Algorithmic Game Theory

Master 2 MMMEF, Université Paris 1 Panthéon-Sorbonne

Antoine Mandel, antoine.mandel@univ-paris1.fr

December 31, 2025

Topics

- Auction and online advertising (1 week)
- Optimal auctions (1 week)
- Mechanism Design with quasi-linear utilty, VCG,(1 week)
- Position auctions (1 week)
- The emergence of RTB and pacing equilibrium (1 week)
- Mechanism design without money: median voter and matching (2 weeks)
- Matching and Pricing in Ride Hailing (1 week)
- Exam

Outline

The online advertising business

- The online advertising business is worth approx 500 billion US dollars in 2025.
- Three main tiers:
  1. Sponsored search (buying impressions/display following search of certain keywords)
  2. Real-time bidding (buying impressions as they are created)
  3. Private contracts (buying impressions in bulk through private agreements)
- Market for 1. and 2. is mostly based on automated auctions and will be our core concern. Market for 3. more like “standard" advertising.

How Sponsored Search Works (Overview)

- **User query:** A user submits a search query (e.g., "running shoes").
- **Advertisers:** Advertisers submit bids for keywords or queries, expressing a *maximum willingness to pay per click*.
- **Auction:** For each query, the platform runs a real-time auction among eligible ads.
- **Ranking:** Ads are ranked by a *score* combining bid and predicted relevance (e.g., expected click-through rate).
- **Allocation:** Top-ranked ads are assigned to limited ad slots with different visibility.
- **Pricing:** Advertisers typically pay *per click*, often determined by a generalized second-price (GSP) rule.
- **Objectives:**
  - Advertisers: maximize value from clicks under budget constraints
  - Platform: maximize revenue, user satisfaction, or long-term engagement
- **Remark:** Ads are personalized: check it.

Players in the Display Advertising / RTB Ecosystem

- **Users:** End users who consume online content and generate ad impressions when visiting webpages or mobile applications.
- **Publishers:** Content providers who own advertising inventory (impressions) and seek to monetize user traffic.
- **Advertisers:** Firms that demand impressions in order to promote products or services to targeted users.

Players in the Display Advertising / RTB Ecosystem

- **Ad Networks and Ad Exchanges (ADX):**
  - Aggregate demand from multiple advertisers and supply from publishers
  - Run real-time auctions to allocate impressions
  - Enable competition across multiple ad networks
- **Supply-Side Platforms (SSP):**
  - Serve publishers by managing and registering their inventories
  - Collect bids from multiple ad networks or exchanges
  - Automatically select winning ads and optimize publisher revenue
- **Demand-Side Platforms (DSP):**
  - Serve advertisers or ad agencies
  - Automatically bid for impressions across multiple ad networks
  - Optimize bids under budget and performance constraints
- **Data Exchanges / Data Management Platforms (DX / DMP):**
  - Provide historical and contextual user data in (near) real time
  - Serve DSPs, SSPs, and ADXs to improve targeting and matching

How RTB Works with Behavioral Targeting

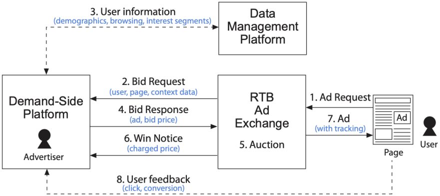

Figure 2.2: How RTB works for behavioural targeting. Source: [Zhang et al., 2014b].

- **0. Impression creation:** A user visits a webpage; an ad impression is created while the page loads.
- **1. Ad request:** The publisher sends an ad request to an ad exchange via an ad network or a Supply-Side Platform (SSP).
- **2. Bid solicitation:** The ad exchange forwards the request to multiple Demand-Side Platforms (DSPs).
- **3. Data enrichment:** DSPs may query Data Exchanges / Data Management Platforms (DMPs) for third-party user data (e.g., past behavior, interests).
- **4. Bid generation:** Advertisers evaluate the user and context and submit bids.
- **5. Auction:** The ad exchange runs a real-time auction (historically second-price; now often first-price). If multiple exchanges are used, the SSP selects the final winner.
- **6. Win notification:** The winning advertiser is notified.
- **7. Ad display:** The winning ad creative (text, image, or video) is returned and displayed to the user.
- **8. Feedback and tracking:** User interactions (clicks, conversions) are recorded and used for future optimization.

The whole process takes place within $\sim$ 100 ms.

Auctions

- Auctions are at the core of the online advertising industry and can be analyzed mathematically (and computationally).
- An auction is a mechanism through which a seller seeks to allocate, against payment, an (ensemble of) objects among a set of buyers on the basis of bids issued by the buyers.
- Formally, an auction is characterized by:
    - A set of players N, a set of feasible allocations $\mathcal{X}$, a set of potential actions (bids, sequence of bids) for each player i, B_i, with $B := \prod_{i \in N} B_i$.
    - An allocation rule $a : B \to X$ that determines an allocation as a function of the submitted bids.
    - A payment rule $p : B \to \mathbb{R}^N$ that determines the payment of each agent as a function of the submitted bids.

Standard Auctions: Single Object

- A **single-object auction** allocates one indivisible item to at most one bidder.
- **Sealed-Bid Auctions**: Bidders submit bids simultaneously and privately.
  - **First-Price Sealed-Bid Auction:** (i) Each bidder submits a single bid, (ii) the highest bidder wins the object. (iii) the winner pays her own bid.
  - **Second-Price (Vickrey) Auction:** (i) Each bidder submits a single bid, (ii) the highest bidder wins the object, (iii) the winner pays the second-highest bid.
- **Open Auctions**: Bidding unfolds over time and prices are publicly observable.
  - **English Auction (Ascending):** (i) The auctioneer starts from a low price and gradually increases it, (ii) bidders remain active while the price is below their valuation, (iii) the last remaining bidder wins and pays the final price.
  - **Dutch Auction (Descending):** (i) The auctioneer starts from a high price and gradually decreases it, (ii) the first bidder to accept the price wins the object, (iiI) the winner pays the accepted price.

Standard Auctions: Multiple Objects

- A **multi-object auction** allocates multiple items or positions, possibly with heterogeneous values and externalities.
- **Multi-Unit Auctions (identical items)**: There are k identical units to be allocated.
  - bidders submit (possibly multi-unit) bids;
  - the k highest marginal bids win units;
  - **uniform-price:** winners pay the same per-unit price (typically the highest losing bid);
  - **first-price:** each winner pays her own winning bid(s) per unit;
  - **Vickrey:** if bidder i wins k_i units, she pays the sum of the k_i highest losing marginal bids submitted by the other bidders.
- **Position Auctions (special case of multi-object auctions)**: Objects are ranked slots with different qualities (e.g., click-through rates).
  - **Generalized Second-Price (GSP) Auction:** (i) bidders submit a bid per click, (ii) bidders are ranked by a score (e.g., bid $\times$ quality), (iii) top bidders are assigned to top positions, (iv) each winner pays (per click) the minimum bid required to retain her position.

