $y \in Y$. The set Y is assumed finite and a stage's y is independent of other stages.

The realized payoff of a given player i is $r_i(a_i, y)$ and can only depend on another player's actions through the effect of their action's on the distribution of y.

The player i's expected stage payoff will be:

$$g_i(a) = \sum_{y \in Y} \pi(y, a) r_i(a_i, y) \qquad (1)$$

2.2 Example: Noisy Prisoner's Dilemma

Public signal: p

Actions: (a_1,a_2), where $a_i \in (C, D)$

Payoffs:
$r_1(C, p) = 1 + p$
$r_1(D, p) = 4 + p$
$r_2(C, p) = 1 + p$
$r_2(D, p) = 4 + p$

Probability law on p:

$$
\begin{aligned}
a_1 = a_2 = C &\rightarrow p = X \\
a_1 \neq a_2 \quad &\rightarrow p = X - 2 \\
a_1 = a_2 = D &\rightarrow p = X - 4
\end{aligned}
$$

Where X is a continuous random variable with cumulative distribution function F and $E[X] = 0$.

	C	D
C	$(1 + X, 1 + X)$	$(-1 + X, 2 + X)$
D	$(2 + X, -1 + X)$	(X, X)

2.2.1 Trigger-price Strategy

Consider the following trigger type strategy for the noisy prisoner's dilemma:

- (I) - Play (C,C) until $p \leq p^*$, then go to Phase II.
- (II) - Play (D,D) for T periods, then go back to Phase I.

Notice the strategy is stationary and symmetric. Also notice the punishment phase uses a static NE.

Choosing appropriate p^* and T, this strategy is an SPE.

Find the continuation payoffs:

In Phase I, if players do not deviate, their payoff will be:

$$v = (1 - \delta)\{(1 + 0) + \delta[F(p^*)\delta^T v + (1 - F(p^*))v]\} \tag{2}$$

Rearranging this equation, we get:

$$v = \frac{1 - \delta}{1 - \delta^{T+1} F(p^*) - \delta(1 - F(p^*))} \tag{3}$$

If the player deviates, his payoff will be:

$$v_d = (1 - \delta)\{(2 + 0) + \delta[F(p^* + 2)\delta^T v + (1 - F(p^* + 2))v]\} \tag{4}$$

The idea is that deviating provides immediate payoff, but increases the probability of entering Phase II. In order for the strategy to be an SPE, the expected difference in payoff from the deviation must not be positive.

Incentive Compatibility Constraint: $v \geq v_d$

$$v \geq (1-\delta)\{2 + F(2+p^*)\delta^{T+1}v + [1-F(2+p^*)]\delta v\} \tag{5}$$

Plugging in v:

$$\frac{1}{1-\delta^{T+1}F(p^*)-\delta(1-F(p^*))} \geq \frac{2}{1-F(2+p^*)\delta^{T+1}-\delta[1-F(2+p^*)]} \tag{6}$$

Any T and p^* that satisfy this constraint would construct an SPE. The best possible trigger-price equilibrium strategy could be found if we could maximize v subject to the incentive compatibility constraint.

Qestion: Are there other equilibria with higher payoffs ?

2.2.2 Equilibrium Payoffs

Reading: Fudenberg & Tirole section 5.5 and paper by Abreu, Pearce and Stachetti.

There's a beautiful characterization of the set of equilibrium payoffs using a multi-agent DP-type operator.

The set of equilibrium payoffs is a fixed point of this operator and the value iteration converges to the set of equilibrium payoffs.

$$V = \{v(\sigma) \mid \sigma \; is \; an \; SPE\} \tag{7}$$

Definition 1 *Let $W \subseteq \mathbb{R}^I$. A pair (a, u) with $a \in A$ and $u : Y \Rightarrow \mathbb{R}^I$ is admissible with respect to W if:*

- $u(y) \in W \; \forall y \in Y$
- $a_i \in \arg\max_{\overline{a}_i \in A_i} \{(1-\delta)g_i(\overline{a}_i, a_{-i}) + \delta \sum_y \pi(y, \overline{a}_i, a_{-i})u_i(y)\}$
  (Incentive Compatibility Constraint)

A profile $a \in A$ is supportable by W if $\exists u_i : Y \to \mathbb{R}^I$ such that (a, u) is admissible with respect to W.

Definition 2 *For some $W \subseteq \mathbb{R}^I$, define:*
$B(W) = \{v \mid \exists \; (a, u) \; admissible \; w.r.t. \; W \; and \; v_i = (1-\delta)g_i(a) + \delta \sum_y \pi(y, a)u_i(y)\}$

Definition 3 *$W \subseteq \mathbb{R}^I$ is self-generating if $W \subseteq B(W)$.*

Theorem 1 *If W is self-generating, then $B(W) \subseteq V$.*

Theorem 2 *Factorization: $V = B(V)$.*

3 Static Games of Incomplete Information

Reading: Fudenberg & Tirole sections 6.1-6.5 and Krishna's "Auction Theory" ch. 1-5
This is a generalization of strategic games to analyze situations where players have incomplete information about payoffs of others.

Interesting games of incomplete of information arise in a variety of domains: bargaining problems, auctions, routing problems and problems with asymmetric information are some examples.

