Algorithmic Game Theory November 18, 2013

Lecture 5

Lecturer: Michal Feldman *Scribe: Ofir Geri, Ofri Ziv, Yonatan Komornik*

1 Coalitions and Nash Equilibria

We introduced the notion of Nash equilibrium, which makes any unilateral deviation unprofitable. We now wish to define an equilibrium that makes a cooperative deviation unprofitable for any coalition of players.

Definition 1 *A cooperative deviation that is profitable for a coalition $T \subseteq N$ is a strategy profile S'_T so that for every $i \in T$,*

$$c_i(S'_T, S_{-T}) < c_i(S_T, S_{-T}).$$

Definition 2 *A strategy profile S is a Strong Equilibrium (SE) if any cooperative deviation is unprofitable.*

Any strong equilibrium is a Nash equilibrium, hence, the Price of Anarchy (PoA) for strong equilibria is at most the PoA for Nash equilibria. The prisoner's dilemma game does not have a strong equilibrium. The only strategy profile that is a Nash equilibrium is when both players choose to defect, but both players can improve their utility if they both choose to cooperate.

1.1 Cost-Sharing Connection Games

Reminder

Let $G = (V, E)$ be a directed graph. Every edge $e \in E$ has a cost, c_e. Player i chooses a path connecting its source node s_i with its node destination t_i. We define the cost of strategy s for player i to be

$$c_i(s) = \sum_{e \in p_i} \frac{c_e}{x_e}$$

where p_i is the path chosen by player i and x_e is the number of players that use edge e.

We saw that the PoA for Nash equilibria is n. We shall now find the PoA for strong equilibria.

Definition 3 *The Price of Anarchy (PoA) for strong equilibria is defined as*

$$PoA = \max_{G} \max_{s \text{ is } SE} \frac{c(s)}{OPT}$$

where OPT is the minimal cost.

First we show an example of a cost-sharing connection game that doesn't have a strong equilibrium. Consider the network presented in Figure 1.

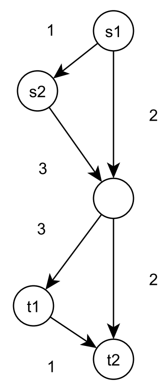

Figure 1

There are two players. We see that each player must use one specific edge, and can choose one of two other edges. The optimum is achieved when both players go to the left (i.e., not use the edges that cost 2). In this case, the costs are:

$$cost(s^*) = 1 + 3 + 3 + 1 = 8$$

$$c_1(s^*) = c_2(s^*) = 4$$

However, this is not a Nash equilibrium. Player 1 has a dominant strategy, which is choosing the edge that costs 2. The best response of player 2 is also to choose the edge that costs 2. The cost of this Nash equilibrium is:

$$cost(s) = 2 + 3 + 3 + 2 = 10$$

$$c_1(s) = c_2(s) = 5$$

This is not a strong equilibrium, and therefore, no strong equilibrium exists in this game.

Symmetric cost-sharing connection games do have a strong equilibrium. The strategy in which all players choose the optimal path from the source to the destination is a strong equilibrium.

Theorem 4 *In cost-sharing connection games, $PoA \leq H_n$ (for strong equilibria, if such an equilibrium exists).*

Proof: Let s^* be the optimal strategy profile, and s be a strong equilibrium. Since s is a strong equilibrium, there is some player i so that $c_i(s) \leq c_i(s^*)$. If there is no such player, all players will deviate to s^*. Let's assume without loss of generality that this player is player n. We get that

$$c_n(s) \leq c_n(s^*).$$

Consider the coalition of players $1, 2, ..., n-1$. s is a strong equilibrium, so we get that without loss of generality,

$$c_{n-1}(s) \leq c_{n-1}(s_1^*, s_2^*, ..., s_{n-1}^*, s_n). \tag{1}$$

We continue iteratively, and get the following:

$$c_{n-2}(s) \leq c_{n-2}(s_1^*, s_2^*, ..., s_{n-2}^*, s_{n-1}, s_n) \tag{2}$$

$$\vdots$$

$$c_1(s) \leq c_1(s_1^*, s_2, ..., s_{n-1}, s_n) \tag{n}$$

Observe that removing a player from the game can only increase the cost for other players. For instance,

$$c_{n-1}(s_1^*, s_2^*, ..., s_{n-1}^*, s_n) \leq c_{n-1}(s_1^*, s_2^*, ..., s_{n-1}^*)$$

where $c_{n-1}(s_1^*, s_2^*, ..., s_{n-1}^*)$ is a strategy profile for a game with $n-1$ players. Summing the inequalities (1), (2), ..., (n) gives

$$\begin{aligned} cost(s) &= c_1(s) + ... + c_n(s) \\ &\leq c_1(s_1^*, s_2, ..., s_{n-1}, s_n) + ... + c_{n-1}(s_1^*, s_2^*, ..., s_{n-1}^*, s_n) + c_n(s^*) \\ &\leq c_1(s_1^*) + ... + c_{n-1}(s_1^*, s_2^*, ..., s_{n-1}^*) + c_n(s^*). \end{aligned}$$