Auction theory: single-object setting

We start our formal analysis with the case of a single object:

- The set of allocation is identically the set of players N.
- The action set of each player is $[0, w]$ for some $w > 0$.
- The allocation rule $a : B := [0, w]^n \to N$ determines the winning player as a function of bids.
- The payment rule $p : B \to \mathbb{R}^N$ determines the payment of each agent as a function of bids.
- Each agent is characterized by type $\theta_i \in \Theta_i$ (signal, , preference, budget,...).
- Valuation can depend on own type only $v_i(\theta_i)$ (private valuation) or also depend on others $v_i(\theta_i, \theta_{-i})$ (interdependent valuation)
- Quasi-linear utility: valuation (if object received) - payment.

Auction theory: desirable properties/research question

- Incentive Compatibility (IC): players bid truthfully (according to own valuation) independently (dominant) or conditional (Bayesian Nash) on others' actions.
- Individual Rationality (IR): each agent obtains non-negative utility by participating in the auction. The most powerful version
- Efficiency: allocation rule maximizes the social utility, e.g. allocates object to bidder with highest valuation.
- Revenue Maximization: seller aims to maximize revenue from auction.

Auction theory: single-object, private-value

- Single object/private valuation: $v_i = \theta_i$.
- Game-theoretic perspective: strategy of player i is function $\beta_i : \Theta_i \to B_i$ from type/valuation to bid/action. How much to bid given valuation ?
- Payoffs in second-price auction:

$$u_i(b) := \begin{cases} \theta_i - \max_{j \neq i} b_j & \text{if } b_i > \max_{j \neq i} b_j \\ 0 & \text{if } b_i < \max_{j \neq i} b_j \end{cases}$$

- A strategy profile β is a Dominant-Nash equilibrium if for all $i \in N$ and all β_i', one has for all $\theta \in \Theta$

$$u_i(\beta_i(\theta_i), \beta_{-i}(\theta_{-i}), \theta) \geq u_i(\beta_i'(\theta_i), \beta_{-i}(\theta_{-i}), \theta)$$

Proposition

In a second price auction with private valuation and quasi-linear utility, truthful bidding $\beta_i(\theta_i) = \theta_i$ is a dominant strategy.

Auction theory: single-object, private-value

- Payoffs in first-price auction:

$$u_i(b) := \begin{cases} \theta_i - b_i & \text{if } b_i > \max_{j\neq i} b_j \\ 0 & \text{if } b_i < \max_{j\neq i} b_j \end{cases}$$

- A strategy profile β is a Bayesian-Nash equilibrium if for all $i \in N$ and all β_i', one has for all $\theta_i \in \Theta_i$:

$$\mathbb{E}_{\Theta_{-i}}(u_i(\beta_i(\theta_i), \beta_{-i}(\theta_{-i}), \theta)) \geq \mathbb{E}_{\Theta_{-i}}(u_i(\beta_i'(\theta_i), \beta_{-i}(\theta_{-i}), \theta))$$

or equivalently:

$$\int_{\Theta_{-i}} u_i(\beta_i(\theta_i), \beta_{-i}(\theta_{-i}), \theta) dF_{-i}(\theta_{-i}) \geq \int_{\Theta_{-i}} u_i(\beta_i'(\theta_i), \beta_{-i}(\theta_{-i}), \theta) dF_{-i}(\theta_{-i}$$

Proposition

In a first price auction, truth telling isn't a dominant strategy and players have incentives to shade their bids.

Reading list for next week

Myerson, R. B. (1981). Optimal auction design. Mathematics of operations research, 6(1), 58-73.

Outline

Auction theory: revenue-equivalence theorem

Theorem

Suppose that a given pair of Bayesian Nash equilibria of two different auction procedures are such that:

1. *For every bidder i, for each possible realization of θ, bidder i has an identical probability of getting the good in the two auctions.*
2. *Every bidder i has the same expected payoff in the two auctions when his valuation for the object is at its lowest possible level.*

Then the two auctions generate the same expected revenue to the seller.

For the proof, see Myerson 1981 and https://gtl.csa.iisc.ac.in/gametheory/ln/web-md9-ret.pdf

Auction theory: revenue-optimal auction (Myerson, 1981)

- Revenue-optimal auctions are characterized through virtual value functions $\phi_i(\theta_i) = \theta_i - {}^{(1-F_i(\theta_i))}/_{f_i(\theta_i)}$ that substract from the true value the informational rent ${}^{(1-F_i(\theta_i))}/_{f_i(\theta_i)}$
- The thicker the tail, the larger the informational rent.

Theorem (Myerson – Regular Case)

1. *If the ϕ_is are non-decreasing and $\max_{j\in N} \ \phi_j(\theta_j) \geq 0$, the revenue-maximizing auction is such that the object is allocated to $i \in \arg\max_{j\in N} \ \phi_j(\theta_j)$ and its expected revenue is*

$$E_\theta(\max_{j\in N} \ \phi_j(\theta_j))$$

2. *If bidders are symmetric, the optimal auction is a second-price auction with a reserve price r such that $\phi(r) = 0$.*

Elements for the proofs of Myerson (1981) results

- Revelation principle: we can focus on incentive compatible mechanisms.
- Myerson Lemmas on characterization of optimal revenue

Comments on Myerson results

- Optimal auction not necessarily socially efficient (as the seller might retain the object though he has no value for it)
- Auctions with reserve price not revenue-equivalent to standard ones.

Reading list for next week

- Nisan, N., Roughgarden, T., Tardos, E., and Vazirani, V. V. (Eds.). (2007). Algorithmic Game Theory. Cambridge University Press.
- Chapter 9: Introduction to Mechanism design for computer scientists.

Outline

Statement of the problem

- A finite set $N = \{1, \cdots, n\}$ of individuals
- A (collective) decision $d \in D$ is to be made (finite or infinite, election, matching, public good, allocation of resources)
- Each individual is characterized by a type $\theta_i \in \Theta_i$ that is private information (valuation, preferences).
- Each individual has preferences $u_i : D \times \Theta_i \to \mathbb{R}_+$ over decisions;

Decision and their properties

- Let $\Theta = \prod_i \theta_i$. A social choice function is a mapping $c : \Theta \to D$ that associates a decision $c(\theta)$ to a profile of types θ (if you knew the types).
- A social choice function is ex-post efficient (Paretian) if for every $\theta \in \Theta$ $c(\theta)$ is Pareto Optimal, i.e. there exists no $d \in D$ such that

$$\exists i \in N\ u_i(d, \theta_i) > u_i(c(\theta), \theta_i) \text{ and } \forall j \in N\ u_j(d, \theta_j) \geq u_j(c(\theta), \theta_j)$$

- Collective choices are usually made indirectly through institutions in which agents interact. A mechanism $M = (S_1, \cdots, S_n, g)$ is the formal representation of such an institution through:
  - A strategy space S_i for each agent.
  - An outcome function $g : \prod_{i=1}^m S_i \to D$

Examples of mechanism design problems

- Auctions: online advertising (google ad), radio spectrum
- Matching: assignment of students to schools, assignment of workers to jobs (Residents to Hospitals, Uber, ...)
- Design of electoral systems,
- Internet routing protocols,
- The 2007 Nobel Memorial Prize in Economic Sciences was awarded to Leonid Hurwicz, Eric Maskin, and Roger Myerson "for having laid the foundations of mechanism design theory".

