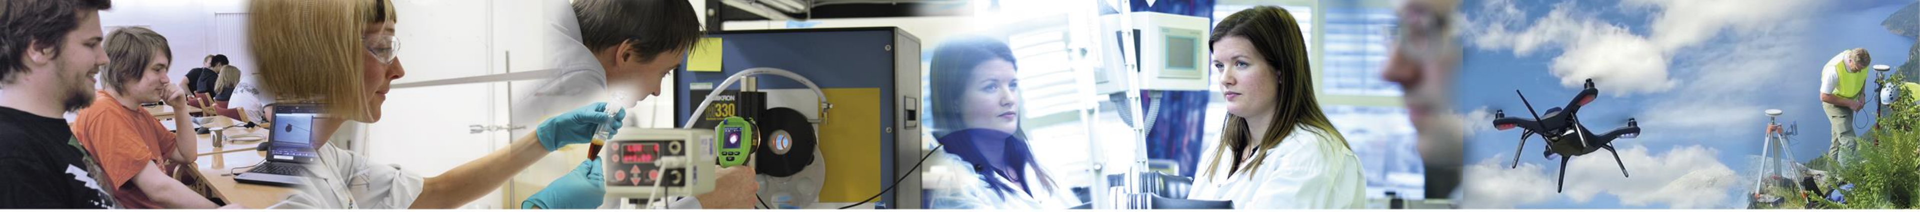

UiO : Department of Technology Systems
University of Oslo

TEK5010/9010 - Multiagent systems 2021
Lecture 11

Auctions

Jonas Moen

Highlights lecture 11 – Auctions*

- Auction terminology

- Singel item auctions
  - English, Dutch, First-price sealed-bid, Vickrey

- Combinatorial auctions
  - VCG mechanism

*Wooldridge, 2009: chapter 14

Auctions

Auctions are mechanisms for allocating scarce resources (items or bundles of items) to agents.

1. Must be a scarce good.
2. Must be desired by more than one agent.

If not, the allocation is trivial.

Auctions

Online auctions of large, international audiences make possible low cost regular auctions in everyday life of people. Earlier auctions were for the sale of expensive objects like paintings, antiques and so forth.

Auctions

Auctions are efficient in the way that the resources are allocated to the agent that value them the most.

Auction terminology

Agents in auctions:

1. Auctioneer is the provider of the scarce resource.
2. Bidders are agents that desire the scarce good.

Auctions

The goal of an auction is for the auctioneer to allocate a good (the scarce resource) to one of the bidders.

- The auctioneer tries to do this through an appropriate auction mechanism, maximizing auctioneer utility.
- The bidders try to achieve their desires through an effective bidding strategy, minimizing cost.

Auction terminology

The value of a good

1. Common value
   The shared value representative of all bidders.
2. Private value
   The hidden value of single bidder that are diffrent from the common value.
3. Correlated value
   Private value depending on other agents’ value. Typically in the case of reselling the goods in the future.

Auction terminology

Winner determination:

1. First-price auction
   The good is allocated to the highest bidder at the price of highest bid.
2. Second-price auction
   The good is allocated to the highest bidder at the price of the second highest bid.

Auction terminology

Information sharing:

1. Open-cry auction
   Every agent can see what the other agents are bidding.
2. Sealed-bid auction
   Agents are not able to determine the bids made by other agents.

Auction terminology

Bidding process:

1. One-shot auction
   Single round of bidding, after where auctioneer allocate the good.
2. Ascending auction
   A reservation price is called out and bidders are allowed to successively increase their bids. The auctioneer allocates the good when there are no more bids.

Auction terminology

Bidding process:

3. Descending auction
   The auctioneer calls out a high price and successively decrease the price until a bid is made (or the reservation price is met). The auctioneer allocates the good to the first bidder.

Auction mechanism design

Auction mechanism design is the analysis of how different auctions affect the agents/auctioneer and their outcomes.

There are two types of auctions:

1. Singel item auctions
   Allocation of a singel good.
2. Combinatorial auctions
   Allocation of a bundle of items.

Singel item auctions

Some well-known single item auction:

1. English auction
2. Dutch auction
3. First-price, sealed-bid auction
4. Vickrey auction

English auction

First-price, open-cry, ascending auction

The most commonly known type of auction.

English auction

Auction process:

1. The auction starts by auctioneer calling out of a reservation price. If no bids are recieved over reservation price the good is allocated to the auctioneer.
2. Agents bid more than current highest bid.
3. When no agent is willing to raise the bid, the good is allocated to the highest bidder at the price of the highest bid.

