# Auction Theory: an Introduction

Dongmo Zhang

Intelligent Systems Lab
University of Western Sydney
Australia

Thematic Trimester on Game Theory @ IRIT, France
6 July 2015

# Objectives of the tutorial

An auction is a process of resource allocation and price discovery on the basis of bids from participants. A huge volume of economic transactions is conducted through auctions, including the ones generated in electronic markets. Auctions have also been widely used as mechanisms for multi- agent interaction, job assignment and resource allocation. This tutorial discusses auction mechanisms from AI perspectives. We will talk about single-sided auctions, double auctions, combinatorial auctions, dynamic auctions and ad auctions. We will show how auction theory can be used in theoretical research of artificial intelligence and the practice of electronic market development.

# Outline

# Auction theory and practice: an overview

- Auctions are trading mechanisms that are used for public sales of goods on the basis of bids from participants.
- The use of auctions can be traced back to 500BC in Babylon.
- Modern use of auctions can be found anywhere in real life: government procurement, foreign exchange, stock exchange, future market, property sell, eBay, ad auction, Groupon, ...
- Auctions also provide simple, well-defined and efficient mechanisms for multi-agent interaction, job assignment, task and resource allocation, ...
- Auction theory is a central prat of economics. The Nobel prize winning theories, such as Incentive Theory and Revenue Equivalence Theorem, were initiated from auction theory (Vickrey 1996, Myerson 2007 Nobel Prizes).

# Auction parameters

- Goods to sell may be
  - single object: one indivisible item only
  - multiple units: multiple units of the same goods
  - multiple items: multiple items of different goods
- Roles in an auction may have:
  - auctioneer: the agent who coordinates the auction.
  - buyer: the agent who wants to buy the goods (placing bids).
  - seller: the agent who wants to sell the goods (placing asks).

An auction may have a single seller (buyer) and multiple buyers (sellers), or multiple buyers and multiple sellers.


# Auction parameters

- Goods can have:
  - private value: how much the goods is worth to you (buyer/seller).
  - public value: how much the goods is worth to everybody.
  - correlated value: how much you'd like to pay for the goods.
- Bidding may be:
  - one shot: one round determine the winner.
  - ascending: start from low price, end with high price
  - descending: start from high price, end with low price
- Bids may be:
  - open cry: every agent can see other agent's bids.
  - sealed bid: only auctioneer can see every agent's bid.
- Winner determination may be:
  - first price: the agent that bids most gets the good and pay its bidding price.
  - second price: the agent that bids most gets the good and pay the second highest bid price.

# Auction mechanisms: an overview

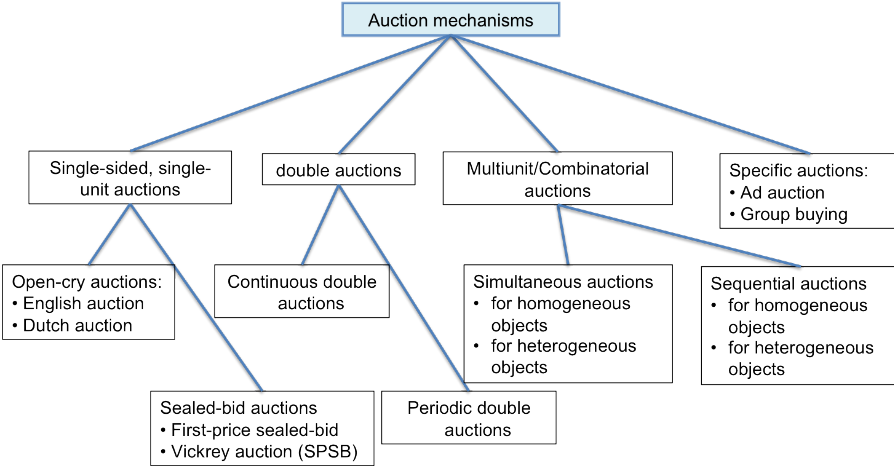

# Outline

# Single-sided single-unit auctions

- A single-sided auction takes place between the auctioneer, acting for a single seller, and a collection of buyers known as the bidders.
- The goal of the auction is for the auctioneer to allocate a single unit of a good to one of the bidders.
- In most settings the auctioneer desires to maximise the price; bidders desire to minimise price.
  - English Auction
  - Dutch Auction
  - First-price sealed-bid (FPSB) auction
  - Second-price sealed-bid (SPSB), aka Vickrey auction

| | First price | Second price |
|---|---|---|
| Open | Dutch | English |
| Closed | FPSB | SPSB |

