Algorithmic Game Theory, Summer 2017 — Lecture 7 (4 pages)

Introduction to Mechanism Design

Instructor: Thomas Kesselheim

In the previous lectures we have adopted a somewhat passive perspective: We have considered a given game, and analyzed strategic outcomes of this game. So in some sense we assumed the rules of the game to be fixed. What if we could change the rules of the game in order to achieve some objective in strategic equilibrium?

This is the grand question of a field called mechanism design, which we will explore next. As a warm-up we will consider single-item auctions. We will identify an auction mechanism with great properties; these properties will henceforth serve as a gold standard against which we will evaluate our solutions.

1 A Motivating Example

In the lecture I challenged you to bid in two different auctions. The items that I auctioned were two Chocolate bars. Thewere so-called *sealed-bid auctions.* Before running the auction I publicly announced that I would determine the winner and his/her payment according to the following rules:

(Sealed-Bid) First-Price Auction

1. Write your bid b_i on one side of a piece of paper, fold it, write your name on it, and hand it to me.
2. Once I have received all bids, I will determine the winner as the the bidder with the highest bid and make him/her pay *what he/she has bid.*

(Sealed-Bid) Second-Price Auction

1. Write your bid b_i on one side of a piece of paper, fold it, write your name on it, and hand it to me.
2. Once I have received all bids, I will determine the winner as the the bidder with the highest bid and make him/her pay *the next highest bid.*

In both cases, I announced that in the case of a *tie*, i.e., more than one highest bid, I will break ties lexicographically favoring Anna over Paul and so on.

The purpose of this experiment was to get you to think about what to bid. What was your reasoning? Did you bid what the chocolate was worth to you? Or did you bid less, hoping to make a bargain? Did you anticipate what the others would bid or not? Did your reasoning depend on the auction format?

2 A Basic Model

In order to reason about what to do in an auction, we need a model of bidder behavior. We will assume that there is a set of n players $\mathcal{N}$ and a single item for sale. Each player $i \in \mathcal{N}$ has a willingness-to-pay (or *value*) $v_i \in \mathbb{R}_{\geq 0}$. We assume players seek to maximize their utility. If for a given bid profile b, player $i \in \mathcal{N}$ wins he/she has a *utility* of

$$u_i(b, v_i) = v_i - b_i \qquad \text{(first price)}$$
$$u_i(b, v_i) = v_i - \max_{j \neq i} b_i \qquad \text{(second price)}$$

and he/she has a utility of $u_i(b, v_i) = 0$ if he/she loses.

3 First-Price Auction

As you probably have realized yourself, bidding in a first-price auction is not easy. Of course, you could just have bid what the chocolate bar was worth to you. But then, no matter what your colleagues were to bid, you would never make a bargain or positive utility in the terminology that we just defined.

How would you bid if your goal was to maximize your utility? Wouldn't you shade your bid in order to achieve a lower price? But by how much should you shade your bid? The problem is that this depends on what you know about the bids of the others!

In the simplest model, the *complete information model*, one assumes that the players know each other's values.

Definition 7.1. *Let $\epsilon > 0$. A bid profile b is a* (pure) ϵ-Nash equilibrium *for value profile v if for every player $i \in \mathcal{N}$ bid b_i is an ϵ-best response to the bids b_{-i} of the other players. A bid b_i of a player i with value v_i is a* (pure) ϵ-best response *to bids b_{-i} by all other players if $u_i((b_i, b_{-i}), v_i) \geq u_i((b'_i, b_{-i}), v_i) - \epsilon$.*

When $\epsilon = 0$ we refer to the bid profile b as a (pure) Nash equilibrium.

Observation 7.2. *In the first-price auction there always exists a pure ϵ-Nash equilibrium in which a player with the highest value wins the item.*

Observation 7.3. *In the first-price auction letting all players bid their true value is generally not a Nash equilibrium.*

4 Second-Price Auction

It turns out that in the second-price auction bidding is much easier. A bit of thinking reveals that bidding your true value is not only a Nash equilibrium it is, in fact, the best you can do, independent of what your colleagues bid.

Definition 7.4. *A bid profile b is a* dominant strategy equilibrium *for value profile v if for each player $i \in \mathcal{N}$ bid b_i is a (weakly) dominant strategy. A bid b_i is a* (weakly) dominant strategy *for player i with value v_i if for all possible bids b'_i by that player and all possible bids b_{-i} of the other players, $u_i((b_i, b_{-i}), v_i) \geq u_i((b'_i, b_{-i}), v_i)$.*

Theorem 7.5 (Vickrey, 1961)**.** *In a second-price auction, for each player $i \in \mathcal{N}$ it is a dominant strategy to bid truthfully.*

Proof. Fix a player i, his/her value v_i, and the bids b_{-i} of the other players. We need to show that player i's utility is maximized by setting $b_i = v_i$.

Let $b_{max} = \max_{j \neq i} b_j$ denote the highest bid by a player other than i. Note that even though there is an infinite number of bids that i could make, only two distinct outcomes can result. For this we can without loss of generality assume that player i loses if he/she bids $b_i = b_{max}$. In this case if $b_i \leq b_{max}$, then i loses and receives utility 0. If $b_i > b_{max}$, then i wins at price b_{max} and receives utility $v_i - b_{max}$.

We now consider two cases. First, if $v_i \leq b_{max}$, the highest utility that bidder i can get is $\max\{0, v_i - b_{max}\} = 0$, and he/she achieves this by bidding truthfully (and losing). Second, if $v_i > b_{max}$, the highest utility that bidder i can get is $\max\{0, v_i - b_{max}\} = v_i - b_{max}$, and he/she achieves this by bidding truthfully (and winning). □

Observation 7.6. *In a second-price auction, if each player bids his/her true value, then he/she never has negative utility.*

5 Desirable Properties of Auctions

The remarkable properties of the second-price auction, or Vickrey auction, will prove as a very useful benchmark against which we can compare other solutions. We also refer to these properties as our "gold standard".