English auction

Bidder strategy:

The dominant strategy of the bidders is for an agent to successively bid a small amount more than current highest bid. This is done until the price bid reaches the private valuation, whereby the agent should withdraw from the auction.

Dutch auction

First-price, open-cry, descending auction

English auction in ‘reverse’.

Dutch auction

Auction process:

1. Auctioneer starts by calling out a high value price asking.
2. Auctioneer then continually lower the price in small increments, until some agent make a bid at current price.
3. The good is allocated to the bidder at the bidder price.

First-price, sealed-bid auctions

Auction process:

1. One-shot bid.
2. The good is allocated to the highest bidder at the price of highest bid.

First-price, sealed-bid auctions

Bidder strategy:

Optimal strategy of the bidder is to bid a little less than the true valuation. How much depends on the other bidders.

Vickrey auctions

Second-price, sealed-bid auction

Have received a lot of attention in the multiagent literature. [Sandholm, 1999]

Vickrey auctions

Auction process:

1. One-shot bid.
2. The good is awarded to the highest bidder at the price of the second highest bid.

Vickrey auctions

Bidder strategy:

The best strategy in Vickrey auction is for the agent to bid the true valuation of the good.

This makes Vickrey auction immune to strategic manipulation.

Vickrey auctions

Proof of bidder strategy (informal):

Assume agent bid more than the true valuation v_0. This strategy runs the risk of agent being awarded the good at a price over v_0.

Assume agent bid less than v_0. The agent then lowers the possibility of receiving the good, but at the same second highest price.

Vickrey auctions

Vickrey auctions can lead to anti-social behaviour if opponents valuation function is ‘known’. Agents could set their bid at just below the highest bid of opponent in order to incerase the cost for opponent if they receive the good.

Winner’s curse

Agents that have been awarder the good on offer, and especially if there are uncertainties about the true value, run the risk of having overvalued the good.

Expected revenue of the auctioneer

The expected revenue of the auctioneer is affected by the risk profile of the bidders and the auctioneer.

1. Risk-neutral bidders
   Under some simple assumptions the expected revenue is equal in all single item auctions.
2. Risk-averse bidders (get the good at slight price increase)
   Dutch and First-price, sealed-bid do better, due to agents are able to ‘insure’ themselves by bidding slightly more than in the risk-neutral case.

Expected revenue of the auctioneer

3. Risk-averse auctioneers
   Do better in English and Vickrey.

Lies and collusion

1. Non of the protocols are immune to collution, but best response is to try to hide the identity of agents from each other.

2. Auctioneer could lie in Vickrey, but not in the other auction types.

3. Shills
   Bogus players.

Counterspeculations

Bidders can engage in activities to obtain information about either:

1. The true value of the good.
2. The true valuation function of the other agents.

This is costly and has to be evaluated against the increased payoff.

Combinatorial auctions

Bidders have preferences over possible bundles of goods offered by the auctioneer.

Let $\mathcal{Z} = \{z_1, \ldots, z_M\}$ be a set of M items to be auctioned.
For all agents $i \in Ag$ we have a valuation function:

$$v_i \colon 2^{\mathcal{Z}} \to \mathbb{R}$$

so that for every bundle of goods $Z \subseteq \mathcal{Z}$ the value $v_i(Z)$ indicates how much Z would be worth to agent i.

Combinatorial auctions

Properties of the valuation function:

1. Normalization
   $$v_i(\emptyset) = 0$$

2. Free disposal
   $Z_1 \subseteq Z_2$ implies $v_i(Z_1) \leq v_i(Z_2)$

Combinatorial auctions

Winner determination problem is that of maximizing the utility of the auction.

In many cases the auctioneer is maximizing the social welfare of the outcomes, i.e. the sum of the valuation functions over all bidders.

Combinatorial auctions

Winner determination problem (assuming maximization of social welfare):

$$(Z_1^*, \dots, Z_N^*) = arg \max_{(Z_1,\dots,Z_N) \in alloc(\mathcal{Z}, Ag)} sw(Z_1, \dots, Z_N, v_1, \dots, v_N)$$

where $Z_1^*, \dots, Z_N^*$ is the optimal allocation
$sw(Z_1, \dots, Z_N, v_1, \dots, v_N) = \sum_{i=1}^{N} v_i(Z_i)$ is the social welfare function

Combinatorial auctions

Winner determination problem (assuming maximization of social welfare):

$$(Z_1^*, \dots, Z_N^*) = arg \max_{(Z_1,\dots,Z_N) \in alloc(\mathcal{Z}, Ag)} sw(Z_1, \dots, Z_N, v_1, \dots, v_N)$$

where $alloc(\mathcal{Z}, Ag)$ is the set of all allocations of goods $\mathcal{Z}$ over agents $i \in Ag$ that satisfies $Z_i \cap Z_j = \emptyset$, meaning that no good is allocated to more than one agent.

