Algorithms, Data Science, and Online Markets

Stefano Leonardi

Department of Computer, Control, and Management Engineering Antonio Ruberti (DIAG)
Sapienza University of Rome

September 2-6, 2019

Data Science Summer School - Pisa

Stefano Leonardi

- Leader of the Research Group on Algorithms and Data Science at DIAG - Sapienza
- ERC Advanced Grant "Algorithms and Mechanism Design Research in Online MArkets" (AMDROMA)
- Chair of the PhD program in Data Science at Sapienza
- Past Chair of the Master's Degree in Data Science at Sapienza
- Fellow of the European Association for Theoretical Computer Science

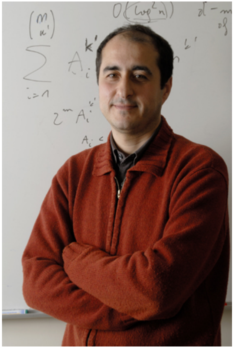

Research Interests:

- Algorithmic Theory
- Algorithmic Data Analysis
- Economics and Computation

Outline

1. Part I: Algorithms, Data Science and Markets
2. Part II: Internet, Equilibria and Games
3. Part III: Games and solution concepts
4. Part IV: The complexity of finding equilibria
5. Part V: The price of Anarchy
6. Part VI: Equilibria in markets
7. Conclusions

Algorithms, Data Science, and Markets

- Digital markets form an important share of the global economy.
- Many classical markets moved to Internet: real-estate, stocks, e-commerce, entarteinment
- New markets with previously unknown features have emerged: web-based advertisement, viral marketing, digital goods, online labour markets, sharing economy

An Economy of Algorithms

- In 2000, we had 600 humans making markets in U.S. stocks. Today, we have two people and a lot of software. One in three Goldman Sachs employees are engineers
  R. Martin Chavez, Chief Financial Officer at Goldman Sachs

  [Data,Dollars,and Algorithms: The Computational Economy, Harvard, 2017]

An Economy of Algorithms

Algorithms take many economic decisions in our life:

- Rank web pages in search engines
- Trade stocks
- Run Ebay auctions
- Price Uber trips
- Kidney exchange
- Internet dating
- Assign interns to hospitals and pupils to schools
- Sell Ads on Webpages
- Price electric power in grids

Success story 1: Internet Advertising

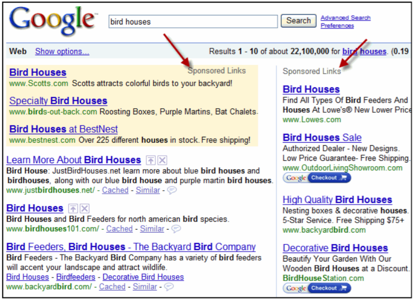

- Provide the major source of revenue of the Internet Industry, more than 90% for Google
- Electronic auctions are executed billions of times a day within the time frame of few hundred milliseconds.
- Many new auction design and big data algorithmic problems are motivated by online markets

Success story 1: Internet Advertising

Selling display ads on the spot market.

Success story 2: Digital Markets

- Need a theory for markets run by algorithms
- Do prices that induce efficient equilibria between buyers and sellers exist?
- Provide incentives to service providers (convince Uber riders to get up at night!) and to consumers to stay in the market.

Success story 2: Digital Markets

- Algorithmic problems in online markets are not standard since they work on inputs that are private information of economic agents
- Algorithmic mechanism design deals with the design of incentives that make agents to report honestly their private information to the algorithm.
- How hard is to find equilibria in markets operated by algorithms? If your laptop cannot find the equilibrium, your system cannot do it either!

Success story 3: Matching Markets

- Goal. Given a set of preferences among hospitals and med-school students, design a self-reinforcing admissions process.
- Unstable pair. Hospital h and student s form an unstable pair if both:
  - h prefers s to one of its admitted students.
  - s prefers h to assigned hospital.
- Stable assignment. Assignment with no unstable pairs.
- Natural and desirable condition.
- Individual self-interest prevents any hospital-student side deal.

Success story 3: Matching Markets

- Gale-Shapley algorithm computes a stable matching
- 2012 Nobel Prize in Economics:
  - Lloyd Shapley. Stable matching theory and GaleShapley algorithm.
  - Alvin Roth: Applied GaleShapley to matching med-school students with hospitals, students with schools, and organ donors with patients.

Algorithms are nowadays running matching markets also on digital platforms, large-scale organ transplants projects.