Gold Standard

1. Strong incentive guarantees. The Vickrey auction is *dominant strategy incentive compatible* (DSIC), i.e., truthtelling is a dominant strategy equilibrium.

2. Strong performance guarantees. At equilibrium, the Vickrey auction maximizes *social welfare* $\sum_{i \in \mathcal{N}} x_i \cdot v_i$, where $x_i \in \{0, 1\}$ indicates whether player i receives the item or not and we require $\sum_{i \in \mathcal{N}} x_i = 1$ for feasibility.

3. Computational efficiency. The Vickrey auction can be computed in *polynomial time.*

We will see that when trying to generalize this result to more complex settings, we often find that obtaining all three properties at once is impossible and we will have to relax one or more of these goals.

6 Single-Parameter Mechanisms

In a *single-parameter mechanism design problem* a set $\mathcal{N}$ of n players (or agents) interacts with a mechanism to select a feasible outcome. Each agent has a private type $\theta_i \in \Theta_i$, which describes her preferences over outcomes. Feasible outcomes correspond to n-dimensional vectors $x \in X$, where $x_i \in \mathbb{R}$ denotes the part of the outcome that player i is interested in. We will restrict attention to settings where the type of agent i can be interpreted as a value v_i or a cost c_i; and player i's value or cost for outcome x is $v_i \cdot x_i$ or $c_i \cdot x_i$.

A (direct) mechanism $M = (f, p)$ consists of an outcome rule $f \colon \Theta \to X$ and a payment rule $p \colon \Theta \to \mathbb{R}^n$.[1] A mechanism asks the players to report their types, which we will denote by b. Think of b_i as player i's bid, reported value in settings where we talk about values and reported cost when players have costs. In the former, $p_i(b)$ will be the payment that the player has to make to the mechanism, in the latter $p_i(b)$ is the payment that the mechanism makes to the player. We make the standard assumption of quasi-linear utilities. That is, player i's utility in the value-case is $u_i^M(b, \theta_i) = v_i \cdot f_i(b) - p_i(b)$ and it is $u_i^M(b, \theta_i) = p_i(b) - c_i \cdot f_i(b)$ when we talk about costs. When it is clear from the context, which mechanism we are referring to we will drop the superscript M.

The basic dilemma of mechanism design is that the mechanism designer (think of a government or company) wants to optimize some global objective such as the social welfare $\sum_{i \in \mathcal{N}} v_i \cdot x_i(b)$ by computing an allocation x based on the bids b, while the players choose their bids b_i so as to maximize their utilities $u_i(b, \theta_i)$.

Example 7.7 (Single-Item Auction)**.** *In a single-item auction n bidders compete for the assignment of an item. Each player can get the item or not, so $X_i = \{0, 1\}$, where we interpret $x_i = 1$ as bidder i gets the item. Then feasible assignments are vectors $x \in X \subseteq \prod_i X_i = \{0, 1\}^n$ with $\sum_i x_i = 1$. Each bidder i has a private value v_i for the item. Our goal is to allocate the item to the bidder with the highest value.*

Example 7.8 (Sponsored Search Auction)**.** *In a sponsored search auction we have n bidders and k positions. Each position has an associated click-through rate α_j, where we assume that positions are sorted such that $\alpha_1 > \alpha_2 > \cdots > \alpha_k > 0$. Feasible allocations are $x \in X$ for which*

[1] Such mechanisms are called direct because they simply ask the players' for their types; we will talk about indirect mechanisms later on.

$x_i \in X_i = \{0, \alpha_k, \ldots, \alpha_1\}$ for all i and for $i \neq j$ we can only have $x_i = x_j > 0$ if $x_i = x_j = 0$. Our goal is to maximize social welfare.

Example 7.9 (Scheduling on Related Machines)**.** *There are n machines, and each player has a private speed s_i. The inverse of the speed $t_i = 1/s_i$ is the time that machine i takes to process a job of unit length. There are m jobs with loads $\ell_1, \ldots, \ell_m$, which need to be allocated to the machines. An allocation induces a work load $W_1, \ldots, W_n$ for each machine. Each machine is interested in maximizing $u_i(b) = p_i(b) - W_i \cdot t_i$, while the mechanism designer wants to minimize the makespan $\max_i W_i \cdot t_i$.*

A very elegant way to resolve the potential conflict of interest between the mechanism designer and the players, is to ensure that it is in each player's interested to bid truthfully. In this case $b_i = \theta_i$ for all $i \in N$ and by choosing an outcome that is optimal for b the mechanism designer chooses an outcome that is optimal for v.

Definition 7.10. *A mechanism $M = (f, p)$ is called dominant strategy incentive compatible (DSIC) (or just truthful), if for each player i bidding $b_i = \theta_i$ is a weakly dominant strategy. That is, for all $i \in \mathcal{N}$, $\theta_i \in \Theta_i$, and all $b \in \Theta$ it holds that*

$$u_i^M((\theta_i, b_{-i}), \theta_i) \geq u_i^M((b_i, b_{-i}), \theta_i) \ .$$

Recommended Literature

- Chapter 9 in the AGT book. (Introduction to the topic)
- Tim Roughgarden's lecture notes `http://theory.stanford.edu/~tim/f13/l/l2.pdf`
- William Vickrey. Counterspeculation, auctions, and competitive sealed tenders. Journal of Finance, 16(1):8–37, 1961.