0368.4483: Algorithmic game theory **Spring 2016**

Lecture 10 : May 25, 2016

Lecturer: *Michal Feldman* *Scribes:* *Ynon Flum, Eyal Golombek, Yuval Lewi Uri Meir, Idan Rejwan & Jonathan Shafer*

Contents

1 Introduction

Up till this point in the course we worked in the *full-information framework*, which assumes that each player knows the true valuations of all the other players. In this lecture we present the *Bayesian framework* which assumes that each players knows a distribution over the valuations of the other players, but does not know the realization of that distribution.

2 The Bayesian model

2.1 Definitions

Each player i has:

- A set T_i of potential types called the *type space*, from which his actual type $t_i \in T_i$ is taken. The type of a player represents his preferences in the game. In the case of auctions, each player's type is his valuations over the items.

- A set A_i of potential actions for the player called the *action space*. In the case of auctions, the actions are the bids the players make.
- A strategy function $\sigma_i : T_i \to A_i$ that determines what action a player will take given his type.

2.2 Assumptions

Each player i knows:

- His own type t_i.
- A prior distribution F over the types of all the players $t = (t_1, ..., t_n)$. Hence, each player i has a posterior distribution $F_{-i}|t_i$ over the types of the other players.

  For example, consider the following distribution over the types of two players:

Pr	t_1	t_2
$\frac{1}{3}$	1	2
$\frac{1}{3}$	1	3
$\frac{1}{3}$	2	5

  Assume that $t_1 = 1$ and $t_2 = 2$. Then player 1 knows that his own type is 1, and therefore his posterior distribution over t_2 is $P(t_2 = 2|t_1 = 1) = P(t_2 = 3|t_1 = 1) = \frac{1}{2}$.

  As a special case, if $F = F_1 \times \cdots \times F_n$ is a product distribution such that $t_i \sim F_i$ then $F_{-i}|t_i = F_{-i}$ and so each player learns nothing about the other players' types from observing his own.
- In a Bayesian equilibrium (to be defined immediately), each player knows the full strategy profile $\sigma = (\sigma_1, ..., \sigma_n)$. This induces a distribution $\sigma_{-i}(t_{-i}|t_i)$ over the actions of the other players.

3 Bayesian Nash Equilibrium

Definition. A Bayesian Nash Equilibrium (BNE) is a strategy profile $\sigma = (\sigma_1, ..., \sigma_n)$ such that for every player i and every type $t_i \in T_i$, the strategy σ_i maximizes the expected utility of player i given t_i, i.e.:

$$\mathbb{E}_{t_{-i} \sim F_{-i}|t_i} \Big[u_i(\sigma_i(t_i), \sigma_{-i}(t_{-i})) \Big]$$

3.1 Example: a BNE for a first-price auction

Consider a first-price auction with two players and a single item in which the players' valuations v_1, v_2 are i.i.d $U([0, 1])$.

Claim: $\sigma_1 = \frac{v_1}{2}$, $\sigma_2 = \frac{v_2}{2}$ is a BNE.

Note: For concreteness, when discussing auctions we talk about player valuations v_i rather than player types t_i, and about player bids b_i rather than actions a_i.

Proof: For player 1, v_1 is known and $v_2 \sim U([0,1])$. Because $\sigma_2 = \frac{v_2}{2}$, it follows that $b_2 \sim U([0, \frac{1}{2}])$. Hence the expected utility of player 1 is:

$$\mathbb{E}_{v_2}\big[u_1(b_1)\big] = 0 \cdot P(\text{player 1 looses}) + (v_1 - b_1) \cdot P(\text{player 1 wins}) =$$
$$= (v_1 - b_1) \cdot P_{v_2}(b_1 > \frac{v_2}{2}) = (v_1 - b_1) \cdot \min\{2b_1, 1\}$$