Recall that the potential function is defined as $\Phi(s) = \sum_{e \in E} c_e \sum_{i=1}^{x_e} \frac{1}{i}$.
By definition, we get that

$$c_i(s_1^*, s_2^*, ..., s_i^*) = \Phi(s_1^*, s_2^*, ..., s_i^*) - \Phi(s_1^*, s_2^*, ..., s_{i-1}^*).$$

And finally,

$$\begin{aligned} cost(s) &\leq \sum_{i=1}^{n} c_i(s_1^*, s_2^*, ..., s_i^*) \\ &= \sum_{i=1}^{n} \Phi(s_1^*, s_2^*, ..., s_i^*) - \Phi(s_1^*, s_2^*, ..., s_{i-1}^*) \\ &= \Phi(s^*) - \Phi(\phi) && \text{[telescoping series]} \\ &= \Phi(s^*). \end{aligned}$$

We showed in a previous lecture that $\Phi(s^*) \leq H_n \cdot cost(s^*)$, hence

$$cost(s) \leq \Phi(s^*) \leq H_n \cdot cost(s^*).$$

Therefore, $PoA \leq H_n$. ■

2 Bayesian Games

In this section we take a first look at Bayesian Games (otherwise known as Incomplete Information Games), and show that the smoothness property we saw still applies.

2.1 Model

- n players.
- Each player i has a type space T_i. Type t_i is sampled for player i from T_i according to a distribution F_i.
  We denote $T = T_1 \times T_2 \times \cdots \times T_n$.
- Each player i has an action space A_i.
  We denote $A = A_1 \times A_2 \times \cdots \times A_n$.
- Let $C^* : T \to A$ be a choice function.
- Let $\Gamma = (T, A, u)$ be a Bayesian Game.

Definition 5 *A **strategy** of a player i is a function S_i from T_i to a distribution of A_i. In other words, $S_i(t_i)$ is a mixed strategy for player i of type t_i.*

Definition 6 *A strategy profile S is a **Bayesian Equilibrium (BNE)** if for each player i and every $t_i \in T_i, a'_i \in A_i$ it holds that:*

$$\begin{aligned} \begin{bmatrix} \text{The expected utility of player} \\ i \text{ given the starategy } S_i(t_i) \end{bmatrix} &\geq \begin{bmatrix} \text{The expected utility of player} \\ i \text{ given the action } a'_i \end{bmatrix} \\ E_{t_{-i} \sim F_{-i}} \left[E_{a \in S(t)} \left[u_i(t_i, a) \right] \right] &\geq E_{t_{-i} \sim F_{-i}} \left[E_{a_{-i} \in S_{-i}(t_{-i})} \left[u_i(t_i, (a'_i, a_{-i})) \right] \right] \end{aligned}$$

where $u_i(t_i, a)$ is the utility of player i of type t_i, and a is a profile of actions.

2.2 Price of Anarchy

- Let $W(\vec{t},\vec{a})$ be the target function. Note that in this lecture the target function will always be the social welfare function.
- Let $OPT(\vec{t})$ be the optimal action profile for types t.

Definition 7 *The* ***Price of Anarchy for Bayesian Games (BPoA)*** *is defined as follows:*

$$BPoA = \min_{S \text{ is } BNE} \frac{E_{t\sim F}\left[E_{a\in S(t)}\left[W(t,a)\right]\right]}{E_{t\sim F}\left[W(t,OPT(t))\right]}.$$

2.3 Smooth Games

Reminder

A complete information game is $(\lambda,\mu)-$*smooth* in regard to an action profile a^* and a target function $W : A \to \mathbb{R}_+$ if it holds that:

$$\sum_{i=1}^{n} u_i(a_i^*, a_{-i}) \geq \lambda \underbrace{W(a^*)}_{OPT} - \mu W(a).$$

We define smooth Bayesian games and show that such games preserve similar PoA properties as smooth complete information games.

