

Auctions and Resource Allocations

Branislav Bošanský and Michal Pěchouček

Artificial Intelligence Center,
Department of Computer Science,
Faculty of Electrical Engineering,
Czech Technical University in Prague

branislav.bosansky@agents.fel.cvut.cz

December 20, 2016

Previously ... on multi-agent systems.

- 1 Social Choice

Auctions and Mechanism Design

Definition (Auction)

An **auction** is a protocol that allows agents (=bidders) to indicate their interests in one or more resources and that uses these indications of interest to determine both an **allocation of resources** and a set of **payments** by the agents. [Shoham & Leyton-Brown 2009]

Auctions are example of **Mechanism Design** – designing mechanisms (rules, protocols, algorithms), where rational agents participate.

Auctions and Mechanism Design

Types of auctions:

- single item
- multi-item
- one-sided
- two-sided

There are other resource allocation mechanisms (e.g., without money) where resources (possibly bundled) are distributed among the agents.

- facility location
- allocation of divisible goods (cake cutting)
- allocation of indivisible goods (CPU, memory)

Resource Allocations

Two key indicators of social welfare. Aspects of **efficiency** (not computational) include:

- The chosen agreement should be such that there is no alternative agreement that would be better for some and not worse for any of the other agents (**Pareto optimality**).
- If preferences are quantitative, the sum of all payoffs should be as high as possible (**utilitarianism**).

Aspects of **fairness** include:

- No agent should prefer to take the bundle allocated to one of its peers rather than keeping their own (**envy-freeness**).
- The agent that is going to be worst off should be as well off as possible (**egalitarianism**).

Resource Allocations

Utilitarian collective utility function (CFU) is defined as the sum of individual utilities.

Egalitarian collective utility function is defined as the minimum of individual utilities.

Simple Auction Mechanisms

single-item one-sided auctions:

- **English** - The auctioneer sets a starting price for the good, and agents then have the option to announce successive bids, each of which must be higher than the previous bid (usually by some minimum increment set by the auctioneer).
- **Japanese** - The auctioneer sets a starting price for the good that is (continuously) increasing and the agents must confirm that they still want to buy the good for that price.
- **Dutch** - The auctioneer begins by announcing a high price and then successively lower the price. The auction ends when the first agent signals the auctioneer that she buys the good for the current price.

Auctions and Mechanism Design

single-item one-sided auctions (continued):

- **1st price sealed-bid** - Each agent submits to the auctioneer a secret, “sealed” bid for the good that is not accessible to any of the other agents. The agent with the highest bid must purchase the good. In a first-price sealed-bid auction (or simply first-price auction) the winning agent pays an amount equal to his own bid.
- **2nd price sealed-bid** - In a second-price auction the winning agent pays an amount equal to the next highest bid (i.e., the highest rejected bid).

Auctions as Games

What is the optimal strategy in auctions? A strategy in a game.

Definition (Bayesian Game)

A **Bayesian game** is a tuple $\langle \mathcal{N}, \mathcal{A}, \Theta, p, u \rangle$ where

- $\mathcal{N} = \{1, \dots, n\}$ is the set of players
- $\Theta = \Theta_1 \times \dots \times \Theta_n$, Θ_i is the type space of player i
- $\mathcal{A} = \mathcal{A}_1 \times \dots \times \mathcal{A}_n$, A_i is the set of actions where A_i is the set of actions for player i
- $p : \Theta \rightarrow [0, 1]$ is a common prior over types
- $u = u_1, \dots, u_n$, where $u_i : \Theta \rightarrow \mathbb{R}$ is the utility function of player i

Bayes-Nash equilibrium (BNE): rational, risk-neutral players are seeking to maximize their expected payoff, given their beliefs about the other players' types.

Auctions as Games

A sealed bid auction under independent private values (IPV) is a Bayesian game in which

- player i 's actions correspond to his bid b_i
- player types Θ_i correspond to player's private valuations v_i over the auctioned item(s)
- the payoff of a player i corresponds to his/her valuation of the item v_i minus payed bid b_i

Truth Telling in Second-Price Sealed-Bid Auctions







The analysis of auctions/mechanisms is difficult. Often, we want to find mechanisms where truth telling (i.e., $b_i = v_i$) is the optimal (dominant) strategy.