# The English Auction

- Procedure: The auctioneer starts off by posting a reserve price. If no buyer bids the reserve price, then no deal is done; otherwise, bids are invited from the buyers, who must bid higher than the current highest bid. When no buyer is willing to raise the bid, the buyer who made the highest bid wins the object and pays the amount of his bid.
- Characteristics: open cry, ascending, price discovery
- Commonly used to sell: arts, antiques, properties, wine, ... Widely used in online markets.
- Susceptible to: winners curse and shills.

# The Dutch Auction

- Procedure: The auctioneer starts out by asking for an artificially high price and continually lowers the asking price by a small value until some buyer makes a bid equal to the current asking price. The buyer who made the bid wins the object and pays the current price.
- Characteristics: open-cry, descending, fast.
- Typically used in: Dutch flower market, Ontario tobacco market, fish markets in Zamibia,
- Susceptible to: winners curse.

# First-Price Sealed-Bid Auctions

- Procedure: Bidders submit bids to the auctioneer. In a buyer-bid auction, the highest bidder wins the auction and pays the amount of his bid. In a seller-bid auction, the lowest bidder sells the object and is paid the amount of her bid.
- Characteristics: first-price, sealed-bid.
- Commonly used for: invited tenders, construction contracting, military procurement and private-firm procurement, refinancing credit, London Gold Exchange, ...
- Susceptible to: winners curse.

# Second-Price Sealed-Bid Auction (Vickrey Auction)

- Procedure: Bidders submit bids to the auctioneer. In a buyer-bid auction, the highest bidder wins the auction and pays the amount of the second highest bid. In a seller-bid auction, the lowest bidder sells the object and is paid the amount of the second lowest bid.
- Characteristics: second-price, sealed-bid, truth telling.
- Usage: not as common as other auctions but eBay's proxy bidding and Google's AD auction use similar idea.

# The Revenue Equivalence Theorem

- As a seller, which type of auction mechanisms should you choose to gain the highest revenue?
- Independent private values(IPV) model:
  - Bidders: $N = \{1, 2, 3, \cdots, n\}$
  - Each bidder $i$ values the object on sale at $v_i$, which is private information to her but it is common knowledge that each $v_i$ is independently drawn from the same continuous distribution $F(v)$ on $[\underline{v}, \overline{v}]$ (so $F(\underline{v}) = 0$, $F(\overline{v}) = 1$) with density $f(v)$.
  - Each bidder is risk-neutral.

# The Revenue Equivalence Theorem

- Revenue Equivalence Theorem: Any mechanism in IPV model yields the same expected revenue if
  - the object always goes to the bidder with highest value;
  - any bidder with lowest value expects zero utility
- The four standard auction mechanisms give the same expected revenue.

Revenue Equivalence Theorem was initially proposed by Vickrey (1962). Myerson (1981) extended it into more general setting, which formed the most fundamental result in mechanism design.

# Outline

# Double auction mechanisms

- A double auction is an auction mechanism that allows buyers and sellers submit bids/asks simultaneously.
- Double auction markets usually feature a large number of buyers and sellers, and thus participants tend to incur lower transaction costs.
- Most financial markets, such NYSE and NASDAQ, use double auction mechanisms.