Decision and their properties

- Let $\Theta = \prod_i \theta_i$. A social choice function is a mapping $c : \Theta \to D$ that associates a decision $c(\theta)$ to a profile of types θ (if you knew the types).
- Collective choices are usually made indirectly through institutions in which agents interact. A mechanism $M = (S_1, \cdots, S_n, g)$ is the formal representation of such an institution through:
  - A strategy space S_i for each agent.
  - An outcome function $g : \prod_{i=1}^{m} S_i \to D$

Dominant strategy implementation

- A mechanism M and a type profile θ induces a game in which a strategy for agent i is S_i and its payoff given a strategy profile $s \in \prod_{i=1}^{m} S_i$ is $u_i(g(s_1, \cdots, s_n), \theta_i)$.
- A strategy s_i is dominant at θ_i if for all $s_i' \in S_I$

$$u_i(g(s_i, s_{-i}), \theta_i) \geq u_i(g(s_i', s_{-i}), \theta_i)$$

- A mechanism M is a dominant strategy implementation of a social choice function c if there exists mappings $\sigma_i : S_i \to M_i$ such that for all $\theta \in \Theta$ and for all $i \in N$, $\sigma_i(\theta_i)$ is a dominant strategy at θ_i and $g(\sigma_1(\theta), \cdots, \sigma_n(\theta)) = c(\theta_1, \cdots, \theta_n)$
- One also says that c is a performance of M.

Revelation principle

- A mechanism M is said to be a direct implementation of a social choice function $c : \prod_{i=1}^{m} \theta_i \to D$ if $S_i = \Theta_I$
- A direct mechanism is said to be dominant strategy incentive compatible (or strategy-proof) if for all $i \in N$, θ_i is a dominant strategy (i.e. $\sigma_i(\theta_i) = \theta_i$).

Proposition

If c can be implemented in dominant strategy (by an arbitrary mechanism) than it can be implemented by an inventive compatible dominant strategy direct mechanism

Proof:

Let $M = (S, g)$ be a dominant strategy implementation of c with associated strategy mappings σ_i. We then define the direct mechanism δ as: $\delta(\theta_1, \cdots, \theta_n) = g(\sigma_1(\theta_1), \cdots, \sigma_n(\theta_n))$.

- It suffices to focus on the space of direct mechanisms !

Bayesian-Nash implementation

- A Bayesian mechanism is a mechanism together with (common) priors F_i on the type of agents.
- In this setting, a Bayes-Nash equilibrium is a strategy profile σ such that for all $i \in N$ and all σ_i', one has:

$$\int_{\Theta_{-i}} u_i(g(\sigma_i(\theta_i), \sigma_{-i}(\theta_{-i})), \theta_i) dF_{-i}(\theta_i) \geq \int_{\Theta_{-i}} u_i(g(\sigma_i'(\theta_i), \sigma_{-i}(\theta_{-i})), \theta_i) dF_{-i}(\theta_i)$$

- One then says that the mechanism is a Bayesian-Nash implementation of the social choice function g.
- The relevation principle holds: if a social choice function is Bayesian-Nash implementable, it is implementable by a direct mechanism.

Impossibiliy theorem

- Consider the case where types are utilities over a finite set of outcomes A (say candidates to an election): $\Theta_i = \mathbb{R}_+^A$
- The decision to be made is the choice of one alternative in A, i.e. $D = A$
- A social choice function c is a dictatorship if there exists a player $i \in N$ such that for all utility profile $(\theta_1, \cdots, \theta_n)$, one has

$$c(\theta_1, \cdots, \theta_n) = \text{argmax}_{a \in A}\, \theta_i(a)$$

- A social choice function c is unanimous if for all utility profile $(\theta_1, \cdots, \theta_n)$, one has

$$(\forall i\ a_0 = \text{argmax}_{a \in A}\, \theta_i(a)) \Rightarrow c(\theta_1, \cdots, \theta_n) = a_0$$

Theorem (Gibbard-Satterthwaite impossibility theorem)

If $|A| \geq 3$, the only unanimous choice functions that can be implemented in dominant strategy are dictatorships

Beyond Impossibiliy theorem

- Impossibility theorem: it appears that the (generic) applicability of mechanism design is limited.
- In the context of auctions: impossibility to assign the object on the basis of utilities/valuations only (or the seller is a dictator)
- But the introduction of money allows to overcome this impossibility.
- More generally, given a social choice function $c : \Theta \to D$, let us consider a payment function $p_i : \Theta \to \mathbb{R}_+$ and quasi-linear utilities of the form:

$$u_i(a, \theta) = \theta_i(a) - p_i(\theta)$$

VCG mechanism

Definition

A direct mechanism with transfer (c, p) is called a Vickrey-Clarkes-Grove (VCG) mechanism if

1. for all $\theta \in \Theta$, $c(\theta) \in \operatorname{argmax}_{a \in A} \sum_{i \in N} \theta_i(a)$
2. for all $i \in N$, there exists $h_i : \Theta_{-i} \to \mathbb{R}$ such that

$$p_i(\theta) = h_i(\theta_{-i}) - \sum_{j \neq i} \theta_j(c(\theta))$$

Theorem

Every VCG mechanism is incentive compatible.

Clarke Pivot rule

Definition

The Clarke pivot rule is such that for all $i \in N$,

$$h_i(\theta_{-i}) = \max_{b \in A} \sum_{j \neq i} \theta_j(b).$$

One then has:

$$p_i(\theta) = \max_{b \in A} \sum_{j \neq i} \theta_j(b) - \sum_{j \neq i} \theta_j(c(\theta))$$

Proposition

Under the Clarke pivot rule, the VCG mechanism makes non-positive transfers and is individually rational provided the θ_is have non-negative values

Interpretation of the VCG mechanism

- The VCG mechanism charges a player for the externality its presence creates on the others.
- In the context of a single-object auction, the VCG mechanism coincides with the second-price auction.