Theorem

Truth-telling is a dominant strategy in a second-price sealed bid auction (assuming independent private values (IPV) model and risk neutral bidders).

Proof We show that the utility for deviating from the true valuation of player i will not be strictly better.

Truth Telling in Second-Price Sealed-Bid Auctions



Dutch and First-price Sealed-bid Auctions

These two mechanisms are strategically equivalent: an agent bids without knowing about the other agents' bids

- a bidder must decide on the amount he's willing to pay, conditioned on having placed the highest bid

Few differences:

- First-price auctions can be held asynchronously
- Dutch auctions are fast, and require minimal communication

There is no dominant strategy, the agents must trade-off between the probability of winning vs. amount paid upon winning.

Individually optimal strategy depends on assumptions about others' valuations (and strategies) – an equilibrium.

Dutch and First-price Sealed-bid Auctions

Equilibria in two-player FPSB Auctions

Assume a first-price auction with two risk-neutral bidders whose valuations are drawn independently and uniformly at random from the interval $[0, 1]$ – what is the equilibrium strategy?

$$\left(\frac{v_1}{2}, \frac{v_2}{2} \right)$$

Can be generalized to the n -player case.

Theorem

In a first-price sealed bid auction with n risk-neutral agents whose valuations v_1, v_2, \dots, v_n are independently drawn from a uniform distribution on the same bounded interval of the real numbers, the unique symmetric equilibrium is given by the strategy profile $\left(\frac{n-1}{n}v_1, \dots, \frac{n-1}{n}v_n \right)$

English and Japanese Auctions

A much more complicated strategy space:

- extensive-form game
- bidders are able to condition their bids on information revealed by others
- in the case of English auctions, the ability to place jump bids

Intuitively, though, the revealed information doesn't make any difference in the independent-private value (IPV) setting.

Theorem

Under the IPV model, it is a dominant strategy for bidders to bid up to (and not beyond) their valuations in both Japanese and English auctions.

Multi-Item (Combinatorial) Auctions

Auctions for bundles of goods.

Let $\mathcal{Z} = \{z_1, \dots, z_m\}$ be a set of items to be auctioned. A valuation function $v_i : 2^{\mathcal{Z}} \rightarrow \mathbb{R}$ indicates how much a bundle $Z \subseteq \mathcal{Z}$ is worth to agent i .

Properties:

- normalization ($v(\emptyset) = 0$)
- free disposal ($Z_1 \subseteq Z_2$ implies $v(Z_1) \leq v(Z_2)$)

Combinatorial auctions do not have to have additive valuation function:

- complementarity: $v(Z_1 \cup Z_2) > v(Z_1) + v(Z_2)$ (e.g. left and right shoe)
- substitutable items: $v(Z_1 \cup Z_2) < v(Z_1) + v(Z_2)$ (e.g. cinema tickets for the same time)

Difficulties in Combinatorial Auctions

There are multiple difficulties that we need to address:

- **Computational complexity:** The allocation problem is computationally hard (NP-complete) even for simple special cases.
- **Representation and communication:** The valuation functions are exponential size objects since they specify a value for each bundle. How can we even represent them? How do we transfer enough information to the auctioneer so that a reasonable allocation can be found?
- **Strategies:** How can we analyze the strategic behavior of the bidders?

Single-Minded Combinatorial Auctions

Definition

A valuation v is called **single minded** if there exists a bundle of items Z^* and a value $v^* \in \mathbb{R}^+$ such that $v(Z) = v^*$ for all $Z \supseteq Z^*$, and $v(Z') = 0$ for all other Z' . A single-minded bid is the pair (Z^*, v^*) .

Theorem

The allocation problem among single-minded bidders is NP-hard. More precisely, the decision problem of whether the optimal allocation has social welfare of at least k (where k is an additional part of the input) is NP-complete.