# Double auction market structure

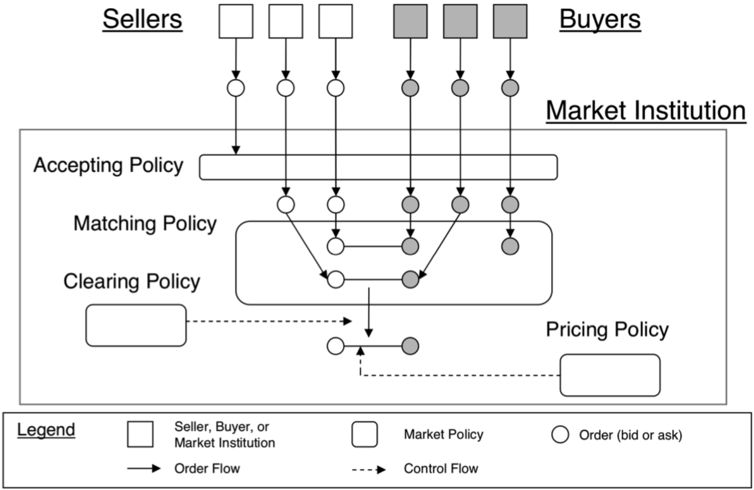

# Shouts

- shout: either a bid (for buying) or an ask (for selling)
- bid shout: the highest price to buy
- ask shout: the lowest price to sell.
- match: An ask shout $p_a$ and a bid shout $p_b$ is matchable if $p_a \leq p_b$
- cleaning price: the transaction price for a matched pair. Can be anything in $[p_a, p_b]$.
- Example:
  1. asks: 50, 44, 52, 80, 55, 48, 60
  2. bids: 34, 36, 52, 40, 63, 47, 48
  3. Matched shouts: (50, 52), (44, 63). Any more matches?
  4. Clearing price for (50, 52) can be anything in between, say 52.

# Traders utility

- Traders: buyers and sellers.
- Each trader $i$ has a private value of the trading good $v_i$.
- For a successful transaction, if the clearing price is $p$, then the utility of trader $i$ ( profit margin) is
  1. $u_i = p - v_i$, if the trader is a seller
  2. $u_i = v_i - p$, if the trader is a buyer
- Trader's utility does not rely on his bidding price if the shout is transacted. However, bidding price determines whether a shout can be matched.
- Bidding prices are determined by the trader's bidding strategies.

# Auctioneers revenue

The auctioneer of a double auction market creates his revenue by:

- charging market registration fees
- charging shout fees
- charging transaction fees
- sharing with the traders profit:
  - for ask shout, profit = clearing price - ask price
  - for bid shout, profit = bid price - clearing price

# Design a double auction market

Design of a double auction market is to specify the following market policies:

- accepting policy: Determine if a shout from a trader should be accepted for further processing.
- matching policy: Determine which two shouts are matched for transaction
- pricing policy: Determine the transaction price for the matched shouts
- clearing policy: Determine when to clear the shouts.
- charging policy: Determine how to charge traders for market services.

# Design a matching policy

Matching can be designed in different ways depending on design criteria:

- Equilibrium matching: maximise auctioneer's profit.
- Maximal matching: maximize liquidity.

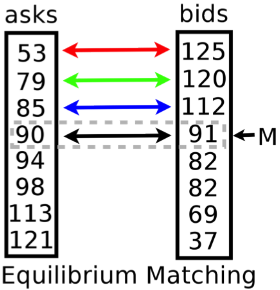


$market\ profit = 141$

$market\ liquidity = 4$

$market\ profit = 113$

$market\ liquidity = 6$

# Design a trader

- Trader's trading strategies:
  - Biding strategy: determine which price to bid.
  - Market selection strategy: determine which market to go.
- Typical bidding strategies:

ZI: Zero Intelligence [Gode and Sunder, 1993]
ZIP: Zero Intelligence Plus [Cliff and Bruten, 1997]
GD: Gjerstad and Dickhaut [Gjerstad and Dickhaut, 2001]
RE: Roth and Erev [Erev and Roth, 1998]

# Outline

# Combinatorial auctions

## Shoe for Sale

Nude High Cut Shoe

- 10cm Natural Leather Stacked Heel.
- Multiple Strap Detail With Gold Western Buckles & Functional Zip.
- Leather Lining & Sock
- Unique

The original price of a pair was $500.
Now the last piece for $5 only.