By taking the derivative, we can see that

$$\frac{v_1}{2} = \text{argmax}_{b_1} \mathbb{E}_{v_2}\big[u_1(b_1)\big]$$

Thus, player 1 maximizes his expected utility by bidding $b_1 = \frac{v_1}{2}$, and a similar reasoning shows that the best response for player 2 is to bid $b_2 = \frac{v_2}{2}$. ■

Note: The seller's expected revenue in this auction would be $\frac{1}{3}$, because $\mathbb{E}\big[\max\{v_1, v_2\}\big] = \frac{2}{3}$, and the revenue is half that amount. Moreover, the seller's expected revenue in a second-price auction with the same two players would be $\frac{1}{3}$ as well, because $\mathbb{E}\big[\min\{v_1, v_2\}\big] = \frac{1}{3}$, the bids are honest, and the revenue is the second-price. This is an example of a general result (that follows from Myerson's theorem).

4 Price of anarchy for BNE

In the previous lecture, we saw that in a second-price auction with submodular valuations, the PoA for NE that satisfy no-over-bidding (NOB) is at most 2.

4.1 Reminder of the proof method

Let b be a NOB NE. We define hypothetical deviations $(b_1^*, ..., b_n^*)$ (using the additive functions a_i^*). Then:

$$\text{SW}(b) = \sum_{i=1}^{n} v_i(s_i(b)) \geq \sum_{i=1}^{n} u_i(b) \overset{(*)}{\geq} \sum_{i=1}^{n} u_i(b_i^*, b_{-i}) \overset{(**)}{\geq} \text{OPT} - \sum_{i=1}^{n} \sum_{j \in s_i(b)} b_{ij} \overset{(\star)}{\geq} \text{OPT} - \sum_{i=1}^{n} v_i(s_i(b))$$

And so

$$\text{SW}(b) \geq \frac{1}{2}\text{OPT}$$

where v_i, u_i and s_i denote the valuation, utility and allocation of player i (respectively), $(*)$ follows from b being a NE, $(**)$ follows from the choice of b^* and $(\star)$ follows from NOB. ■

We would like to generalize this result to BNE.

Definition. Let F be a prior distribution over the valuations of the players. The Bayesian Price of Anarchy (BPoA) with respect to F is

$$\text{BPoA}(F) = \frac{\mathbb{E}_{v \sim F}\big[OPT(v)\big]}{\inf_{\sigma \in \text{BNE}} \mathbb{E}_{v \sim F}\big[\text{SW}(\sigma(v))\big]}$$

4.2 Main result

From here on we will assume that the prior F is a product distribution $F = F_1 \times \cdots \times F_n$ such that $v_i \sim F_i$ for all i.

Theorem. For every G a bayesian simultaneous second price auction, if $\mathcal{F}$ the prior-distribution of G is a product distribution, and if the support of $\mathcal{F}$ is contained in the set of all XOS functions, then for every bayesian Nash Equilibrium $\sigma = (\sigma_1, ..., \sigma_n)$ satisfying NOB it's true that:

$$\mathbb{E}_{V\sim\ \mathcal{F}}[SW(\sigma(V)] \geq \frac{1}{2}\mathbb{E}_{V\sim\ \mathcal{F}}[OPT(V)]$$

Remark. Simply put, it means that under the condition that our prior distribution is a product distribution, we get the same result as in a full information environment, which is an upper bound to the PoA of simultaneous second price auctions with XOS valuation functions.

Remark. It's clear that we can't expect to get a better result, since Full Information Environment can be seen a special case of Bayesian Environment where the distribution is degenrate, and we have seen that the same bound on PoA in full information environment is tight.