Outline

Reading list for next week

Edelman and al. (2007). "Internet advertising and the generalized second-price auction: Selling billions of dollars worth of keywords". American economic review, 97(1), 242-259.

Position auctions (sponsored search)

- Sponsored search can be understood as a multi-object auction in which the seller (the platform) sells K slots/objects of decreasing quality.
- Slot k yields an (expected) click rate of α_k and $k < \ell \Rightarrow \alpha_k > \alpha_\ell$.
- Each advertiser i has a value per click of θ_i.
- Hence if the gets slot k at cost c per click, advertiser i has utility $\alpha_k(\theta_i - c)$
- We consider auctions where each advertiser i submits a bid b_i for a single slot, let $g(j)$ be the jth highest bidder.

VCG for position auction

Proposition

In the VCG auction:

- *Bidder $g(j)$ gets assigned the jth slot.*
- *If $N \geq K$, bidder $g(K)$ pays 0. If $N < K$, he pays $p_N^V := \alpha_N b_{g(N+1)}$*
- *By recursion, one has for all $i \leq \min(N-1, K-1)$*

$$p_i^V = p_{i+1}^V + (\alpha_i - \alpha_{i+1}) b_{g(i+1)}$$

Generalized second price auction

Definition

The most common auction used in practice is the generalized second price auction whereby:

- Bidder $g(j)$ gets assigned the jth slot.
- If $N \geq K$, bidder $g(K)$ pays 0. If $N < K$, he pays $p_N^V := \alpha_N b_{g(N+1)}$
- For $i \leq \min(N-1, K-1)$, the payment is $\alpha_i b_{g(i+1)}$

Proposition

Truth-telling is not a dominant strategy under GSP.

Long-term properties of GSP

- Consider GSP repeated over time for multiple searches.
- Eventually types become common-knowledge.
- What are properties of long-term stable vector of bids in the resulting game Γ ?

Definition

An equilibrium of the simultaneous-move game induced by GSP is locally envy-free if a player cannot improve his payoff by exchanging bids with the player ranked one position above him. More formally, in a locally envy-free equilibrium, for any $i \leq \min(N-1, K)$ one has

$$\alpha_i \theta_{g(i)} - p^{g(i)} \geq \alpha_{i-1} \theta_{g(i-1)} - p^{g(i-1)}$$

Long-term properties of GSP

- Assume advertisers are ordered by decreasing valuation and use the following vector of bids.
- $b_1^* = \theta_1$ and, for all $2 \leq j \leq \min(N-1, K)$, $b_j^* = {}^{p^V_{j-1}}/_{\alpha_{j-1}}$

Theorem

The strategy profile B^* is a locally envy-free equilibrium of the game induced by GSP. In this equilibrium, each advertiser's position and payment are equal to those in the dominant-strategy equilibrium of the game induced by VCG.

Reading list for next week

- Conitzer and al. (2022). "Pacing equilibrium in first price auction markets." Management Science 68.12 (2022): 8515-8535.
- Despotakis, and al. (2021). First-price auctions in online display advertising. Journal of Marketing Research, 58(5), 888-907.

Outline

Sponsored Search vs. RTB: Conceptual Comparison

	Sponsored Search	**RTB / Display**
Object sold	Ad positions (slots)	Individual impressions
Auction type	Position auction (GSP)	Single-item auction
Market structure	Few ordered slots, repeated advertisers	Massive, heterogeneous impressions
Intermediation	Single platform	Multiple intermediaries (SSP, DSP, ADX)
Time scale	Repeated but slow	Real-time, per-impression
Values and information	Per-click private values	Data-driven, learned, multi-dimensional
Pricing model	Cost per Click (CPC)	Cost per Mille (CPM)
Learning	Limited	Continuous online learning
Risk borne by	Platform (CTR uncertainty)	Advertiser (response risk)
Bidding unit	Value per click	Value per impression
Role of budgets	Secondary	Central (pacing across auctions)

RTB Auctions: From Second-Price to First-Price

- **Early RTB (late 2000s – early 2010s): Second-price auctions**
  - Inspired by Vickrey auctions and their truthfulness guarantees.
  - Advertisers (DSPs) bid their estimated value per impression.
  - Conceptually simple and aligned with auction theory.
- **Practical frictions emerge**
  - Presence of multiple intermediaries (SSP, ADX, DSP).
  - Reserve prices and *soft floors* distort payments.
  - Budget constraints and repeated participation break truthfulness.
  - Advertisers shade bids despite nominal second-price rules.
- **Header bidding and competition among exchanges**
  - Publishers solicit bids from multiple exchanges simultaneously.
  - Effective pricing increasingly resembles first-price auctions.
  - Transparency and predictability of second-price auctions deteriorate.
- **Industry transition (mid–late 2010s): First-price auctions**
  - Major platforms explicitly adopt first-price auctions.
  - Bidders adjust via bid shading and budget pacing.
  - First-price auctions better reflect actual payments and incentives.

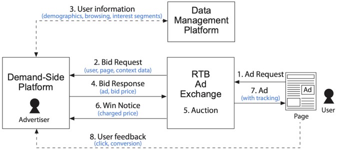

Figure 2.2: How RTB works for behavioural targeting. Source: [Zhang et al., 2014b].

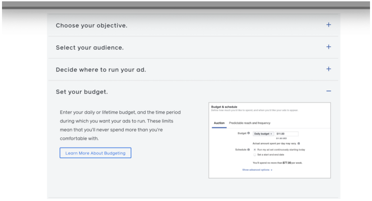

Figure 1 **Flow to create a Facebook ad as described at** `https://www.facebook.com/business/ads`.

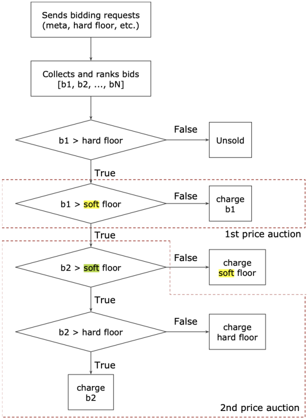

Figure 6: An illustration of the auction process in modern ad exchanges. The auction mechanisms are mixed and extended by introducing floor prices. The soft floor price is also in popularity but the impact of such mixture is largely unaware of. The complicated setting puts advertisers in an unfavourable position and could damage the eco-system.