Definition 8 *A Bayesian game $\Gamma = (T,A,u)$ is $(\lambda,\mu)-$ smooth in regard to the choice function c^* if for every $s,t \in T, a \in A$ (such that a is a valid action vector for s), it holds that:*

$$\sum_{i=1}^{n} u_i(t_i,(c_i^*(t),a_{-i})) \geq \lambda W(t,c^*(t)) - \mu W(s,a).$$

Theorem 9 *Let $\Gamma = (T,A,u)$ be $(\lambda,\mu)-$ smooth in regard to the social welfare function W. Then:*

$$W(BNE) \geq \frac{\lambda}{1+\mu} W(OPT).$$

Proof: Let Γ be a $(\lambda,\mu)-$*smooth* game in regard to an optimal choice function c^*, and s be a Bayesian equilibrium. The best strategy for each player depends on the types of other players. Therefore, the best strategy that player i can choose is to sample the types of the rest of the players, and then act according to c_i^*. Formally, we define $\widehat{s}_i(t_i)$ to be the action of player i with type t_i. The action will be determined by sampling $\sigma_{-i}^{(i)} \sim F_{-i}$, and then playing $c_i^*(t_i, \sigma_{-i}^{(i)})$. We wish to prove a relation that takes the following form:

$$\underbrace{E_{t\sim F}\left[E_{a\sim s(t)}\left[W(t,a)\right]\right]}_{BNE} \geq \textbf{?} \cdot \underbrace{E_{t\sim F}\left[W(t,c^*(t))\right]}_{OPT}.$$

We show that:

$$\begin{aligned}
E_{t\sim F}\left[E_{a\sim s(t)}\left[W(t,a)\right]\right] &= E_{t\sim F}\left[E_{a\sim s(t)}\left[\sum_{i=1}^{n} u_i(t_i,a)\right]\right] && \text{[definition of W]}\\
&= \sum_{i=1}^{n} E_{t\sim F}\left[E_{a\sim s(t)}\left[u_i(t_i,a)\right]\right] && \text{[linearity of expectation]}\\
&\geq \sum_{i=1}^{n} E_{t\sim F}\left[E_{\substack{\widehat{a}_i\sim\widehat{s}_i(t_i)\\ a\sim s(t)}}\left[u_i(t_i,(\widehat{a}_i,a_{-i}))\right]\right] && \text{[}s(t)\text{ is a BNE]}\\
&\geq \sum_{i=1}^{n} E_{t\sim F}\left[E_{\substack{\sigma_{-i}^{(i)}\sim F\\ a\sim s(t)}}\left[u_i(t_i,(c_i^*(t_i,\sigma_{-i}^{(i)}),a_{-i}))\right]\right]. && \text{[definition of }\widehat{s}_i\text{]}
\end{aligned}$$

Since each player samples $\sigma_{-i}^{(i)}$ from independent distributions for the rest of the players, we can sample σ once for all players and use it instead. We continue:

$$\begin{aligned}
E_{t\sim F}\left[E_{a\sim s(t)}\left[W(t,a)\right]\right] &\geq \sum_{i=1}^{n} E_{t,\sigma\sim F}\left[E_{a\sim s(t)}\left[u_i(t_i,(c_i^*(t_i,\sigma_{-i}),a_{-i}))\right]\right]\\
&= \sum_{i=1}^{n} E_{t,\sigma\sim F}\left[E_{a\sim s(\sigma)}\left[u_i(t_i,(c_i^*(t_i,t_{-i}),a_{-i}))\right]\right]\\
&\geq E_{t,\sigma\sim F}\left[E_{a\sim s(\sigma)}\left[\lambda\cdot W(t,c^*(t)) - \mu\cdot W(\sigma,a)\right]\right] && \text{[smoothness]}\\
&= \lambda\cdot E_{t,\sigma\sim F}\left[E_{a\sim s(\sigma)}\left[W(t,c^*(t))\right]\right] - \mu\cdot E_{t,\sigma\sim F}\left[E_{a\sim s(\sigma)}\left[W(\sigma,a)\right]\right]\\
&= \lambda\cdot E_{t\sim F}\left[W(t,c^*(t))\right] - \mu\cdot E_{t\sim F}\left[E_{a\sim s(t)}\left[W(t,a)\right]\right].
\end{aligned}$$

Finally, we get

$$W(BNE) \geq \frac{\lambda}{1+\mu} W(OPT).$$

■