Proof: Let σ be some BNE. Let i be some player. We will define a deviation for i from the Equilibrium σ in the following way: We sample $W_{-i} \sim\ \mathcal{F}_{-i}$, now player i has a full valuation (v_i, W_{-i}) since he obviously knows his own valuation. Given this complete valuation (v_i, W_{-i}) we can define σ_i^* the deviation:

Let's mark S_i^* the package i gets under the optimal allocation, according to the sampled valuation (v_i, W_{-i}). $\sigma_i^*(v_i)$ is a bidding vector $b_i^* = (b_{i,1}^*, ..., b_{i,m}^*)$ defined by:

- if $j \in S_i^*$ then $b_{i,j}^* = a_i^*(j)$ where a_i^* is an additive function derived from the fact that the valuation is XOS. (See the same proof for Full Information environment for explanation)
- if $j \notin S_i^*$ then $b_{i,j}^* = 0$

We mark $\sigma_i^*(v_i) = b_i^*(v_i, W_{-i})$ since it's dependent on the sample W. Since σ is a BNE, player i can't benefit from any deviation, including σ_i^*:

$$\mathbb{E}_{V_{-i}\sim\ \mathcal{F}_{-i}}[u_i(\sigma(V)] \geq \mathbb{E}_{V_{-i}\sim\ \mathcal{F}_{-i}}[u_i(\sigma_i^*(V_i), \sigma_{-i}(V_{-i})]$$

From the way we defined the deviation we get:

$$\mathbb{E}_{V_{-i}\sim\ \mathcal{F}_{-i}}[u_i(\sigma_i^*(v_i), \sigma_{-i}(v_{-i})] = \mathbb{E}_{V_{-i}\sim\ \mathcal{F}_{-i};W_{-i}\sim\ \mathcal{F}}[u_i(b_i^*(v_i, W_{-i}), \sigma_{-i}(v_{-i})]$$

All of this analysis was done in the point of view of a single player: i. But v_i is also chosen at random by $\mathcal{F}$ in the point of view of the mechanisem designer. We can take expactancy over v_i as well. Since $\mathcal{F}$ is a product distribution we get:

$$\mathbb{E}_{V\sim\ \mathcal{F}}[u_i(\sigma(V)] \geq \mathbb{E}_{V\sim\ \mathcal{F};W\sim\ \mathcal{F}}[u_i(b_i^*(W), \sigma_{-i}(V_{-i})]$$

Using the linearity of expectation we will sum this inequallity over all players to get:

$$\mathbb{E}_{V\sim\ \mathcal{F}}[\sum_{i=1}^n u_i(\sigma(V)] \geq \mathbb{E}_{V\sim\ \mathcal{F};W\sim\ \mathcal{F}}[\sum_{i=1}^n u_i(b_i^*(W), \sigma_{-i}(V_{-i})]$$

Since $b_i^*(W)$, the deviation bid of i, was crafted in a very specific way, we know that:

$$\sum_{i=1}^{n} u_i(b_i^*(W), \sigma_{-i}(V_{-i})) \geq OPT(W) - \sum_{i=1}^{n} \sum_{j \in S_i} \sigma_i(V_i)_j$$

Where S_i is the package i would get from the allocation of the simultaneous second price mechanism, if everyone played according to the N.E σ and if the valuations were all according to V (See the same proof for full information environment for reference). Finally:

$$\mathbb{E}_{V\sim \mathcal{F}}[\sum_{i=1}^{n} u_i(\sigma(V)] \geq \mathbb{E}_{V\sim \mathcal{F};W\sim \mathcal{F}}[OPT(W) - \sum_{i=1}^{n} \sum_{j \in S_i} \sigma_i(V_i)_j]$$
$$\mathbb{E}_{V\sim \mathcal{F}}[\sum_{i=1}^{n} u_i(\sigma(V)] \geq \mathbb{E}_{W\sim \mathcal{F}}[OPT(W)] - \sum_{i=1}^{n} \mathbb{E}_{V\sim \mathcal{F}}[\sum_{j \in S_i} \sigma_i(V_i)_j]$$

Since σ the Nash Equilibrium upholds no-overbiding we know that:

$$V_i(S_i) \geq \sum_{j \in S_i} \sigma_i(V_i)_j$$

Moreover, since all the prices are non-negative, it's clear that

$$\mathbb{E}_{V\sim \mathcal{F}}[SW(\sigma(V)] \geq \mathbb{E}_{V\sim \mathcal{F}}[\sum_{i=1}^{n} u_i(\sigma(V)]$$

