CS 6840: Algorithmic Game Theory **Spring 2017**

Lecture 8: February 10

Lecturer: Éva Tardos *Scribe: Victor Reis*

8.1 Mixed Nash Equilibria

Given a game with k players and a vector of strategies $\mathbf{s} = (s_1, \ldots, s_k)$, let $c_i(\mathbf{s})$ denote the cost of player i. For each i, let σ_i be a probability distribution on strategies, and let $\sigma = \prod_i \sigma_i$ sample strategies $s_i \sim \sigma_i$ for each player independently, so that the expected cost for player i is $\mathbb{E}_{\mathbf{s}\sim\sigma}\left[c_i(\mathbf{s})\right]$.

We say that σ is a mixed Nash equilibrium if for all players i, and for all strategies s_i', we have

$$\mathbb{E}_{\mathbf{s}\sim\sigma}\left[c_i(\mathbf{s})\right] \leq \mathbb{E}_{\mathbf{s}\sim\sigma}\left[c_i(s_i', \mathbf{s}_{-i})\right].$$

Lemma 8.1 *If σ is a mixed Nash equilibrium, then for any probability distribution σ_i', we have*

$$\mathbb{E}_{\boldsymbol{s}\sim\sigma}\left[c_i(\boldsymbol{s})\right] \leq \mathbb{E}_{s_i'\sim\sigma_i', \boldsymbol{s}\sim\sigma}\left[c_i(s_i', \boldsymbol{s}_{-i})\right].$$

Proof: Indeed, we have

$$\begin{aligned}
\mathbb{E}_{s_i'\sim\sigma_i', \mathbf{s}\sim\sigma}\left[c_i(s_i', \mathbf{s}_{-i})\right] &= \mathbb{E}_{s_i'\sim\sigma_i'}\mathbb{E}_{\mathbf{s}\sim\sigma}\left[c_i(s_i', \mathbf{s}_{-i})\right] \\
&\geq \mathbb{E}_{s_i'\sim\sigma_i'}\mathbb{E}_{\mathbf{s}\sim\sigma}\left[c_i(\mathbf{s})\right] \\
&= \mathbb{E}_{\mathbf{s}\sim\sigma}\left[c_i(\mathbf{s})\right],
\end{aligned}$$

as desired. ■

8.2 Smooth Games

A game is said to be (λ, μ)-*smooth* if for any two strategy vectors $\mathbf{s}, \mathbf{s}'$ we have

$$\sum_i c_i(s_i', \mathbf{s}_{-i}) \leq \lambda \sum_i c_i(\mathbf{s}') + \mu \sum_i c_i(\mathbf{s}).$$

Theorem 8.2 *Given a (λ, μ)-smooth game with $\mu < 1$ and σ a mixed Nash equilibrium, we have*

$$\mathbb{E}_{\boldsymbol{s}\sim\sigma}\left[\sum_i c_i(\boldsymbol{s})\right] \leq \frac{\lambda}{1-\mu} \cdot \min_{\boldsymbol{s}'} \sum_i c_i(\boldsymbol{s}').$$

Proof: Since σ is a mixed Nash, we have for any player i,

$$\mathbb{E}_{\mathbf{s}\sim\sigma}\left[c_i(\mathbf{s})\right] \leq \mathbb{E}_{\mathbf{s}\sim\sigma}\left[c_i(s_i', \mathbf{s}_{-i})\right].$$

Summing over all i gives

$$\sum_i \mathbb{E}_{\mathbf{s}\sim\sigma}\left[c_i(\mathbf{s})\right] \le \sum_i \mathbb{E}_{\mathbf{s}\sim\sigma}\left[c_i(s_i', \mathbf{s}_{-i})\right].$$

By linearity of expectation, we can write this as

$$\mathbb{E}_{\mathbf{s}\sim\sigma}\left[\sum_i c_i(\mathbf{s})\right] \le \mathbb{E}_{\mathbf{s}\sim\sigma}\left[\sum_i c_i(s_i', \mathbf{s}_{-i})\right].$$