Pacing Equilibrium in First Price Auction Markets

- With RTB, strategic focus shifts to budget and audience reached rather than bid per se.
- Proxy/automatic bidders and substantial learning efforts on conversion probability.
- In a system where advertisers' primary or only strategic lever is the budget, what is the best way to allocate and price ads ?
- Proxy bidders shade bids on behalf of advertisers to maximize their utilities over the course of the campaign.
- When the proxy bidder is designed to shade bids so the advertiser's budget is exhausted at the end of the budget horizon, would the system as a whole perform better if individual impressions are sold through a first price or a second price mechanism?
- First price pacing equilibrium (FPPE) vs second price pacing equilibrium (SPPE)

Pacing Equilibrium in First Price Auction Markets

- We consider a set of N bidders and a set of M (divisible) goods.
- Each bidder i has a valuation $\theta_{i,j}$ for each good j and a total budget B_i
- An allocation $x \in [0,1]^{N \times M}$ specifies the share $x_{i,j}$ of object j allocated to i (or the probability of allocation).
- A pacing multiplier for agent i is a real number $\alpha_i \in [0,1]$ that is used (by the platform) to scale the bids of agent i.

First price pacing equilibrium

Definition

A pair (α, x) is a budget-feasible first price pacing multipliers (BFPM) if

- For all j, $p_j := \max_{i \in N} \alpha_i v_{i,j}$ (first price pacing)
- If $x_{i,j} > 0$ then $\alpha_i v_{i,j} = \max_{k \in N} \alpha_k v_{k,j}$ (goods go to highest bidders)
- For all i, $\sum_{j \in M} x_{i,j} p_j \leq B_i$ (budget constraint).
- For all j, $\sum_{i \in N} x_{i,j} \leq 1$
- For all j if $p_j > 0$ then $\sum_{i \in N} x_{i,j} = 1$

Definition

A BFPM is a first price pacing equilibrium (FPPE) is for all $i \in N$, one has:

$$\sum_{j \in M} x_{i,j} p_j \leq B_i \Rightarrow \alpha_i = 1$$

Second price pacing equilibrium

Definition

A pair (α, x) is a second price pacing equilibrium (SPPE)

- For all j, $p_j := \max_{\{i \in N | \alpha_i v_{i,j} < \max_{k \in N} \alpha_k v_{k,j}\}} \alpha_i v_{i,j}$ (second price pacing)
- If $x_{i,j} > 0$ then $\alpha_i v_{i,j} = \max_{k \in N} \alpha_k v_{k,j}$ (goods go to highest bidders)
- For all i, $\sum_{j \in M} x_{i,j} p_j \leq B_i$ (budget constraint).
- For all j, $\sum_{i \in N} x_{i,j} \leq 1$
- For all j if $p_j > 0$ then $\sum_{i \in N} x_{i,j} = 1$

Properties of first price pacing equilibrium

Theorem

There exists an unique FPPE. The pacing multipliers are uniquely determined and such that revenue is maximized over all BFPM.

Theorem

Any FPPE $(\alpha, \overline{x}, \overline{p})$ *is a competitive equilibrium in the sense that*

1. *For all* $i \in M$, x_i *is a solution to*

$$\max_{\{x_i | \sum_{j\in M} p_j x_{i,j} \leq B_i\}} \sum_{j\in M} (\theta_{i,j} - \overline{p}_j) x_{i,j}$$

2. *For all* j *such that* $p_j > 0$, *one has* $\sum_{i\in N} x_{i,j} = 1$.

Proposition

An FPPE is shill proof (the platform doesn't have incentives to add fake bids).

FPPE vs SPPE

	SPPE	FPPE
Exists?	Yes	Yes
Buyers best responding?	Yes	No
Is market eq.?	Yes	Yes (even for supply-unaware buyers)
Is unique?	No	Yes, in utilities, multipliers, and prices
Is efficiently computable?	PPAD-complete	Convex program
Is welfare monotone?	No	Yes, in goods
Is revenue monotone?	No	Yes, in goods/bidders/budgets
Is shill proof?	No	Yes
Simulated regret/IC	No regret, very small IC	small regret, small IC
Simulated revenue	SPPE $\leq$ FPPE	
Simulated welfare	Ambiguous	

Table 1 A comparison of FPPE and SPPE. In the second row, we note that under SPPE, each bidder's multiplier is a mutual best response given other-agent bids are fixed—and would be even if the set of deviating strategies were expanded to allow the bidder to submit arbitrary per-auction bids (for a proof, see Proposition 1 in Conitzer et al. 2021). Under FPPE, a buyer may increase its utility by shading its multiplier (or shading individual bids, if it could submit arbitrary bids for each auction).

TODOS

- Check litterature on optimisation of reserve prices using bandit algorithms.
- Finally automatic bidding (Aggarwal and al. 2024): there discuss bandit algorithms....

Reading list for next week

- Nisan, N., Roughgarden, T., Tardos, E., and Vazirani, V. V. (Eds.). (2007). Algorithmic Game Theory. Cambridge University Press.
- Chapter 10: Mechanism design without money

Outline

Outline

Motivation

- A number of socio-economic situations involves the matching of actors and/or objects:
  - Assignments of students to schools
  - Assignments of tasks to teams
  - Assignments of public housing to demanders
  - Marriage
  - Formation of pairs among group of players.
- Characteristics of matching problems:
  - one-to-one, many-to-one, many-to-many,
  - one-sided or two-sided preferences,
  - bipartite our unipartite.

Statement of the problem

A “marriage" problem is given by

- A set of men $M = \{m_1, \cdots, m_p\}$
- A set of women $W = \{w_1, \cdots, w_q\}$
- A preference profile $(\succ_{m_1}, \cdots, \succ_{m_p}, \succ_{w_1}, \cdots, \succ_{w_q})$ where:
  - For $m \in M$, $\succ_m$ is a strict preference over $W \cup \{m\}$,
  - For $w \in W$, $\succ_w$ is a strict preference over $M \cup \{w\}$

  A solution to the problem is a matching function $\mu : M \cup W \to M \cup W$ such that:
  - for all $m \in M$, $\mu(m) \in W \cup \{m\}$
  - for all $w \in W$, $\mu(w) \in M \cup \{w\}$
  - for all $(m, w) \in M \times W$, $\mu(m) = w \Leftrightarrow \mu(w) = m$
- A mechanism design problem with preference as types and matching as allocations.

Stable matching

- A matching is blocked by an agent $i \in M \cup W$ if $iP_i\mu(i)$
- A matching is blocked by a pair $(m, w) \in M \times W$ if $w \succ_m \mu(m)$ and $m \succ_w \mu(w)$, i.e. both would be better of being matched together than with their current partners.
- A matching is stable if it is not blocked by an agent or a pair of agent.

Stable matching and the core

Proposition

The set of stable matchings coincide with the core of the associated cooperative game.

Proof:

N.B. The core of a cooperative game is the set of allocations that cannot be improved upon by a coalition of players.
It is clear than an element in the core is a stable matching.
Conversely, if μ is stable. Assume there exists a matching ν with which a coalition S blocks μ. Then for some $i \in S$ $\nu(i) \succ_i \mu(i)$. Either $\nu(i) = i$ in which case i blocks μ and the matching is not stable. Either $\nu(i) = j$ and then one must have $j \in S$ and $\nu(j) = i \succ_j \mu(j)$ in which case (i, j) blocks μ.