Success story 4: Online Labour Marketplaces

- Outsource complex tasks to workforce recruited on the cloud
- Algorithmic methods for job scheduling, task allocation, team formation, and distributed coordination.
- Incorporate fairness and diversity in the algorithms

Success story 4: Online Labour Marketplaces

- How can we form teams of experts online when compatibility between workers is modelled by a social network?
- How can we decide online when to use outsourced workers, when to hire workers in a team and when to fire inactive workers?
- How to limit the disparate impact of machine learning systems in online labor marketplaces and impose equality of gender and ethnic groups?
- How to provide the right incentives to workers and charge the right payments to outsourcing companies?

Outline

Internet, Equilibria in games

- The Internet is a socio-economic system formed by a multitude of agents (buyers, sellers, publishers, ISP, political organizations,..)
- The strategic interaction among Internet agents is regulated by algorithms
- The central notion of Game theory and Market economics is the one of equilibrium
- An equilibrium is an outcome of a game such that no agent has any incentive to deviate

Example 1: GPS Car Navigation

- A GPS car navigator chooses at any time the shortest path to destination
- Does this converge to an equilibrium or does it oscillate?
- Does it produce low congestion traffic?

Game theoretical and Algorithmic questions

- Does an equilibrium state exist?
- Does an efficient algorithm exist?
- How fast is the convergence to an equilibrium state?
- How efficient is the equilibrium state with respect to an optimal centralised solution
- How good is the market's invisible hand?
- Which type of incentives are needed to motivate agents to act in the global interest

Outline

1. Part I: Algorithms, Data and Markets
2. Part II: Algorithms, Complexity and games
3. Part III: Games and solution concepts
4. Part IV: The complexity of finding equilibria
5. Part V: The price of Anarchy
6. Part VI: Equilibria in markets
7. Conclusions

Prisoner's Dilemma - Dominant Strategies

- The most desirable notion of equilibrium is the dominant strategy equlibrium: each player has a best strategy to be played whatever strategy is played by the others
- The prisoner's dilemma has a dominant strategy: confess,confess
- A dominant strategy can be computed by analysing all the strategies of each player

c_1, c_2	confess	silent
confess	4, 4	1, 5
silent	5, 1	2, 2

Games in Strategic Normal Form

- A game is defined by a set of strategies for each agent.
- We consider one shot games
- The state of a game is the combination of strategies played by the agents
- In each state there is a payoff for each agent
- Players are rationals and selfish, their only goal is to maximise individual utility
- A game with two players is called a two-player game
- A game with sum of payoffs equal to 0 in each state is called zero-sum game

[Von Neumann and Morgenstern, 1944]
Many more definitions and practical settings

Battle of the Sexes - Pure Nash Equilibria

- There is no dominant strategy: the strategy played depends on the choice of the other agent
- There are two *Pure Nash Equilibria*: there is no incentive to deviate if the other player does not deviate
- To find a Pure Nash equilibrium it is required to analyse all the states of the game.

Game		Player 2	
		Boxing	Ballet
Player 1	Boxing	(2,1)	(0,0)
	Ballet	(0,0)	(1,2)

Rock Scissors Paper - Mixed Nash Equilibria

- It does not exist any Pure Nash Equilibria
- A mixed strategy is a probability distribution over a set of strategies, e.g., $1/3, 1/3, 1/3$.
- A *Mixed Nash Equilibrium* is a collection of mixed strategies - one for agent - such that no agent has any incentive to deviate.

Theorem (Nash, 1951)

It always exists a Mixed Nash Equilibrium in game in strategic normal form.

u_1, u_2	rock	paper	scissors
rock	0, 0	-1, 1	1, -1
paper	1, -1	0, 0	-1, 1
scissors	-1, 1	1, -1	0, 0

Outline

1. Part I: Algorithms, Data and Markets
2. Part II: Algorithms, Complexity and games
3. Part III: Games and solution concepts
4. Part IV: The complexity of finding equilibria
5. Part V: The price of Anarchy
6. Part VI: Equilibria in markets
7. Conclusions

Zero-sum games

The Mixed-Nash Equilibrium can be found efficiently in a two-player zero sum game [Von Neumann, 1928].
Application of the min-max principle:

- Assume the column player knows the strategy played by the row player.
- The column player will respond with the strategy that maximises her payoff
- Then, the row player will play the strategy that can be responded with the minimum maximum payoff of the opponent.

