Introduction to Game Theory

Isabel Amigo

IMT Atlantique

October 2017

Notes:

Exercise: More Information might hurt

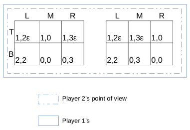

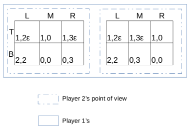

- Consider these games where ϵ verifies $0 < \epsilon < 1/2$
- Compute the BNE expected payoff for player 2 in both variants of the game

Notes:

- In single-person decision problems, a person cannot be worse off with more information. In strategic interactions, a player may be worse off if she has more information and other players know that she has more information.
- This example is taken from An Introduction to Game Theory, by Martin Osborne

1 Introduction
Examples
Basic assumptions
Course organization
2 Strategic Games
Definition and First solution concept
Solution concept: NE in Mixed Strategies
3 Games in Extensive form with perfect information
4 Pricing
Discussion: Pricing in the Internet
5 Games with incomplete Information
Bayesian Games
Auctions
Case Study: Adwords Auctions
6 Cooperative Game Theory
7 Conclusion
8 Acknowledgments

Notes:

Notes:

- Photo: Flora Holland Flower Auction by Scott Ableman, taken on june 26th 2017 `https://flic.kr/p/VNZEew`
- Auctions have been used for ages for selling public goods, art, etc.
- As a pricing mechanism: situations where the market price for a good is not known in advance
- As a resource allocation mechanism: how to allocate a resource when there is more demand than offer
- So what is an auction?
- Some definitions describe the as a selling institution where the winner and price are only based on bids
- This implies they are *universal* (can be used to sell any good) and *anonymous* (winner and price do not depend on *who* placed the bid but on the amount of the bids)
- A larger selling institution are Mechanisms
  - Mechanisms differ from auctions in that they are not necessarily universal or anonymous

Different flavors of auctions

Among the most classical types auctions we have

- Open
  - English auction
  - Dutch auction
- Sealed-bid
  - First-price auction
  - Second-price auction

Notes:

- The Dutch open descending price auction is strategically equivalent to the first-price sealed-bid auction
- When values are private, the English open ascending auction is also equivalent to the second-price sealed-bid auction, but in a weaker sense than noted earlier

Different flavors of auctions

But we also have

- Ad-words
- Double-auctions
- “deadline” auctions
- Other mechanisms ..

Notes:

Sealed-bid auctions

1. Each bidder i privately communicates a bid b_i to the auctioneer
2. The auctioneer decides who gets the good (if anyone)
3. The auctioneer decides on a selling price

Notes:

- There is a natural way to implement step 2: give the object to the bider placing the highest bid. Today, we will focus on this allocation rule that we study.
- How to implement step 3 is less intuitive. This step has a huge impact on bidder behavior. For example, imagine we simply decide to charge nothing for the winning bidder. Then bidders will have incentive to increase indefinitely their bids.

Questions/Objectives

- From the bidder point of view
  - How to set the bid?
  - How to use available information?
  - How to win without paying *too much*
- From the seller, or auction designer point of view
  - Revenue maximization
  - Social optimum
  - Incentive compatibility
  - individual Rationality
  - Complexity, distributed implementation
  - Collusion avoidance

Notes:

- Revenue maximization mechanisms are called *optimal* mechanism
- Mechanisms maximizing the aggregated utilities are called *efficient* mechanisms
- A mechanism is incentive compatible if bidding *truthfully* is a dominant strategy
- An individually rational mechanisms guarantees that the expected payoff for any player is greater or equal to zero

Some Vocabulary

- Valuation: willingness to pay of a bidder for the object on sale
- Bid: offer done by a buyer
- Price: the price determined for the object, after auction takes place
- Utility: quasi-linear model usually adopted
- Allocation rule (step 2 two slides ago)
- Payment rule (step 3 two slides ago)

Notes:

- the utility model we are assuming is the so-called quasi linear one. there are other models that could also be pertinent
- under this model the utility for player i is the difference between i's valuation and the paid price , if win, or zero otherwise
- We shall adopt the following notations: valuation for player i is v_i, bid for player i is b_i utility for player i is $u_i = \mathbb{1}_{\text{win}}(v_i - b_i)$

Some assumptions to start with

- Single-object
- Private and independent values
- Common prior
- Bidders seek to maximize their expected profits

Notes:

- Private values: each player's valuation for the object is a private information for him/her
- Common prior: all private values are drawn from a same known distribution, number of bidders are known by all
- Independent values: valuations are independently drown from the common distribution
- Each players valuation is a random variable (RV) noted V_i. Each V_i is independently and identically distributed on some interval $[0, \omega]$ according to the increasing distribution function F.
- Models for multiple-objects, interdependent values also exist

Preferences towards risk - Definitions

Risk Neutral
Given the choice between a guaranteed payoff of R and a gamble with expected payoff also equal to R , the bidder is completely indifferent.