Combinatorial auctions

Winner determination problem (assuming maximization of social welfare):

$$(Z_1^*, \dots, Z_N^*) = arg \max_{(Z_1,\dots,Z_N) \in alloc(\mathcal{Z}, Ag)} sw(Z_1, \dots, Z_N, v_1, \dots, v_N)$$

This is a classic combinatorial optimization problem.

1. NP-hard to represent and calculate
2. Private valuation function $v_i \rightarrow \hat{v}_i$, i.e. strategic manipulation.

Combinatorial auctions

Representational complexity

The representational complexity of the valuation function can be high. For instance in a bandwidth auction of 1122 licenses we need 2^{1122} entries to represent the valuation function $v(Z)$.

Bidding languages can make a succinct and complete representation of the valuation function.

Combinatorial auctions

Atomic bids of a bundle $Z \subseteq \mathcal{Z}$ and $p \in \mathbb{R}_+$representing the price an agent i is willing to pay for Z:

$$\beta = (Z, p)$$

defines a valuation function of a bundle Z' such that

$$v_\beta(Z') = \begin{cases} p \text{ if } \mathrm{Z}^{'} \text{satisfies } (Z, p), \text{meaning } Z \subseteq Z' \\ 0 \text{ otherwise} \end{cases}$$

Combinatorial auctions

Example:

$$\beta = (\{a, b\}, p)$$

is satisfied for $\{a, b, c\}$ but not for $\{a\}$ or $\{a, c, d\}$

Combinatorial auctions

XOR bids

In *exclusive or bids* we must pay for at most one of the atomic bids to be satisfied:

$$\beta = (Z_1, p_1) XOR \cdots XOR (Z_k, p_k)$$

Then such a bid defines a valuation function

$$v_\beta(Z') = \begin{cases} 0 \text{ if } Z' \text{does not satisfy any of the atomic bids} \\ max\{p_i | Z_i \subseteq Z'\} \end{cases}$$

Combinatorial auctions

XOR bids are fully expressive, meaning that any valuation function can be expressed. Calculation of the valuation function can be done in polynomial time. [Nisan, 2006]

Combinatorial auctions

Example:

$$\beta_1 = (\{a,b\},3\ )XOR(\{c,d\},5\ )$$

$$v_{\beta_1}(\{a\}) = 0$$
$$v_{\beta_1}(\{b\}) = 0$$
$$v_{\beta_1}(\{a,b\}) = 3$$
$$v_{\beta_1}(\{c,d\}) = 5$$
$$v_{\beta_1}(\{a,b,c,d\}) = 5$$

Combinatorial auctions

OR bids

In *or bids* we are willing to pay for more than one bundle:

$$\beta = (Z_1, p_1) OR \cdots OR (Z_k, p_k)$$

An atomic bid offer W is such that every atomic bid in β is satisfied by $Z' \subseteq \mathcal{Z}$ and mutually disjoint $Z_i \cap Z_j = \emptyset$ for all agents. There is also no subset W' better than the atomic bid W.

Combinatorial auctions

Example:

$$\beta_2 = (\{a,b\},3\,)OR(\{c,d\},5\,)$$

$$v_{\beta_2}(\{a\}) = 0$$
$$v_{\beta_2}(\{b\}) = 0$$
$$v_{\beta_2}(\{a,b\}) = 3$$
$$v_{\beta_2}(\{c,d\}) = 5$$
$$v_{\beta_2}(\{a,b,c,d\}) = 8$$

Combinatorial auctions

Integer linear programming (ILP)

$$\begin{array}{ll} \text{max}imize & f(x_1, \dots, x_k) \\ subject\ to\ constraints & \varphi_1(x_1, \dots, x_k) \\ & \vdots \\ & \varphi_l(x_1, \dots, x_k) \end{array}$$

where $f(x_1, \dots, x_k)$ is the linear objective function
$\varphi's$ are the linear constraints on decision variables x.

Combinatorial auctions

Combinatorial problems as ILPs [Blumrose and Nisan, 2007].

$$\max \quad \sum_{i \in Ag, Z \subseteq \mathcal{Z}} x_{i,Z} v_i(Z)$$

$s.t.$

1. $\sum_{i \in Ag, Z \subseteq \mathcal{Z} | z \in Z} x_{i,Z} \leq 1$ for all $z \in Z$
2. $\sum_{Z \subseteq \mathcal{Z}} x_{i,Z} \leq 1$ for all $i \in Ag$
3. $x_{i,Z} \geq 0$ for all $i \in Ag, Z \subseteq \mathcal{Z}$

Combinatorial auctions

Combinatorial problems as ILPs [Blumrose and Nisan, 2007].