Deferred acceptance algorithm

Theorem

There exists a stable matching

Proof:

see deferred acceptance algorithm below.

Deferred acceptance algorithm:Informal description.

- Each man proposes to his preferred woman. Each woman tentatively accepts her most preferred acceptable proposal (if any).
- Then, each unmatched man proposes to his preferred woman among those that have not yet rejected him.
- The algorithm ends when there are no new proposals. Each woman is matched with her current proposal. Any unmatched agent remains single.

Algorithm also known as Gale-Shapley algorithm.

Deferred acceptance algorithm: formal description.

- Initialization:
  - $U = M$, (set of unmatched men at time 0)
  - $W_m = \{w \in W \mid w \succ_m m\}$ (set of potential partner for m at time 0)
  - for all w, $\mu_w = w$ (initial matching of women)
  - for all m, $\mu_m = m$ (initial matching of men)
  - $t = 0$
- While $U \neq \emptyset$ and $\cup_{m \in U} W_m \neq \emptyset$ do
  - For all $m \in U$: $w_m = \max_{\succ_m} W_m$ and $W_m = W_m / \{w_m\}$
  - For all w :
    - $M_w = \{m \mid w_m = w\}$ and $m_w = \max_{\succ_w} M_w$
    - If $m_w \succ_W \mu_w$ then $m'_w = \mu_w$; $\mu_w = m_w$; $\mu_{m_w} = w$; $U = U / \{m_w\}$
    - If $m'_w \neq w$ then $U = U \cup \{m'_w\}$
- The algorithm stops in finite time because $\cup_{m \in U} W_m$ is strictly decreasing during each iteration.
- The women get (weakly) better of and the mean (weakly) worst of as the algorithm proceeds.

Stability of Gale-Shapley algorithm

The resulting matching is stable:

- No woman w has accepted a men worst than w, no man m has proposed to a woman worst than m
- If a pair (m, w) blocks μ, one must have $w \succ_m \mu(m)$ and $m \succ_w \mu(w)$. Yet, $w \succ_m \mu(m)$ implies m has proposed to w during the algorithm. But as $m \neq \mu(w)$, this implies that w has rejected m at some point in the algorithm and one should thus have $\mu(w) \succ_w m$. This yields a contradiction.

Properties of stable matchings I

- In general, there exist multiple stable matchings, e.g. men-proposed stable matching μ^m and women-proposed stable matchings μ^w.

Theorem

The outcome of the men proposed stable matching is weakly better than that of any other stable matching for every men.

Proof:

- *Let us say that a woman w is achievable for a man m if there exists a stable matching μ such that $\mu(m) = w$.*
- *Assume $\overline{\mu}$ (from Gale-Shapley) does not assign achievable woman to at least a man, i.e. is not weakly better than every other stable matching.*
- *Let us consider the first step of the Gale-Shapley algorithm in which a man m is rejected by an attainable woman w . Then, there must be m' such that $m' \succ_w m$ (and m' and w get matched at this stage of the algorithm). As this is a first time a men gets rejected, one also has $w \succ_{m'} \mu(m')$ (otherwise m' would have already been rejected by admissible woman $\mu(m')$). Thus (m', w) blocks μ contradicting the fact that μ is stable.*

Properties of stable matchings II

Interests of men and women are in fact "opposed" in the matching process

Theorem

The outcome of the men proposed stable matching is weakly worst than that of any other stable matching for every woman.

Proof:

- *Let μ^M be the men-proposed stable matching.*
- *Suppose μ is stable and such that for some w $\mu^M(w) \succ_w \mu(w)$.*
- *Then w can not be single at μ^M. Let $m = \mu^M(w)$.*
- *One has $m \neq \mu(w)$ and $w \neq \mu(m)$. Given μ^M is men-optimal, one has $\mu^M(m) = w \succ_m \mu(m)$.*
- *Hence $w \succ_m \mu(m)$ and $m \succ_w \mu(w)$. Thus (m, w) blocks μ.*

Properties of stable matchings III

Theorem

The set of matched agents is identical at every stable matching.

Proof:

- *Let $W_\mu = W \cap \mu(M)$ (resp. $M_\mu = M \cap \mu(W)$) be the set of women (resp. men) matched at matching μ.*
- *One has using man-optimality of μ^M and man-pessimality of μ^W that for all stable μ,*

$$|M_{\mu^M}| \geq |M_\mu| \geq |M_{\mu^W}|$$

- *Similarly, for all stable μ,*

$$|W_{\mu^W}| \geq |W_\mu| \geq |W_{\mu^M}|$$

- *As $|M_{\mu^M}| = |W_{\mu^M}|$ and $|W_{\mu^W}| = |M_{\mu^W}|$, all these inequalities are equalities.*
- *In particular $|M_{\mu^M}| = |M_\mu|$ but, by man-optimality, any man matched at $:$ mu must be matched at μ^M, thus $M_{\mu^M} = M_\mu$*
- *A similar argument holds for woman*

Matching and strategic behavior I

A matching problem (M, W, P), can be embedded in game in which:

- The strategy space of each agent is the set $\mathcal{P}_i$ of preference relations and $\mathcal{P} = \prod_i \mathcal{P}_i$
- A matching mechanism ϕ associates to a stated preference profile R, a matching $\phi(R)$.
- The preference of agent i over outcomes is given by its true preferences P_i.

Theorem

There is no matching mechanism that is both stable and strategy-proof.

- Nevertheless, truthtelling is a (weakly) dominant strategy for men in the men-proposing Gale-Shapley algorithm (as the outcome if the men-optimal matching).

Matching and strategic behavior II

Consider the following strategy profile for 2 men and 2 women:

$$P = \begin{array}{|c|ccc|}
\hline
\succ_{m_1} & w_1 & w_2 & m_1 \\
\hline
\succ_{m_2} & w_2 & w_1 & m_2 \\
\hline
\succ_{w_1} & m_2 & m_1 & w_1 \\
\hline
\succ_{w_2} & m_1 & m_2 & w_2 \\
\hline
\end{array}$$

- μ^M and μ^W are the only stable mechanisms. Assume wlog $\phi(P) = \mu^M$.
- Then consider P' where w_1 misrepresents her preference as (m_2, w_1, m_1).
- the only stable matching then is μ^W, one has $\phi(P') = \mu^W$ and w_1 is better-of.
- Hence ϕ is not strategy-proof.

Many to one matching

- Many to one matching
- Settings where one kind of agent can be matched with many agents of the other kind.
- School choice, assignment of/to tasks...