Risk Averse
In the previous scenario, prefers the former

Notes:

- For instance, a risk neutral person is indifferent between receiving 100 for sure and 0 with probability 9/10 and 1000 with probability 1/10.
- A risk averse person will prefer the 100 for sure
- We will usually assume risk neutral bidders, but not always

Second-Price (or Vickrey) auctions [6]

- Bidders submit a sealed bid
- Allocation: the bidder with the highest bid is awarded the object
- Payment: the winner pays the second-highest bid

Theorem
Bidding truthfully is a weakly dominant strategy

Notes:

- utility for player i, $u_i = v_i - max_{j \neq i}$ if $b_i > max_{j \neq i}$ (i.e. i wins); 0 otherwise.
- proof intuition:
  - Bidders have incentives to submit their true valuation (incentive-compatible mechanism):
  - over-bidding creates the risk of paying more than one's valuation
  - under-bidding leads to the risk of losing the auction and getting utility 0 in some cases when bidding truthfully would have led to a strictly positive utility.
- The proof can be readily formalized
- Such proof does not relay in two of the assumptions we have previously done (valuations independently and independently distributed). Only the assumption of private values is important, so the Theorem holds as long as we have private values.

Second-Price (or Vickrey) auctions, some properties

- It is incentive-compatible
- It is individually rational
- If bidders bid truthfully, then it maximizes the social surplus
- Can be implemented in linear time

Notes:

- Second-price auctions fulfill a series of our previously mentioned desired properties
- the two first properties imply truthful reporting is a dominant strategy and never leads to negative utility, so it is easy for a bidder to chose a bid
- we define social surplus as $\sum_i v_i \mathbb{1}_{iwins}$

Second-Price (or Vickrey) auctions

- How much a bidder expects to pay in equilibrium?

Expected Payment
For player i bidding b_i is the probability that i wins times the expected value of the 2nd highest bid, given that i wins with b_i

Notes:

- Let Y_1 be the RV defined as the maximum among N-1 randomly and independently distributed values $V_i, i \neq 1$ Y_1 is known as the *highest-order statistic* of $V_2, V_3, \ldots, V_N$.
- Let G denote the probability distribution of Y_1 it can be easily deduced that $G(y) = F(y)^{N-1}$ (remember F is the dist. of V_is)
- We can then compute the expected utility payment of bidder i under second-price auction as: which yields to $p_i^{2nd}(b_i)$=Prob[i wins]$\times$ E[2nd highest bid $|b_i$ is the highest bid] which is equivalent to $p_i^{2nd}(b_i) = Prob[Y_1 < b_i] \times E[max_{j \neq i} b_i | Y_1 < b_i]$.
- Finally, using the fact that in equilibrium $v_i = b_i$ we obtain $p_i^{2nd}(b_i) = G(v_i) \times E[Y_1 | Y_1 < v_i]$

First price auction

- Bidders submit a sealed bid
- Allocation: the bidder with the highest bid is awarded the object
- Payment: the winner pays the highest bid
- Bidding strategy is less intuitively deduced than in second-price auctions
- Bidders have incentives to submit a bid lower than their valuation

Notes:

- Utility for player i, $u_i = v_i - b_i$ if $b_i > max_{j \neq i}$ (i.e. i wins); 0 otherwise.
- Bidding strategy obtained through the equilibrium of the game (finding the best response of each player)
- Best response $s_i(v_i) = E[Y_1 < v_i]$
- We can think it as if each bidder considers for bidding all the cases in which his/her valuation is the highest, and in that case computes the expectation of the highest of the other players' valuations. This expectation is the amount she/he will bid.
- In the particular case of N bidders with valuations are randomly and independently distributed according to a uniform distribution between 0 and 1. $s_i(v_i) = \frac{N-1}{N} v_i$

First-price auctions: expected payment

Expected payment
For i is Prob[i wins] $\times$ Amount bid

Notes:

- $p_i^{1st}(v_i) = G(v_i) E[Y_1 | Y_1 < v_i]$
- Note that is the same as in second-price auctions

First price auction - some properties

- Are not incentive compatible
- Are efficient

Notes:

Revenue-equivalence theorem

Theorem

Consider two auction mechanisms such that

- bidders are risk-neutral;
- bidder valuations are independently distributed over a given interval, with a finite and strictly positive density;
- bidder with the lowest possible valuation expects a null utility;
- bidder with the highest valuation always wins the item.

