

Game Theory 2024

Ulle Endriss

Institute for Logic, Language and Computation

University of Amsterdam

Plan for Today

This and the next lecture are going to be about *mechanism design*, which has been called “inverse game theory” by some.

The goal is to design a game that encourages certain player behaviour.

Today we are going to focus on the design of *auctions*:

- players are bidders in an auction trying to obtain certain goods
- we want to incentivise them to be truthful about their valuations

A good starting point for finding out more about auctions is the survey by McAfee and McMillan (1987).

R.P. McAfee and J. McMillan. Auctions and Bidding. *Journal of Economic Literature*, 25(2):699–738, 1987.

Auctions for a Single Item

General setting for the most basic type of auction:

- one seller (the *auctioneer*)
- many potential buyers (the *bidders*)
- one single item to be sold:
 - auctioneer has a *reservation price* (under which she won't sell)
 - each bidder i has her own private *valuation* v_i of the item

Except for the reservation price (could be 0), we ignore the preferences of the auctioneer and regard this as a game played by the bidders.

We assume *quasi-linear utilities*. Thus, if you win the auction, then:

$$\text{your utility} = \text{your valuation of the item} - \text{price paid for it}$$

The Winner's Curse

Our assumption that you just 'have' a valuation is a simplifying one.

Imagine you want to buy a house:

- If you want to live in the house, then it's a *private-value auction* and you can (maybe) be assumed to know your own valuation.
- If you see it as an investment, then it's a *common-value auction* and your true valuation depends on the valuations of others.

Winner's curse: If you win but have been uncertain about the true value of the item, should you actually be happy?

Auctions, Mechanisms, Games

An *auction* is a *mechanism* for selling items from one seller to many potential buyers (or for one buyer to buy from many potential sellers).

Think of a mechanism as a *game* played by the bidders (the players).

The rules of the game are *designed* by the auctioneer, who has a particular objective in mind (e.g., to maximise revenue or to obtain truthful information from the bidders).

Remark: Today we are going to avoid formal definitions of the specific mechanisms we study. They are special cases of the (formally defined) mechanisms that are going to be the subject of the next lecture.

English Auctions

Well-known protocol to auction off paintings, antiques, etc.:

- auctioneer opens by announcing her reservation price
- in each round, every bidder can increment the price (by $\geq \epsilon$)
- final bid wins

Exercise: *What's the best strategy to use here? Is it dominant?*



Dutch Auctions

Dutch auctions have the advantage of being very fast:

- auctioneer opens by announcing an overly high initial price
- price is lowered a little bit in each round
- first bidder to accept to buy at the current price wins the auction

Exercise: *What's the best strategy to use here? Is it dominant?*



Fun Fact: Used to sell $\sim 43\text{M}$ flowers/day in Aalsmeer near Amsterdam.

Sealed-Bid Auctions

Sealed-bid (as opposed to *open-cry*) auctions are used for awarding public building contracts and the like:

- each bidder submits a bid in a sealed envelope
- the highest bid wins (unless it's below the reservation price)


This is a *direct-revelation mechanism*, as opposed to the indirect ones we've seen before: bids are simply—truthful or fake—valuations.

Exercise: *What's the best strategy to use here? Is it dominant?*

Let's Play: Sealed-Bid Auction

You will be assigned to a *group* and receive your *true valuation*.

All valuations were drawn from the binomial distribution $B(20, 0.5)$:



Rules: You submit a bid for any amount you like (any real number). You win your group's auction if you submit the highest bid. If you do, your utility is the difference between your true valuation and your bid. Any group winner with *strictly positive utility* gets €10 and the one with the highest utility gets an extra €50. Equal shares in case of ties.

Analysis

It is difficult to figure out a good strategy for bidding:

- If you bid too low, you might lose, even if your valuation is high.
- If you bid too high, even if you do win, your utility will be small.

Dutch Auction vs. Sealed-Bid Auction

The Dutch and the sealed-bid auction are *strategically equivalent*.

As a bidder i with valuation v_i , in both cases, you reason as follows:

- Try to estimate the highest bid \hat{v}^* of your competitors.
- Bid $\hat{v}^* + \epsilon$ if $\hat{v}^* + \epsilon < v_i$, i.e., in case winning would be profitable.
Bid $v_i - \epsilon'$ otherwise (just in case your estimate was wrong).

In practice, there still are some differences between the two auctions (e.g., time pressure in a Dutch auction might affect how people bid).

Vickrey Auctions

A *Vickrey auction* is a sealed-bid *second-price* auction:

- each bidder submits a bid in a sealed envelope
- the bidder with the highest bid wins, but pays the price of the second highest bid (unless it's below the reservation price)

Exercise: *What's the best strategy to use here? Is it dominant?*

Incentive Compatibility of Vickrey Auctions

A direct-revelation auction mechanism (in which you bid by submitting a valuation) is called *incentive-compatible* (or *strategyproof*) in case submitting your *true valuation* is always a *dominant strategy*.

Theorem 1 (Vickrey, 1961) *Vickrey auctions are incentive-compatible.*

Proof: Consider bidder i with true valuation v_i who is choosing between reporting v_i and some fake \hat{v}_i :

- Suppose v_i would be a winning bid. This would be profitable for i . Then switching to \hat{v}_i either makes no difference or results in a losing bid. ✓
- Suppose some $\hat{v}^* > v_i$ would win. Then any $\hat{v}_i < \hat{v}^*$ would still lose, while any $\hat{v}_i \geq \hat{v}^*$ would win and result in a loss in utility. ✓



William Vickrey
(1914–1996)

W. Vickrey. Counterspeculation, Auctions, and Competitive Sealed Tenders. *Journal of Finance*, 16(1):8–37, 1961.

The Revenue Equivalence Theorem

Which mechanism should the auctioneer choose? Surprisingly, in some sense and under certain assumptions, it does not matter:

Theorem 2 (Vickrey, 1961) *Our **four auction mechanisms** all give the **same expected revenue**, provided bidders are risk-neutral and valuations are drawn independently from the same uniform distribution.*

Proof sketch: expected revenue = second highest (true) valuation.