Statement of the problem

A “school choice" problem is given by

- A set of students $S = \{s_1, \cdots, s_m\}$
- A set of colleges $C = \{c_1, \cdots, c_n\}$
- A vector of capacities for colleges $Q = (q_1, \cdots, q_n)$
- A strict preference profile $\succ_s$ over $C \cup \{s\}$ for each student s.
- A strict preference profile $\succ_c$ over $S \cup \{c\}$ for each college c. This is assumed to induce a preference profile over groups of students 2^S under the following assumption (responsiveness):
  - For any $T \subset S$ such that $|T| < q_c$, and any $s \in S/T$:
    $$T \cup \{s\} \succ_c T \Leftrightarrow s \succ_c c$$
  - For any $T \subset S$ such that $|T| < q_c$, and $s, s' \in S/T$:
    $$T \cup \{s\} \succ_c T \cup \{s'\} \Leftrightarrow s \succ_c s'$$

Matching for the school choice problem

A solution to the problem is a matching correspondence $\mu : M \cup W \to M \cup W$ such that:

- for all $c \in C$, $\mu(c) \subset S$ and $|\mu(c)| \leq q_c$
- for all $s \in S$, $\mu(s) \subset C \cup \{s\}$ and $|\mu(s)| = 1$
- for all $(c, s) \in C \times S$, $s \in \mu(c) \Leftrightarrow \mu(s) = \{c\}$

Stable matching

- A matching μ is blocked by a college c if there exists $s \in \mu(c)$ such that $cP_c s$.
- A matching μ is blocked by a student s if $sP_s\mu(s)$.
- A matching μ is blocked by a pair $(s, c) \in S \times c$ if $cP_s\mu(s)$ and either
  - There exists $s' \in \mu(c)$ such that $sP_c s'$
  - $|\mu(c)| < q_c$ and $sP_c c$.

Proposition

The set of stable matchings coincide with the weak domination core.

Deferred acceptance algorithm

One can use the one-to-one algorithm and treat each seat in a college as the copy of one agent with a given preference. This results in the following algorithm:

- Each student applies to his preferred college. Each college tentatively accepts students up to its capacity.
- Each unmatched student applies to his most preferred college that has not yet rejected him.
- The algorithm ends when there are no new applications.

Proposition

The outcome of the deferred acceptance algorithm is a stable matching.

Properties of Gale-Shapley algorithm in the school-choice problem

- The student proposing algorithm is student-optimal and strategy proof for students.
- The student proposing algorithm is the worst for colleges.
- Same sets of students and colleges get matched at every stable matching.
- There are no strategy-proofs mechanisms for colleges (colleges act like a coalition of players and have incentives to misreport capacity).

Outline

Setting the problem

- M sellers with an identical object to sale and limit prices $(a_1, \cdots, a_M)$
- N buyers willing to buy the object with limit prices $(b_1, \cdots, b_N)$
- How do buyers and sellers get match and at what price do they trade ?
- Applications to stock market (NYSE, Euronext,..) but also online advertising, electricity markets, allocation frequency spectrum

Auction design

- Which information do the traders have ?
- How are buyers and sellers matched ?
- Centralized or decentralized mechanism ?
- At what price does trade takes place ?
- Continuous or periodic clearing of orders ?

Double auction at the NYSE

- At opening and at closing, call market/periodic auction:
  - orders for the opening and closing market are entered during a certain time-interval preceding opening and closing.
  - All feasible trade are performed at a common “reference" price.
- After opening, continuous double auction.
  - A trade is made every time a match is made between a seller and a buyer.

Double auction mechanisms

- A double auction mechanism is a mapping $a : \mathbb{R}_+^M \times \mathbb{R}_+^N \to \{0,1\}^M \times \{0,1\}^N \times \mathbb{R}_+^M \times \mathbb{R}_+^N$ that associates to a vector of ask price for sellers $s \in \mathbb{R}_+^M$ and bid price for buyers $b \in \mathbb{R}_+^N$, an allocation $x \in \{0,1\}^M$ and a payment $p \in \mathbb{R}^M$ for sellers as well as an allocation $y \in \{0,1\}^N$ and a payment $q \in \mathbb{R}^N$ for buyers with the following characteristics.
  - $x_i = 1$ if the seller sold the object and then $p_i \geq s_i$ is the selling price (otherwise $x_i = 0$ and $p_i = 0$)
  - $y_j = 1$ if the buyer bought the object and then $q_j \leq b_j$ is the buying price (otherwise $y_j = 0$ and $q_j = 0$)
  - One has $\sum_i x_i = \sum_j y_j$.

Equilibrium mechanism

- Assume bid prices are ordered by decreasing value $b_1 \geq b_2 \geq \cdots \geq b_N$
- Sell prices are ordered by increasing value $s_1 \leq s_2 \leq \cdots \leq s_M$
- Let k be the largest index such that $b_k \geq s_k$
- Any price in $[\max(s_k, b_{k+1}), \min(b_k, s_{k+1})]$ is an equilibrium price since supply equals demand.
- In the following we assume to simplify notation that $b_{k+1} < s_k$ and $b_k < s_{k+1}$ so that the set of equilibrium price is $[s_k, b_k]$

Profit and efficiency

- The profit of a seller i that sells the object at price p_i is $p_i - s_i$.
- The profit of a buyer j that buys the object at price q_j is $b_j - q_j$.
- All other buyers and sellers have profit zero.
- Total profit/welfare at an allocation (x, y, p, q) is given by $W_a = \sum_{\{i|x_i=1\}} p_i - s_i + \sum_{\{j|y_j=1\}} b_j - q_j$.
- At any equilibrium/welfare allocation with price $\overline{p}$, the total profit is given by

$$W_e = \sum_{\{i|x_i \in X_E\}} \overline{p} - s_i + \sum_{\{j|y_j \in Y_e\}} b_j - \overline{p} := \sum_{\{j|y_j \in Y_e\}} b_j - \sum_{\{i|x_i \in X_E\}} s_i$$

  where X_e and Y_e denote respectively the set of buyers and sellers matched at equilibrium.
- Total profit is the same at every equilibrium

Equilibrium properties

- Equilibrium is individually rational: no agent loses from participation.
- Equilibrium is efficient: it maximizes total profit (among individually rational allocations) because the agents with the highest profit get matched.
- The quantity traded at equilibrium is the optimal traded quantity.
- N.B: relation to welfare theorems

Desirable properties of double auction mechanisms

- Efficiency
- Individual rationality
- Strategy proofness: no agent has incentives to misreport his price.
- Budget balance: no external source of funding required (all exchanges are between buyers and sellers).

α- double auction

- The α- double auction mechanism corresponds to the equilibrium allocation with the price set at $\alpha s_k + (1 - \alpha) b_k$
- The α-double allocation is efficient, individually rational and budget balanced.
- But not strategy proof: buyers have incentives to under-report their value, sellers to over-report.