Then both schemes yield the same expected revenue to the seller at equilibrium, and each bidder gets the same utility.

Notes:

- In particular, first and second price auctions are equivalent from the point of view of the revenue obtained by the auctioneer
- For instance, a third-price auction, verifying the hypothesis of the theorem yields also the same revenue

Vickrey-Clarke-Groves (VCG) auction

Setup

- Multi-unit auction
- Consider a set X of all possible outcomes (
- Bids: each bidder submits a bid for each possible outcome
- Allocation rule: such that social welfare is maximized (according to the declared bids)
- Payment rule: opportunity cost (the total loss of (declared) value his/her presence imposes on the others)

Notes:

- Each bid submits a bid $b_i(x)$ for each $x \in X$
- allocation rule : $x^{VCG} = \text{argmax}_x \sum_{i \in N} b_i(x)$
- payment rule: $p_i^{VCG} = \max_{x \in X} \sum_{j \neq i} b_j(w) - \sum_{j \neq i} b_j(x^{VCG})$

VCG properties

- Bidding truthfully is a dominant strategy

Theorem

Only mechanism to be jointly incentive compatible, individually rational and efficient.

Notes:

- Note that for a single object, VCG coincides with a second-price auction
- Can you think on inconveniences of this auction?
- Proof intuition of truthfulness: consider player i and fix b_i
  - $u_i = v_i(x^{VCG}) - p_i^{VCG} =$
    $v_i(x^{VCG}) + \sum_{j \neq i} b_j(x^{VCG}) - \max_x \sum_{j \neq i} \hat{V}_j(x_{-i})$

Beyond auctions: Mechanism design - The Science of Rule-Making

- Mechanism: set of rules chosen by the decision maker to optimize a global objective
- It includes
  - A set of available strategies for each agent,
  - An outcome rule, that maps the strategy profiles of agents to an outcome (ex: allocation of resource).
- The rules have to be designed so that a game played by selfish agents "naturally" reaches the expected outcome.

Notes:

- Is a Stackelberg situation, where the leader is the designer of the game played by the followers.
- Bayesian games: all agents reason on a common prior distribution of the types of the others.

1. Introduction
   - Examples
   - Basic assumptions
   - Course organization
2. Strategic Games
   - Definition and First solution concept
   - Solution concept: NE in Mixed Strategies
3. Games in Extensive form with perfect information
4. Pricing
   - Discussion: Pricing in the Internet
5. Games with incomplete Information
   - Bayesian Games
   - Auctions
   - Case Study: Adwords Auctions
6. Cooperative Game Theory
7. Conclusion
8. Acknowledgments

Notes:

Case Study: adword auctions

- Search engines play a crucial role in the Internet.
- They get revenue mainly through advertising slots, usually displayed at the top or right of the search page.
- Advertisers submit bids for relevant keywords only.

Notes:

- Content in this subsection comes from the work of B. Tuffin (INRIA, France) and P. Maillé (IMT Atlantique, France)

Auction principle (single keyword, K slots)

- Advertisers submit bids for specific keywords.
- Each time there is a search on that keyword:
- advertisers are ranked and allocated slots according to a prespecified criterion:
  - bid value (initially for Yahoo!)
  - the revenue they will generate (more or less Google), taking into account the (learned) click-through rate (CTR).

Notes:

Auction principle (single keyword, K slots)

Notes:

- Possible payment rules:
  - *Pay-Per-Impression* (PPI): advertisers charged every time their ad is displayed
  - *Pay-Per-Click* (PPC): advertisers is charged only when the ad is clicked
  - *Pay-Per-Transaction* (PPT): advertisers charged when the click results in a real sell.
- Amount to be paid each time?
  - First Price
  - Generalized Second Price (GSP): they pay the bid of advertiser below them in the ranking
  - VCG auctions
- In use: PPC and GSP.

More details on auctions for ads

Notes:

- N advertisers, $k(< N)$ advertisement slots
- v_i: valuation of Advertiser i for the "considered action" (impression, click, sale): maximum price i is willing to pay
- b_i: bid submitted by i (not necessarily equal to v_i)
- $\mathbf{b} = (b_1, \dots, b_{|N|})$ bid profile
- But the k slots do not have the same probability to be looked at
  Ads at the top more seen than those at the bottom.
- Different interest for N advertisers depending on the search too.