And so all togther we get:

$$\mathbb{E}_{V\sim \mathcal{F}}[SW(\sigma(V)] \geq \mathbb{E}_{W\sim \mathcal{F}}[OPT(W)] - \sum_{i=1}^{n} \mathbb{E}_{V\sim \mathcal{F}}[V_i(S_i)]$$
$$\mathbb{E}_{V\sim \mathcal{F}}[SW(\sigma(V)] \geq \mathbb{E}_{W\sim \mathcal{F}}[OPT(W)] - \mathbb{E}_{V\sim \mathcal{F}}[\sum_{i=1}^{n} V_i(S_i)]$$
$$2\mathbb{E}_{V\sim \mathcal{F}}[SW(\sigma(V)] \geq \mathbb{E}_{W\sim \mathcal{F}}[OPT(W)]$$

The left hand side is twice the expected SW of some BNE, and right hand side is the expected value of the optimal allocation. Since σ is an arbitrary BNE, we get our bound for BPoA. ■

Remark. In the full-information framework we proved that if for every profile v there exist bids $b_1^*, ..., b_n^*$ such that:

$$\sum_{i=1}^{n} u_i(b_i^*, b_{-i}) \geq \lambda * OPT(v) - \mu * \sum_{i=1}^{n} p_i(b)$$

Then:

$$PoA \leq \frac{\lambda}{1+\mu}$$

Using the "sample trick" of the last theorem's proof, we get the same result for BNE

5 PoA in single-item first price auctions

Our objective is getting an upper bound for the price of anarchy in first price auctions.
In this lesson we cover a bound for pure N.E.

Remark. In a single-item second price auction, for every valuation profile v,
there exist bids $b_1^*, ..., b_n^*$ such that: $\sum_{i=1}^{n} u_i(b_i^*, b_{-i}) \geq \max_{i=1}^{n} v_i - \max_{i=1}^{n} b_i$
We would like to get a similar connection in first price auctions in order to get our bound.

Lemma. In a first price auction, if the highest bidder bids $b_i^* = \frac{v_i}{2}$,
and then rest of the bidders bid $b_i^* = 0$ then:

$$(*)\ \sum_{i=1}^{n} u_i(b_i^*, b_{-i}) \geq \frac{1}{2} \max_{i=1}^{n} v_i - \max_{i=1}^{n} b_i$$

Proof: We will divide our proof into two cases:

- If $\frac{1}{2} \max_{i=1}^{n} v_i - \max_{i=1}^{n} b_i < 0$, then since for every player i, $u_i(b_i^*, b_{-i}) \geq 0$, we get that (*) holds.
- Otherwise $\frac{1}{2} \max_{i=1}^{n} v_i > \max_{i=1}^{n} b_i$,
  let $i^* = argmax\, v_i$, i^* wins the auction, yielding utility $\frac{1}{2} v_{i^*} \geq \frac{1}{2} v_{i^*} - \max_{i=1}^{n} b_i$.
  This verifies (*).

■

Corollary. Every pure Nash equilibrium of a first-price single item auction has SW of at least $\frac{1}{2}$ of OPT.

Remark. It can be shown that each pure NE of a first price auction has optimal SW.

Proof of corollary: Let $v_i(b)$ be the SW contributed by player i in the outcome of the profile b, and let $p_i(b)$ be i's payment and $u_i(b) = v_i(b) - p_i(b)$ his utility.
Then the following holds:

$$\sum_{i=1}^{n} v_i(b) = \sum_{i=1}^{n} u_i(b) + \sum_{i=1}^{n} p_i(b) \geq \sum_{i=1}^{n} u_i(b_i^*, b_{-i}) + \max_i b_i \geq \tfrac{1}{2} \max_i v_i - \max_i b_i + \max_i b_i = \frac{1}{2} \max_i v_i$$

■