	C	D
A	2	−1
B	1	3

The min-max principle

[von Neumann 1928]

- The problem reduces to finding the extreme point of a polyedra described by a set of linear equations that maximises the minimum payoff.
- The problem can be solved efficiently (polynomial time) by a Linear Programming solver.

$$
\begin{aligned}
\max\ & v \\
& 2a + b \geq v \\
& -a + 3b \geq v \\
& a + b = 1 \\
& a, b \geq 0.
\end{aligned}
$$

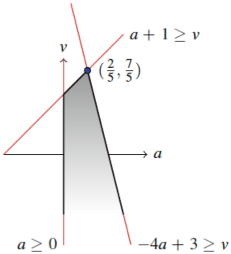

Mixed Nash Equilibria in General Games

- The complexity of the problem of computing a MNE in a two-player non zero-sum game has been open till very recently
- One possibility to reach an equilibrium state is to let the two players to play a best response game till they reach an equlibrium
- A MNE can be seen as the fixed point of a best response function $F(a1, a2) = (a1, a2)$ with (a1,a2) the two mixed strategies of the two players.
- The existence of a Nash Equilibrium can be demonstrated by using the Sperner's Lemma on the coloring of an arbitrarily dense triangle decomposition

Sperner's Lemma, 1928

- Vertices A,B and C have different colors
- All vertices on one side (e.g., AB) do not have the colour of the opposite vertex (e.g., C)
- the remaining vertices can have any colour
- Sperner's Lemma claims the existence of a triangle with the three vertices coloured differently
- A best response dynamic navigating the decomposition by only crossing black/white edges will eventually reach the triangle with three colours.

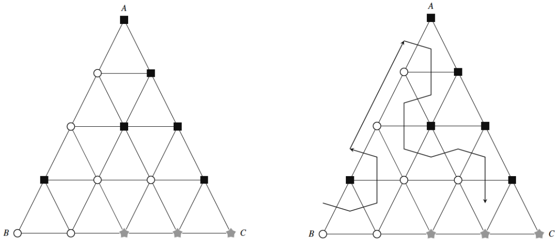

Complexity of finding a MNE

- The problem can be solved by enumerating all possible subset of strategies forming the support of the two mixed strategies.
- There are $2^{|S|}$ different supports for a set S of strategies
- The problem of finding an efficient algorithm for finding a MNE was opened for decades.
- The best response dynamic may take an exponential number of steps before to converge even in a two-player game.
  [Daskalakis, Goldberg and Papadimitriou, 2005, Chen and Deng, 2005]

Outline

The price of Anarchy

- Rational agents are only driven by their own interest
- They respond in any state outside equilibrium with a strategy which improves the individual utility.
- How good is the social welfare achieved at the equilibrium?
- Social welfare is defined as the sum of the payoffs of the agents.
- The tragedy of commons: the social welfare of an equilibrium is much worst that the optimum social welfare.
- The price of Anarchy [Koutsoupias and Papadimitriou, 1998] is a quantitative measure of this degradation.

Driving with a navigator

- Which is the impact on traffic of a GPS navigator that routes each car on a lowest latency path?
- Does it reach an equilibrium? Yes, it is a potential game! [Monderer and Shapley, 1996]
- How bad is the equilibrium with respect to an optimum routing scheme with cars obeying to a central coordinator?

Routing games

- Each agent needs to move a car move from source to destination
- The set of strategies is given by the different itineraries
- The travel time (latency) depends from the number of cars (flow) that choose the same itinerary
- The only equilibrium is the one with one unit of traffic on the bottom edge. It has cost $1x1 = 1$
- The optimal solution will split the traffic between the two itineraries, with a total cost $1/2x1 + 1/2x1/2 = 3/4$
- The Price of Anarchy is equal to $1/(3/4) = 4/3$.

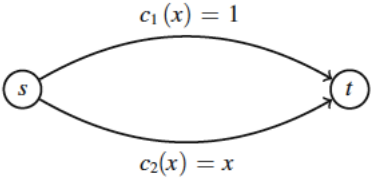

Braess' Paradox

- In the first network, the 1 unit flow splits at the equilibrium between the two paths with cost $0.5x(1 + 1/2) + 0.5(1 + 1/2) = 3/2$
- We now add a superfast link (0 cost) to improve our network
- In the second network, the whole traffic goes through the superfast link with a cost $1x(1 + 1) = 2$
- The price of Anarchy is equal to $2/(3/2) = 4/3$
- Tim Roughgarden and va Tardos [2001] proved that for any arbitrarily complicated network with linear delay costs on the links $(ax + b)$ the Price of Anarchy is never worst that $4/3$!

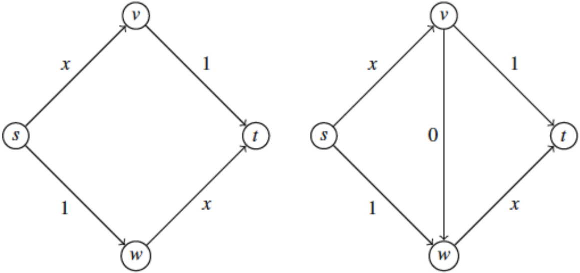

Outline

Internet Advertising

- Search Ads are sold with online electronic auction
- Goods on Ebay are sold with online electronic auctions
- Prices are set in order to bring markets to equilibria: Demand = Offer
- Prices are decided by algorithms for the Internet markets, the sharing economy and many other economic activities