Click-Through-Rate (CTR)

Notes:

Definition (Click-Through-Rate)

Probability that a given ad will be clicked when displayed.

CTR $w_{i,s}$ for advertiser i at slot s.

CTR often assumed separable:

$$w_{i,s} = q_i r_s$$

- q_i: attractiveness of Advertiser i
- r_s: probability that a user considers the ad on slot s. (Slots ordered $r_1 \geq \dots \geq r_{|N|}$.)

Ranking slot allocations

Notes:

- Rank according to bids
- Rank according to $w_{i,s} b_i$. More exactly
  1. Ranked first: the advertiser maximizing $w_{i,1} b_i$
  2. Ranked second: advertiser maximizing $w_{i,2} b_i$ (excluding the first)
  3. ...
- Generalizes ranking per bid: just consider $w_{i,s} = 1 \; \forall i, s$
- If the charge is p_s at the s-th slot,revenue generated with a pay-per-click scheme: $\sum_{s=1}^{k} w_{(s)s} p_s$.

Charging rule: GSP

- First price auction could be considered
- ... or VCG, but
- In practice GSP: Generalized Second-Price

Generalized Second-Price
you pay not what you have declared, but a price equivalent to the minimum bid to maintain your position in the ranking

Notes:

GSP

Explicitly:

- if ranking by bid, the winner of slot $s \leq k$ is charged $b_{(s+1)}$, because bidding less would mean losing the s-th slot.
- If ranking by revenue,
  - revenue associated to slot s: $w_{(s)s}p_s$
  - under the separability assumption, price p_s charged such that bidding less than p_s would make you lose the slot : $q_{(s)}r_s p_s \geq q_{(s+1)}r_s b_{(s+1)}$. This gives

$$p_s = b_{(s+1)}\frac{q_{(s+1)}}{q_{(s)}}.$$

  - Intuition: some advertisers with low CTR q would generate a low revenue even if their bids are high.

Notes:

Example: $k = 3$ slots, $n = 5$ advertisers with $r_1 = 1/2$, $r_2 = 1/4$ and $r_3 = 1/5$

Advertiser i	Bid b_i	CTR q_i	Product $b_i q_i$
1	10	0.05	0.5
2	9	0.1	0.9
3	6	0.12	0.72
4	5	0.15	0.75
5	4	0.2	0.8

Notes:

- Ranking per bid:
  - The three slots are allocated to first three advertisers
  - $p_1 = b_2 = 9$, $p_2 = b_3 = 6$, $p_3 = b_4 = 5$
  - Expected revenue $\sum_{s=1}^{3} r_s q_{(s)} p_s = \sum_{s=1}^{3} r_s q_{(s)} b_{(s+1)} = \frac{1}{2}0.45 + \frac{1}{4}0.6 + \frac{1}{5}0.6 = 0.52$.
- Ranking per revenue:
  - Advertiser 2 is allocated the first slot, Advertiser 5 is allocated the second, and Advertiser 4 the third
  - $p_1 = 8$, $p_2 = 3.75$ and $p_3 = 4.8$ (with $p_s = b_{(s+1)}\frac{q_{(s+1)}}{q_{(s)}}$)
  - Revenue $\sum_{s=1}^{3} r_s q_{(s)} p_s = \sum_{s=1}^{3} r_s q_{(s+1)} b_{(s+1)} = \frac{1}{2}0.8 + \frac{1}{4}0.75 + \frac{1}{5}0.72 = 0.7315$.

GSP and VCG

Proposition
In the case of a single slot, VCG and GSP are equivalent.

VCG procedure:

- Maximizes the the declared valuation (bid) of the winner for the bid-based ranking
  - hence selects the largest bidder like GSP.
  - Price paid: loss of declared value, second highest bid;
- Maximizes the (declared) generated revenue for the revenue-based ranking
  - hence advertiser maximizing $q_i b_i$ like GSP.
  - *Total charge*: loss of declared revenue of other players, value $r_1 q_{(2)} b_{(2)}$. Idem.