Myerson–Satterthwaite Theorem I: bayesian game

- Double auction as a Bayesian game. Let $a = (x, y, p, q)$ be a double allocation mechanism.
- Values of sellers are assumed to be independently drawn (in $[0, 1]$) according to measure μ
- Values of buyers are assumed to be independently drawn (in $[0, 1]$) according to measure ν
- Each seller/buyer knows its true value $\bar{s}_i / \bar{b}_j$
- Each seller/buyer has $\Sigma : \mathbb{R}_+ \to \mathbb{R}_+$ as strategy space and reports value $S_i(\bar{s}_i) / B_j(\bar{b}_j)$ as functions of the true value.
- Allocation $x_{S,B}, y_{S,B}, p_{S,B}, q_{S,B}$ are random variables that depend on the actual realization of values.
- Expected payoff given strategies

$$\pi_{i,(S,B)}(\bar{s}_i) = \int x_{i,(S,B)} \left(p_{i,(S,B)} - \bar{s}_i\right) d(\delta_{\bar{s}_i} \otimes \mu_{-i}) \otimes \nu$$

$$\pi_{j,(S,B)}(\bar{b}_j) = \int y_{j,(S,B)} \left(\bar{b}_j - q_{i,(S,B)}\right) d\mu \otimes (\delta_{\bar{b}_j} \otimes \nu_{-j})$$

- Bayesian-Nash equilibrium (S^*, B^*): each agent maximizes expected payoff given strategies of other players: for all $(\bar{s}, \bar{b})$:

$$\pi_{i,(S^*,B^*)}(\bar{s}_i) \geq \pi_{i,((S_i,S^*_{-i}),B^*)}(\bar{s}_i) \text{ and } \pi_{j(S^*,B^*)}(\bar{b}_j) \geq \pi_{j,(S^*,(B_j,B^*_{-j}))}(\bar{b}_j)$$

Desired properties

- Truthful strategies are $\mathcal{S}_i(s_i) = s_i$ and $\mathcal{B}_j(b_j) = b_j$
- A double auction mechanism is:
  - Budget balance is for all (s, b)
    $\sum x_i(s, b) p_i(s, b) \leq \sum y_j(s, b) q_j(s, b)$
  - Ex-ante individually rational, if for all (s, b), and all i, j
    $\pi_{i,(\mathcal{S},\mathcal{B})}(\bar{s}_i) \geq 0$ and $\pi_{j,(\mathcal{S},\mathcal{B})}(\bar{b}_j) \geq 0$.
  - Efficient if for all (s, b), (x, y, p, q) maximizes
    $\sum_i p_i(s, b) - x_i(s, b) s_i + \sum_j y_j(s, b) b_j(s, b) - q_j(s, b)$
  - Bayesian incentive compatible if for all (s, b), $(\mathcal{S}, \mathcal{B})$ is an equilibrium of the associated bayesian game.
- Theorem (Myerson-Satherwaite+ Rustichini): there is no double auction mechanism that is budget balance, individually rational, efficient and incentive compatible

Proof of Myerson-Satherwaite Theorem

- Case of 1 buyer and 1 seller with priors whose support interact.
- See Myerson and Satterthwaite (1983) for details.

Price of anarchy and asymptotic efficiency

- Given a double-auction mechanism that is incentive compatible and budget balanced, let $W_a(s, b)$ be the total welfare at the worst Bayesian Nash allocation associated to a (given (s, b)).
- Let $W_e(s, b)$ be the maximum welfare (i.e. the welfare at the competitive equilibrium).
- The ratio $W_e(s, b)/W_a(s, b)$ is called the price of anarchy: the price society has to pay for agents being selfish.

Asymptotic and experimental efficiency of the α-double auction

- In α double auction Rustichini et al. (1994) show that price of anarchy tends to 1 as $\mathcal{O}(1/n^2)$ and that difference between true value and actual bids is $\mathcal{O}(1/n)$
- Behavioral experiments, starting from Smith (1962), have shown that double auction are efficient even for small numbers.

Trade reduction and McAfee mechanism

- Exemples of incentive compatible mechanism
- Reduced trade mechanism:
  - Let k be the largest index such that $b_k \geq s_k$.
  - The $k-1$ highest bidders and the $k-1$ lowest sellers trade at prices b_k for bidders and s_k for sellers
- McAfee mechanism
  - Let $p = {b_{k+1}+s_{k+1}}/{2}$. If buyer k and seller k accept to trade at p than k lowest sellers and k highest buyers trade at p.
  - Otherwise revert to the reduced trade mechanism.
- Price of anarchy is limited: at most one trade is lost.

Profit of the auctioneer

- Revenue of the auctioneer for a mechanism $a = (x, y, p, q)$ is given by $\sum y_j(s, b)q_j(s, b) - \sum x_i(s, b)p_i(s, b) \geq 0$.
- Alternative design question (from the point of view of the auctionneer, e.g. ebay): how to maximize revenue ?
- The question somehow relates to the profit of a market maker on a stock market.
- Maximum potential revenue (see Deshmukh et al. 2002)

Outline

References I

- Nisan, N., Roughgarden, T., Tardos, E., and Vazirani, V. V. (Eds.). (2007). Algorithmic game theory. Cambridge University Press.
- Gale, D., and Shapley, L. S. (1962). College admissions and the stability of marriage. The American Mathematical Monthly, 69(1), 9-15.
- Edelman, B., Ostrovsky, M., and Schwarz, M. (2007). Internet advertising and the generalized second-price auction: Selling billions of dollars worth of keywords. American economic review, 97(1), 242-259.
- Deshmukh, K., Goldberg, A. V., Hartline, J. D., and Karlin, A. R. (2002, September). Truthful and competitive double auctions. In European Symposium on Algorithms (pp. 361-373). Springer, Berlin, Heidelberg.
- McAfee, R. P. (1992). "A dominant strategy double auction". Journal of Economic Theory. 56 (2): 434

References I

- M. Babaioff; N. Nisan (2004). "Concurrent Auctions Across The Supply Chain". Journal of Artificial Intelligence Research. 21: 595–629.
- Myerson, R. (1981), “Optimal Auction Design,” Mathematics of Operations Research, Vol. 6, pp. 58–73.
- Myerson, Roger B.; Mark A. Satterthwaite (1983). "Efficient Mechanisms for Bilateral Trading". Journal of Economic Theory. 29 (2): 265–281
- Rustichini, Aldo; Satterthwaite, Mark A.; Williams, Steven R. (1994). "Convergence to Efficiency in a Simple Market with Incomplete Information". Econometrica. 62 (5): 1041–1063
- Smith, Vernon L. (1962). "An Experimental Study of Competitive Market Behavior". Journal of Political Economy. 70 (2): 111–137. doi:10.1086/258609. JSTOR 1861810.