- **Vickrey**: dominant strategy is to be truthful ✓
- **English**: bidding stops when second highest valuation is reached ✓
- **Dutch/first-price sealed-bid**: *in expectation*, bidders correctly estimate each others' val's, so winner bids second highest val ✓

Remark: Above assumptions are too strong to be satisfied in practice.

W. Vickrey. Counterspeculation, Auctions, and Competitive Sealed Tenders. *Journal of Finance*, 16(1):8–37, 1961.

Combinatorial Auctions

Selling several items through several single-item auctions in sequence is not a good idea (imagine auctioning off a left and a right shoe) ...

A *combinatorial auction* setting is a tuple $\langle N, G, \mathbf{v} \rangle$, where:

- $N = \{1, \dots, n\}$ is a finite set of *bidders*,
- $G = \{1, \dots, \ell\}$ is a finite set of *goods*, and
- $\mathbf{v} = (v_1, \dots, v_n)$ is a profile of (true) *valuations* $v_i : 2^G \rightarrow \mathbb{R}_{\geq 0}$.

The valuations are private information. Each bidder i submits a *bid* by reporting a—true or fake—valuation $\hat{v}_i : 2^G \rightarrow \mathbb{R}_{\geq 0}$.

A combinatorial auction *mechanism* consists of (i) an *allocation rule*, mapping \hat{v} to bundles (B_1, \dots, B_n) s.t. $B_i \subseteq G$ and $B_i \cap B_j = \emptyset$, and (ii) a *pricing rule*, mapping \hat{v} to a price vector $(p_1, \dots, p_n) \in \mathbb{R}^n$.

The (quasi-linear) *utility* of bidder i derived from such an outcome is:

$$u_i(B_i, p_i) = v_i(B_i) - p_i$$

Topics in Combinatorial Auctions

We postpone the game-theoretical analysis of CA's to the next lecture, and briefly review some other topics of interest in this context . . .

P. Cramton, Y. Shoham, and R. Steinberg (eds.). *Combinatorial Auctions*. MIT Press, 2006.

Aside: Bidding Languages

Bidding in a CA requires reporting a valuation $\hat{v}_i : 2^G \rightarrow \mathbb{R}_{\geq 0}$, i.e., declaring a value for *every possible subset* of the set of goods G ...

In practice, bidders will only care about certain bundles of goods.

Use a *bidding language* to encode valuations. Examples:

- *Single-minded bids*: report one bundle B_i and a price p_i to express $\hat{v}_i : B' \mapsto p_i$ for all $B' \supseteq B_i$ and $\hat{v}_i : B' \mapsto 0$ for all other bundles.
- *Weighted formulas*: report several weighted formulas (φ, w) , with propositional variables in G , to express $\hat{v}_i : B \mapsto \sum_{(\varphi, w)} w \cdot \mathbb{1}_{B \models \varphi}$

More generally, the design and analysis of formal languages for the compact *representation of preferences* is an important research topic.

N. Nisan. Bidding Languages. In P. Cramton et al. (eds.), *Combinatorial Auctions*. MIT Press, 2006.

Y. Chevaleyre, U. Endriss, J. Lang, and N. Maudet. Preference Handling in Combinatorial Domains: From AI to Social Choice. *AI Magazine*, 29(4):37–46, 2008.

Aside: Complexity of Combinatorial Auctions

Typically, the *allocation rule* requires us to compute an allocation that maximises utilitarian *social welfare* relative to the valuations reported:

$$(B_1^*, \dots, B_n^*) \in \underset{\substack{(B_1, \dots, B_n) \subseteq G^n \\ \text{s.t. } B_i \cap B_j = \emptyset \text{ for } i \neq j}}{\operatorname{argmax}} \sum_{i \in N} \hat{v}_i(B_i)$$

Unfortunately, this problem of **WINNER DETERMINATION** is *NP-hard*, even for the case of single-minded bidders (Rothkopf et al., 1998).

The proof is by reduction from SET PACKING.

An important research topic is to identify restrictions of the general setting of practical relevance where this problem is tractable.

M.H. Rothkopf, A. Pekeč, and R.M. Harstad. Computationally Manageable Combinational Auctions. *Management Science*, 44(8):1131–1147, 1998.

Aside: Solving Combinatorial Auctions

Sometimes the kind of WINNER DETERMINATION problem we face in practice is (theoretically) intractable. Yet, we still want to solve it.

Designing *algorithms* for solving large combinatorial auction instances, for various bidding languages, is yet another important research topic.

Most existing work is based on *mathematical programming* techniques (mixed integer programming) developed in Operations Research and *heuristic-guided search techniques* developed in AI.

T. Sandholm. Optimal Winner Determination Algorithms. In P. Cramton et al. (eds.), *Combinatorial Auctions*. MIT Press, 2006.

Summary

This has been an introduction to basic auction theory:

- four basic mechanisms: English, Dutch, FPSB, Vickrey auction
- Vickrey auction: simple protocol + incentive-compatible
- all four mechanisms generate the same expected revenue
- combinatorial auctions for selling bundles of goods

We've also discussed issues regarding combinatorial auctions that are not about game theory but rather about knowledge representation, complexity theory, and algorithm design.

What next? Generalising Vickrey's idea to combinatorial auctions and even more general mechanisms for collectively choosing an outcome.