1. We do not allow any good to be distributed to more than one agent.
2. Each agent is only allocated one bundle of goods.
3. $x_{i,Z}$ is the decision variable and is either 0 or 1.

Combinatorial auctions

Combinatorial problems as ILPs [Blumrose and Nisan, 2007].

ILPs are NP-hard (at least in worst case).

Alternative methods are

1. Approximate methods (often imprecise in result).
2. Heuristics like greedy search, evolutionary optimization, etc.

Combinatorial auctions

The Vickrey-Clarke-Groves (VCG) mechanism

Strategic manipulation is easy in auctions. A mechanism to deal with this problem is needed.

The VCG mechanism is a generalization of the Vickrey mechanism extended to combinatorial auctions.

A VCG mechanism is *incentive compatible*, meaning that telling the true validation function is rational.

Combinatorial auctions

The Vickrey-Clarke-Groves (VCG) mechanism

The basic idea:

The winners pays a compensation (to the mechanism) equal to the «lost utility» of all other players by winning the auction.

Combinatorial auctions

The Vickrey-Clarke-Groves (VCG) mechanism

1. Every agent $i \in Ag$ declares its valuation function $\hat{v}_i$.
2. Calculate the optimal allocation $Z_1^*, \ldots, Z_N^*$ using the sw-function.
3. Every agent $i \in Ag$ pays p_i to the mechanism:

$$p_i = sw_{-i}(Z_1', \ldots, Z_N', \hat{v}_1, \ldots, v_i^0, \ldots, \hat{v}_N)$$
$$-sw_{-i}(Z_1^*, \ldots, Z_N^*, \hat{v}_1, \ldots, \hat{v}_i, \ldots, \hat{v}_N)$$

Combinatorial auctions

The Vickrey-Clarke-Groves (VCG) mechanism

No agent can benefit by declaring anything other than its true valuation function.

Easy to see when one single good is to be allocated using VCG, i.e. classic Vickrey auction.

The VCG is NP-hard [Müller 2006, Lavi 2007].

Combinatorial auctions

Example: VCG auction of Facebook ads*

Advertisers $Ag = \{1,2,3,4\}$ are bidding for two impressions $\mathcal{Z} = \{a, b\}$ on Facebook:

$$\beta_1 = (\{a\}, 11) XOR (\{b\}, 11)$$
$$\beta_2 = (\{a\}, 7) XOR (\{b\}, 7)$$
$$\beta_3 = (\{a\}, 5) XOR (\{b\}, 5)$$
$$\beta_4 = (\{a\}, 3) XOR (\{b\}, 3)$$

*https://www.linkedin.com/pulse/how-vcg-auction-mechanism-works-behind-facebook-ad-santosh-yaduvanshi

Combinatorial auctions

Example: VCG auction of Facebook ads

The winners are the highest bidders, giving the allocation:

Goods	Player 1	Player 2	Player 3	Player 4
$\{a\}$	11	0	0	0
$\{b\}$	0	7	0	0

Combinatorial auctions

Example: VCG Auction of Facebook ads

Remove player 1 from the auction gives the allocation:

Goods	Player 1	Player 2	Player 3	Player 4
$\{a\}$	11	7	0	0
$\{b\}$	0	0	5	0

$$p_1 = (7 - 7) + (5 - 0) = 5.$$

Combinatorial auctions

Example: VCG Auction of Facebook ads

Remove player 2 from the auction gives the allocation:

Goods	Player 1	Player 2	Player 3	Player 4
$\{a\}$	11	0	0	0
$\{b\}$	0	7	5	0

$$p_2 = (11 - 11) + (5 - 0) = 5.$$

Online auctions - auctions in practice

Auction bots, web agents that buy and sell goods on behalf of users. User specify time of closure, desired price, limiting prices, ascend/descend bid and asking functions, etc

1. Kasbah [Chavez and Maes, 1996] with its own agentbased market.
2. Ebay using proxy bidding.
3. Adwords auction by companies like Facebook, Google, etc
4. The Trading Agent Competition [Wellman et *al.*, 2007] uses simulated markets.

Highlights lecture 11 – Auctions*

- Auction terminology

- Singel item auctions
  - English, Dutch, First-price sealed-bid, Vickrey

- Combinatorial auctions
  - VCG mechanism

*Wooldridge, 2009: chapter 14