Notes:

But not true for more than one slot

Back to our previous example, focusing on revenue-based ranking:

- Allocations for VCG are the same
- Payments:
  - For Advertiser 2, winner of the first slot: loss of (declared) revenue due to his presence:

$$b_3 q_3(r_3 - 0) + b_4 q_4(r_2 - r_3) + b_5 q_5(r_1 - r_2)$$
$$= \frac{0.72}{5} + 0.75\left(\frac{1}{4} - \frac{1}{5}\right) + 0.8\left(\frac{1}{2} - \frac{1}{4}\right) = 0.3815,$$

  - For Advertiser 5,winner of the second slot:

$$b_3 q_3(r_3 - 0) + b_4 q_4(r_2 - r_3) = 0.72\frac{1}{5} + 0.75\left(\frac{1}{4} - \frac{1}{5}\right) = 0.1815;$$

  - For Advertiser 4,winner of the third slot:

$$b_3 q_3(r_3 - 0) = 0.72\frac{1}{5} = 0.144.$$

- Expected revenue 0.707, less than the 0.7315 when using GSP.

Notes:

Why GSP instead of VCG?

- GSP does not satisfy the incentive compatibility property in general (exercise)
  - VCG prices unique truthful prices
  - But at least verifies properties such as "every bidder allocated position s has no incentive to switch to positions $s-1$ or $s+1$ through a bid change";
- GSP more "complicated" in terms of strategy and resulting equilibrium
- And payment rule simpler to understand.

Notes:

- Comparison of expected revenue GSP vs VCG .
  Let $p_j^{(GSP)}$ and $p_j^{(VCG)}$ charges per click of GSP and VCG for slot j.
  Our induction assumption is $p_{s+1}^{(GSP)} \geq p_{s+1}^{(VCG)}$ (equal for the last slot).
  Then with VCG is the difference of opportunity costs between s and $s+1$:

$$\begin{aligned} r_s q_{(s)} p_s^{(VCG)} - r_{s+1} q_{(s+1)} p_{s+1}^{(VCG)} &= b_{(s+1)} q_{(s+1)}(r_s - r_{s+1}) \\ &\leq b_{(s+1)} q_{(s+1)} r_s - b_{(s+2)} q_{(s+2)} r_{s+1} \\ &= r_s q_{(s)} p_s^{(GSP)} - r_{s+1} q_{(s+1)} p_{s+1}^{(GSP)}. \end{aligned}$$

  Therefore
  $r_s q_{(s)} p_s^{(GSP)} \geq r_s q_{(s)} p_s^{(VCG)} + r_{s+1} q_{(s+1)}(p_{s+1}^{(GSP)} - p_{s+1}^{(VCG)}) \geq r_s q_{(s)} p_s^{(VCG)}$
- The analysis we have made is based on pay-per-click
- There also exist on pay-per-view
- Some advertisers are more interested in *brand awareness* –not related to clicks–. Ex: Coca-Cola; no direct sale from clicks

Learning

- CTR have to be learned
- The advertiser has to trust the publisher: some advertisers filed lawsuits claiming to be victims of overcharging by lying (increasing) the real CTR
- Statistical tools to estimate the CTR.

Notes:

Quizz Lecture 4

Exercise 1 (Second-price Auctions) *Probe that in second-price auctions, with private values and a single object, bidding truthfully is a dominant strategy.*

Exercise 2 (VCG for one object) *Show that for the case of a single object, VCG coincides with a second-price auction.*

Exercise 3 *In this problem we will ask how the number of bidders in a second-price, sealed-bid auction affects how much the seller can expect to receive for his object. Assume that there are two bidders who have independent, private values v_i which are either 1 or 3.*

For each bidder, the probabilities of 1 and 3 are both $1/2$. (If there is a tie at a bid of x for the highest bid the winner is selected at random from among the highest bidders and the price is x.)

1. *Show that the seller's expected revenue is $3/2$.*
2. *Now let's suppose that there are three bidders who have independent, private values v_i which are either 1 or 3. For each bidder, the probabilities of 1 and 3 are both $1/2$. What is the seller?s expected revenue in this case?*
3. *Briefly explain why changing the number of bidders affects the seller?s expected revenue.*

Exercise 4 (Adwords auction) *Consider the example considered in the class slides with a small variation. Here, $k = 3$ slots, where $r_1 = 1/2$,$r_2 = 1/4$ and $r_3 = 1/5$, and $n = 5$ advertisers with bids and CTRs given in the table. The modification is: q_2 is increased to 1.*

Advertiser	*Bid* b_i	*CTR* q_i
1	*10*	*0.05*
2	*9*	*1*
3	*6*	*0.12*
4	*5*	*0.15*
5	*4*	*0.2*

1. *Determine the winners of the slots for the ranking based on bids. Compute the paid prices and expected revenue.*

2. *Determine the winners of the slots for the ranking based on revenue. Compute the paid prices and expected revenue.*

3. *Show that ranking per revenue does not always produce the largest revenue*