# Combinatorial auctions

- A combinatorial auction sells multiple objects semitanuously.
- Bidders can place bids on combinations of items in bundles.
- Bidder's valuation on a bundle may be different from the sum of the valuations of all the items in the bundle.
  - Complementary: the value of a combination of items is worth more than the sum of the values of the separate items.
  - Substitutable: the value of a combination of items is less than the sum of the values of the separate items.

# The model of combinatorial auctions

$E = (N \cup \{0\}, X, \{v_i\}_{i \in N})$ is a combinatorial auction if

- $N = \{1, 2, \cdots, n\}$ is the set of buyers
- 0 represents the seller
- $X$ is the set of items
- $v_i : 2^X \to \mathbb{Z}^+$ the buyer i's value function

**Example:**

$N = \{1, 2\}$, $X = \{a, b\}$.
$v_1(\emptyset) = 0$, $v_1(\{a\}) = v_1(\{b\}) = v_1(\{a, b\}) = 1$.
$v_2(\emptyset) = 0$, $v_2(\{a\}) = v_2(\{b\}) = 1$, $v_2(\{a, b\}) = 3$.

**Question:** How to allocate the items to the buyers so that each item goes to the buyer who gives it the highest value?

# Efficient allocations

- Allocation: $\pi : N \cup \{0\} \to 2^X$ such that
  - $\pi(i) \cap \pi(j) = \emptyset$ for any $i \neq j$.
  - $\bigcup_{i \in N \cup \{0\}} \pi(i) = X$.

  which allocate all the items to the buyers, each buyer can have a bundle but one item can only be allocated to at most one buyer.
- Efficient allocation $\pi^*$: $\pi^*(0) = \emptyset$ and for every allocation $\pi$ of $X$,

$$\sum_{i \in N} v_i(\pi^*(i)) \geq \sum_{i \in N} v_i(\pi(i))$$

**Question:** How to find an efficient allocation?

# Walrasian equilibria

- Price vector $\mathbf{p}$: assign a non-negative real number to each item in $X$.
- Demand correspondence:

$$D_i(\mathbf{p}) = \arg\max_{A \subseteq X} (V_i(A) - \sum_{a \in A} p_a)$$

representing all the bundles that give i the highest utility based on the current market price. For instance, if $\mathbf{p} = (0.5, 0.5)$,

$$D_1(\mathbf{p}) = \{\{a\}, \{b\}\}. \; D_2(\mathbf{p}) = \{\{a, b\}\}$$

- Walrasian equilibrium $(\mathbf{p}, \pi)$: $\mathbf{p}$ is a price vector and $\pi$ is an allocation of $X$ such that $\pi(0) = \emptyset$ and $\pi(i) \in D_i(\mathbf{p})$ for all $i \in N$.
- Any Walrasian equilibrium determines an efficient allocation.

**Question:** How to find a Walrasian equilibrium?

# Dynamic auctions

A dynamic auction procedure:

1. Initially set the price vector $\mathbf{p}$ to a starting price vector $\mathbf{p}^0$.
2. Ask each buyer i to report her demand correspondence $D_i(\mathbf{p})$.
3. Test if there is excess demand. If no, stop. Otherwise, raise the price of each item with excess demand by one and go to step (2).

**Question:** If the procedure guarantees to converge to an Walrasian equilibrium?

# Some results

- Walrasian equilibria do not always exist [Gul and Stacchetti 1999].
- Walrasian equilibria exist if each buyer's valuation function satisfies the following GS condition [Gul and Stacchetti 2000]:

  Gross substitutes condition: *if the price of some items was increased, the demand for the items which price has not been increased remains the same.*
- This condition implies substitutability.

**Question:** If each bidder's valuation satisfies GS condition, whether does a dynamic auction procedure converges to an Walrasian equilibrium? **Yes**

# Gross substitutes and complements

- Gross substitutes and complements (GSC) condition: *if all the items can be divided into two categories, say software and hardware, increasing the prices of items in one category and decreasing the prices of items in the other category would not affect the demand of the items that prices are not changed.* [Sun and Yang 2006].