Applying smoothness, we have

$$\mathbb{E}_{\mathbf{s}\sim\sigma}\left[\sum_i c_i(\mathbf{s})\right] \le \mathbb{E}_{\mathbf{s}\sim\sigma}\left[\lambda \sum_i c_i(\mathbf{s}') + \mu \sum_i c_i(\mathbf{s})\right] = \lambda \sum_i c_i(\mathbf{s}') + \mu \cdot \mathbb{E}_{\mathbf{s}\sim\sigma}\left[\sum_i c_i(\mathbf{s})\right].$$

Rearranging terms, we get the desired result. ■

8.3 Coarse-correlated equilibria

In the proof of the above theorem, we never used the fact that σ_i are independent distributions. This motivates a more general notion of equilibrium: given a probability distribution σ on vectors of strategies, we say that σ is a coarse-correlated equilibrium if for all players i, and for all strategies s_i', we have

$$\mathbb{E}_{\mathbf{s}\sim\sigma}\left[c_i(\mathbf{s})\right] \le \mathbb{E}_{\mathbf{s}\sim\sigma}\left[c_i(s_i', \mathbf{s}_{-i})\right].$$

This is exactly the same definition as for mixed Nash, but here we don't have the condition that σ is a product of distributions σ_i. However, we didn't use this condition in our proof, so we have

Theorem 8.3 *Given a (λ, μ)-smooth game with $\mu < 1$ and σ a coarse correlated equilibrium, we have*

$$\mathbb{E}_{\boldsymbol{s}\sim\sigma}\left[\sum_i c_i(\boldsymbol{s})\right] \le \frac{\lambda}{1-\mu} \cdot \min_{\boldsymbol{s}'} \sum_i c_i(\boldsymbol{s}').$$

Any mixed Nash is therefore a coarse-correlated equilibrium, but not vice-versa, as we'll show next.

8.3.1 Example: Jobs and Machines

Consider the game described in Problem Set 1, Problem 3, in the special case where we have 4 jobs and 6 machines. Each job has weight 1, and a strategy consists of picking one of the machines. The cost of a job is equal to the number of jobs which picked the same machine.

This game has a trivial pure Nash equilibrium: each player can pick one of the machines deterministically, each paying cost 1, with no incentive to deviate.

It also has a simple mixed Nash equilibrium: if all players choose a machine uniformly at random, each player will have an expected cost of $1 + \frac{3}{6} = 1.5$, to see why, consider a player. Choosing a machine, the expected number of other jobs on the same machine is 3/6, as each jobs has probability 1/2 of being there. Any deterministic deviation from any fixed player does not affect the expected cost.

Now consider a distribution σ which picks two players at random and puts them in a uniformly randomly chosen machine, and puts the other two players alone in uniformly randomly chosen empty machines. Since the distributions are no longer independent, this is not a mixed Nash, but it is still a coarse-correlated equilibrium: the expected cost for a player is $\frac{1}{2} \cdot 2 + \frac{1}{2} \cdot 1 = 1.5$ (two jobs have cost 2 as they share a link and the other two have cost 1), and the cost of a player who deviates and picks a fixed machine is

$$\frac{1}{2} \cdot \left(\frac{1}{2} \cdot 1 + \frac{1}{2} \cdot 2\right) + \frac{1}{2} \cdot \left(\frac{1}{6} \cdot 3 + \frac{1}{6} \cdot 2 + \frac{4}{6} \cdot 1\right) = 1.5,$$

where the first term corresponds to the case when the player who deviated was supposed to have been paired by σ, so when he is trying to deviate, the other two are on 3 separate machines, so he has probability $1/2$ of colliding with one of them. The second term corresponds to the case when the player who deviated was one of the two other players who would have surely been placed in an empty machine, with probability $1/6$ he will share with the pair, with probability $1/6$ with the other singleton, and with the remaining probability $4/6$ he'll be alone. Since there is no incentive to deviate, σ is indeed a coarse-correlated equilibrium.