3.1 Example: Variant of Battle of the Sexes

Imagine a situation where player 1 is unsure if player 2 wants to meet her or avoid her (she assigns 50% probability to each case), but player 2 knows player 1's preferences.

If player 2 wants to meet player 1, the payoffs will be:

	B	S
B	$(2,1)$	$(0,0)$
S	$(0,0)$	$(1,2)$

On the other hand, if player 2 wants to avoid player 1:

	C	D
C	$(2,0)$	$(0,2)$
D	$(0,1)$	$(1,0)$

From player 1's point of view, player 2 has two possible types.

We can also describe the problem by saying the world has two possible states (each with 50% probability) and only player 2 knows the actual state.

Equilibrium Notion: Use NE concept in an expanded game. Players will form conjectures about other player's actions in each state and act optimally given these conjectures.

Payoffs

Given the 50% conjecture, player 1 would do better by playing B:

$$E[B,(B,S)] = \frac{1}{2} * 2 + \frac{1}{2} * 0 = 1$$

$$E[S,(B,S)] = \frac{1}{2} * 0 + \frac{1}{2} * 1 = \frac{1}{2}$$

Clearly, playing B is optimal for type 1 (meet) of player 2 and playing S is optimal for type 2 (avoid) of player 2.

Therefore $(B,(B,S))$ is a NE.

3.2 Bayesian Game

An incomplete information game can be modelled as follows:

- A set of players $\mathcal{I}$
- A set of actions (pure strategy space) for each player i: S_i
- A set of types for each player i: $\theta_i \in \Theta_i$ (assume finite for simplicity)
- A payoff function for each player i: $u_i(s_1, \dots, s_I, \theta_1, \dots, \theta_I)$

Assume the types are drawn from some probability distribution $p(\theta_1, \dots, \theta_I)$. Given player i knows her own type, she will condition the distribution on her actual type: $p(\theta_{-i}|\theta_i)$.

Note that all information, except the actual types of the players, is common knowledge.

Also note that $s_i : \Theta_i \to S_i$, which is very similar to a strategy in a Correlated Equilibrium.

Definition 4 *The strategy profile* $s(\cdot)$ *is a pure strategy Bayesian NE if:* $\forall i \in \mathcal{I}, \forall \theta_i \in \Theta_i$, $s_i(\theta_i) \in \arg\max_{s'_i \in S_i} \sum_{\theta_{-i}} p(\theta_{-i}|\theta_i) u_i(s'_i, s_{-i}(\theta_{-i}), \theta_i, \theta_{-i})$.

3.3 Example: Incomplete Information Cournot Competition

Two firms produce at constant unit cost. Both know firm 1's unit cost is C. Only 2 knows its unit cost. Firm 1 believes firm 2's unit cost to be C_L with probability θ and to be C_H with probability $(1-\theta)$, where $C_L < C_H$.

This game has 2 players, 2 states (L and H) and the possible actions of each player are $q_i \in [0, \infty)$. Firm two has two possible types.

Payoffs Functions:

$$u_1((q_1, q_2), I) = q_1(p(q_1 + q_2) - C)$$
$$u_2((q_1, q_2), I) = q_2(p(q_1 + q_2) - C_I)$$

BNE: (q_1^*, q_L^*, q_H^*) which could also be represented as $(q_1^*, q_2^*(\theta_2))$.

To find the BNE, find the best response functions and find the intersection.

$$B_1(q_L, q_H) = \arg\max_{q_1 \geq 0} \{\theta(p(q_1 + q_L) - C)q_1 + (1-\theta)(p(q_1 + q_H) - C)q_1\}$$

$$B_L(q_1) = \arg\max_{q_L \geq 0} \{(p(q_1 + q_L) - C_L)q_L\}$$

$$B_H(q_1) = \arg\max_{q_H \geq 0} \{(p(q_1 + q_H) - C_H)q_H\}$$

BNE: (q_1^*, q_L^*, q_H^*) such that $B_1(q_L^*, q_H^*) = q_1^*$, $B_L(q_1^*) = q_L^*$ and $B_H(q_1^*) = q_H^*$.

Exercise: $P(Q) = \alpha - Q$, $Q \leq \alpha$. Solve.

$$q_1^* = \frac{1}{3}(\alpha - 2C + \theta C_L + (1-\theta)C_H)$$

$$q_L^* = \frac{1}{3}(\alpha - 2C_L + C) - \frac{1}{6}(1-\theta)(C_H - C_L)$$

$$q_H^* = \frac{1}{3}(\alpha - 2C_H + C) + \frac{1}{6}\theta(C_H - C_L)$$

If both firms know each other's unit cost, then in a NE the output of firm i is $\frac{1}{3}(\alpha - 2C_i + C_j)$.

With incomplete information, firm 2's output is less if its cost is C_L and more if its cost is C_H.

Assuming $q_L^* < q_1^* < q_H^*$:

If firm 1 knew firm 2's cost is high, then it would produce more. Since it doesn't, it only produces a moderate output. If 2's cost is indeed high, it benefits from the lack of knowledge of firm 1 and it produces more than if 1 knew his actual cost.