# Double-track auction

1. Let $X = X_1 \cup X_2$ and $X_1 \cap X_2 = \emptyset$.
2. Set initial price of each item in $X_1$ to be zero and the initial price of each item in $X_2$ to be an artificial high price.
3. Ask each buyer i to submit her demand correspondence $D_i(\mathbf{p})$.
4. If there is no excess demand for items in $X_1$ and there are no residual items in $X_2$, stop. Otherwise, increase the price of each item in $X_1$ with excess demand by 1 and decrease the price of each item in $X_2$ with no demand by 1. Go to step (2).

**Theorem**

If each bidder's valuation satisfies GSC with the items of two categories $X_1$ and $X_2$, the above algorithm converges to a Walrasian equilibrium [Sun and Yang 2009]. The complexity of the algorithm is polynomial [Zhang *et al.* 2010]

# Outline

# Early Internet Advertising

- Advertisers place ads (banners) on a website.
- The owner of the website charges the advertisers on a per-impression basis, typically a flat fee to show their ads a fixed number of times (one thousand impressions).
- Major problems:
  1. Contracts were negotiated on a case-by-case basis, mostly are quite small (few thousand dollars per month).
  2. The capacity of a website is limited. A banner ad is shown to every one visiting the website, no matter whether it is relevant or not. Obviously the web resource is largely wasted.

# Sponsored search and first-price ad auction

- An advertiser chooses a set of keywords that are related to the product it wishes to sell.
- Each advertiser submits a bid for each keyword, i.e., the amount to be paid per click.
- When a user's search query matches a keyword, a set of ads is displayed. These ads are ranked by bids. The ad with the highest bid receives the best position; i.e., the position that is mostly likely to be clicked on by the user.
- Advertisers pay by clicks. Every time a consumer clicked on a sponsored link, the associated advertisers account is automatically billed the amount of the advertisers most recent bid.
- This approach was introduced by Overture Services (now part of Yahoo) in 1997.

# Weakness of first-price ad auctions

Suppose there are two slots on a page and three bidders. An ad in the first slot receives 200 clicks per hour while the second slot gets 100. The bidders valuations on each click are

| | Bidder 1 | Bidder 2 | Bidder 3 |
|---|---|---|---|
| valuation | $10 | $4 | $2 |

After a period of time when bidder 2 knew that bidder 3's bid was $2, he reduced his bid to $2.01 to guarantee that he got a slot. Then bidder 1 would not want to bid more than $2.02. But then bidder 2 would want to revise his bid to $2.03 to get the top spot, bidder 1 would in turn raise his bid to $2.04, and so on.

# Weakness of first-price ad auctions

- There is no pure strategic equilibrium in the one-shot auction therefore the game is unstable.
- The bidder who could react to its competitors moves fastest had a substantial advantage. The mechanism therefore encouraged inefficient investments in gaming the system.
- It also created volatile prices that in turn caused allocative inefficiencies.

# Ad auction: second-price

- Advertisers are invited to submit bids.
- Advertisers' links are arranged on the page in descending order of their bids.
- The advertiser in the first position pays a price per click that equals to the bid of the second advertiser plus an increment (say 1 cent); the second advertiser pays the price offered by the third advertiser plus an increment and so forth
- This auction mechanism is known as Ad auction, name by Google, or Generalized Second-Price Auction (GSP auction) by economists.
- Ad auction was introduced by Google in 2002. In 2004, Over 98% of Googles total revenue was generated from Ad auctions ($3.189 billion).

# Ad auctions: second-price

| | Bidder 1 | Bidder 2 | Bidder 3 |
|---|---|---|---|
| value | \$10 | \$4 | \$2 |

If all advertisers bid truthfully, then bids are \$10, \$4, \$2. Under second-price ad auction, payments will be \$4.01 and \$2.01. Note that total payments of the bidders are around \$800 and \$200, respectively while the first-price auction gives about \$400 and \$200 at its worst case, respectively. Google says their "unique auction model uses Nobel Prize-winning economic theory to eliminate … that feeling that you've paid too much".