General Valued Combinatorial Auctions

Consider a linear pricing rule, where a price per each item is available, and the price of each bundle is the sum of the prices of the items in this bundle.

The winner determination problem in combinatorial auctions can be formulated as an integer program (IP).

Definition

Given a set of prices, **the demand** of each bidder is the bundle that maximizes her utility. (There may be more than one such bundle, in which case each of them is called a demand.) Formally, for a given bidder valuation v_i and given item prices p_1, \dots, p_m , a bundle Z is called a demand of bidder i if for every other bundle $Z' \subseteq \mathcal{Z}$ we have that $v_i(Z') - \sum_{j \in Z'} p_j \leq v_i(Z) - \sum_{j \in Z} p_j$.

Walrasian Equilibrium

Definition

A set of nonnegative prices p_1^*, \dots, p_m^* and an allocation S_1^*, \dots, S_n^* of the items is a **Walrasian equilibrium** if for every player i , S_i^* is a demand of bidder i at prices p_1^*, \dots, p_m^* and for any item j that is not allocated we have $p_j = 0$.

Walrasian equilibrium does not always have to exist.

Theorem

A Walrasian equilibrium exists in a combinatorial-auction environment if and only if the corresponding linear programming relaxation of integer program (IP) admits an integral optimal solution.

Auctions on Internet

Web operators (search engines, ad services, ...) sell the position for advertisement based on the user, keyword, web-page, etc.

There are multiple positions, we assume that the higher positions have higher utility (more users will click on the ad, will be influenced by the ad, etc.)

To preserve privacy and allow asynchronous protocols, generalized sealed-bid auctions can be used (we are selling multiple items of different quality).

The game is naturally repeated (the players observe their outcome (their position for their bid) and can update their bids accordingly).

Auctions on Internet


Question

Can we use generalized first-price sealed bid auction?

Advantages:

- simple, easy to understand (each agents pays what they bid)
- cannot be manipulated by the auctioneer by setting (a hidden) reserve price to increase the revenue

Main disadvantages is instability:



(b) 1 week

Fig. from Edelman, Ostrovsky: "Strategic Bidder Behavior in Sponsored Search Auctions"

Generalized Second Price Auctions

Generalization of the second-sealed bid auction.

Each bidder places a bid. The highest bidder gets the first slot, the second-highest, the second slot etc. The highest bidder pays the price bid by the second-highest bidder, the second-highest pays the price bid by the third-highest, etc.

Not a truthful mechanism in general – submitting true valuations is not dominant or equilibrium strategy, there may be multiple equilibria.

Generalized Second Price Auctions

Example:

clickthrough rates	slots	advertisers	revenues per click
10	(a)	(x)	7
4	(b)	(y)	6
0	(c)	(z)	1

Assume that all agents bid their true values.

Agent x 's utility for the first slot is $10 \cdot (7 - 6) = 10$.

However, decreasing their bid to 5 and getting the second slot increases the utility to $4 \cdot (7 - 1) = 24$.

Auctions in a Practical Approach

Consider the agents have a limited budget B_i for a given time-period.

Even GSP is not ideal – consider, for example, a **vindictive bidding** – one can increase the payment of the bidder in the slot above by raising ones bid without affecting ones own payment.

How do the auctioneer allocate the slots?

If we can estimate the future (who is coming, what people are searching) – a simple optimization problem. But this is often difficult (imprecise).

Practical Approach - Dynamic Procedure

Greedy approach: among the bidders whose budgets are not exhausted, allocate the query to the one with the highest bid.

Algorithm 1

Every time a query i arrives, allocate its advertisement space to the bidder j , who maximizes $b_{ij}\varphi(f_j)$, where b_{ij} is the bid for item i , f_j is the fraction of the bidder j 's budget which has been spent so far, and $\varphi(x) = 1 - e^{x-1}$.

Theorem

The revenue of Algorithm 1 is at least $1 - 1/e$ of the optimum revenue.