# The model of GSP auctions

- $K$ bidders: the advertisers $k = 1, 2, \cdots, K$.
- $N$ items: the positions on the screen, where ads related to a keyword can be displayed, $i = 1, 2, \cdots, N$.
- Click Through Rates (CTR) $\alpha_i$: the expected clicked if an ad placed in position $i$ during a period. Without loss of generality, we assume that $\alpha_1 > \alpha_2 > \alpha_3 > \cdots > \alpha_N$.
- Value per click $s_k$: the value to bidder $k$ if it receives a click.
- Payoff $u_k(i)$: bidder $k$'s payoff from an ad placed at position $i$.

$$u_k(i) = \alpha_i s_k$$

# Properties of GSP auctions

- Truth-telling is not a dominant strategy under GSP
- This means that a bidder does not have to bid using its true value to maximize its profit.
- For a market mechanism, if it is not a truth-telling mechanism, bidders would be able to manipulate the mechanism (play tricks) to get better outcomes.

# Why GSP is not a truth-telling mechanism?

**Lower your position:**

Assume that three bidders bid for two ad slots with private values \$10, \$4 and \$2, respectively. The two slots receive a similar clicks per hour, say 200 and 199, respectively. The first bidder bids with its true value receiving a payoff: $(\$10 - \$4) * 200 = \$1200$. If it bids \$3, it receives a payoff: $(\$10 - \$2) * 199 = \$1592 > \$1200$.

# Why GSP is stable?

**Envy Free Equilibrium:**

- Suppose that you won a slot at position $l$ and your competitive opponent won the slot immediate above yours and pays your bidding price. To give your opponent a lesson, you raise your bid a bit, which does not affect your payoff but will reduce the payoff of your opponent.
- However, your opponent can easily retaliate you with the following strategy. Because she knows your truth valuation, she can simply reduce her bid to your valuation, which will force you move up one position (therefore you will have to pay exactly your valuation plus an increment). This means that you will receive a negative utility.
- Such a property is called envy free. A GSP procedure guarantees to reach an envy free equilibrium.

In the reality, Ad auction has been combined with many other factors, such as link quality, to make it work more smoothly.

# Recommended readings


Paul Klemperer.
Auction Theory: A Guide to the Literature.
*Journal of Economic Surveys*, 13(3):227–286, 1999.

Ken Hendricks and Robert H. Porter.
An empirical perspective on auctions.
In Mark Armstrong and Robert Porter, editors, *Handbook of Industrial Organization*, volume 3, chapter 32, pages 2073–2143. Elsevier, 1 edition, 2007.

Paul Milgrom.
Putting auction theory to work : Ascending auctions with package bidding.
*Journal of Political Economy*, 108(2):245–272, 2000.

# Recommended readings


- Peter Cramton, Yoav Shoham, and Richard Steinberg.
  An overview of combinatorial auctions.
  *SIGecom Exch.*, 7:3–14, December 2007.

- Benjamin Edelman, Michael Ostrovsky, and Michael Schwarz.
  Internet advertising and the generalized second-price auction: selling billions of dollars worth of keywords.
  *American Economic Review*, 97(1):242–259, March 2007.

- Ning Sun and Zaifu Yang.
  A double-track adjustment process for discrete markets with substitutes and complements.
  *Econometrica*, 77(3):933–952, 2009.

- Michael P. Wellman, Amy Greenwald, and Peter Stone.
  *Autonomous Bidding Agents: Strategies and Lessons from the Trading Agent Competition*.
  MIT Press, 2007.

# Recommended readings

Dongmo Zhang, Laurent Perrussel, and Wei Huang.
Dynamic auction: A tractable auction procedure.
In *Proceedings of the Twenty-Fourth AAAI Conference on Artificial Intelligence (AAAI 10)*, pages 935–940, 2010.

Michael Thielscher and Dongmo Zhang.
A formal market specification language for general trading agents.
In E. David et al., editor, *Agent-Mediated Electronic Commerce: Designing Trading Strategies and Mechanisms for Electronic Markets,*, LNBIP 59, pages 259–274. Springer, 2010.

Dengji Zhao, Dongmo Zhang, and Laurent Perrussel.
Mechanism design for double auctions with temporal constraints,
In *Proceedings of the 22nd International Joint Conference on Artificial Intelligence (IJCAI-11)*, 472-477,2011.

# The End

Thank you for your attention!