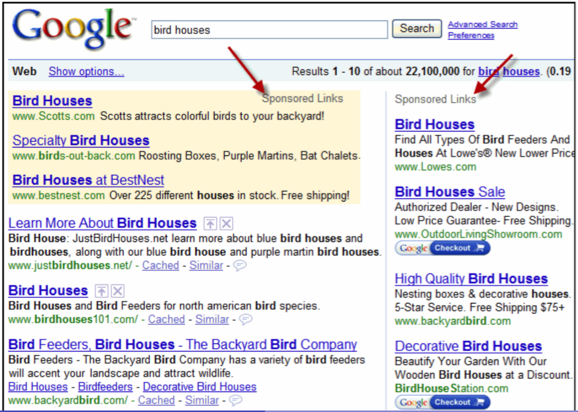

Auction design

- The internet advertising economy boomed since Google decided in 2004 to use the second price auction
- In second price auction the item is given to the bidder with highest bid at price equal to the second highest bid
- Before 2014, search ads were sold using the first price auction: the price is the highest bid
- First price auction does not posses a dominant strategy equilibrium

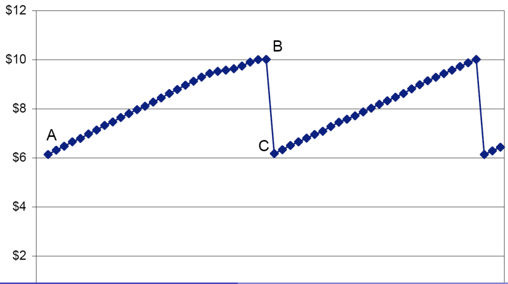

Vickrey Second Price Sealed Bid Auction [1961]

- Bidder i has valuation v_i for the good on sale
- Bidder i communicates bid b_i to the auctioneer in a sealed envelope
- The item is sold to the bidder with highest bid at price p equal to the second highest bid
- The utility of bidder i is $u_i - p$ if he gets the item, 0 otherwise

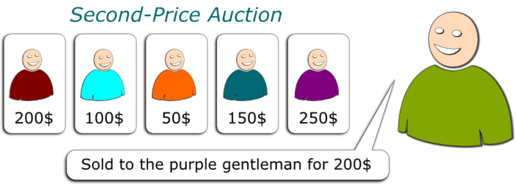

Equilibria in Second Price Auction

- Second price auction has a dominant strategy equilibrium for each agent: bid the true value $b_i = v_i$
- A similar auction is called Dominant strategy incentive compatible
- Bidding higher than v_i can lead to buy at price higher than valuation
- Bidding lower than v_i can lead to loose the item when it is sold at price lower than v_i
- The mechanism can be generalised to many other auction settings [Vickrey, Clarke, Groves, 1973]

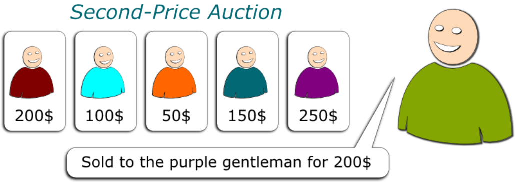

Outline

Conclusions of the Introduction...

- Economic decisions are taken more and more often by algorithms
- There are several barriers to the reach of good equilibria between agents:
  - computational complexity
  - coordination between agents
  - selfish behaviour
- In the last two decades Economics and Computer Science have made huge progresses in modelling and quantifying these phenomena

Coming next

- I. Algorithmic Mechanism Design for Two-sided Markets
- II. Algorithms for Online Labour marketplaces

I. Algorithmic Mechanism Design for Two-sided Markets

Based on joint work with Riccardo Colini Baldeschi (Facebook), Paul Goldberg (Oxford), Bart de Keijzer (King's College), Tim Roughgarden (Columbia), Stefano Turchetta (Twente & NTT DATA)

Outline

1. Part I: Mechanism Design in Two-sided Markets
2. Part II: Bilateral Trade
3. Part III: Two-sided Auctions
4. Conclusions

One-sided vs Two-sided Markets

One-sided markets:

- Adword auctions
- Ebay auctions

Two-sided market:

- Ad Exchange for display ads
- Online labor marketplaces
- Sharing economy (Uber, Airbnb, Lift, ..)
- Electricity market
- Stock exchange

Two-sided auctions

- Selling display ads is an example of a two-sided market
- Need to provide incentives to both buyers/advertisers and sellers/publishers that act strategically

Mechanism Design for One-sided Markets

Suppose we have k items and n interested buyers. We want to sell the items by interacting with the buyers.

- Each buyer $i \in [n] = \{1, \ldots, n\}$ holds a private valuation $v_i \in \mathbb{R}_{\geq 0}$ with distribution $F_i(v_i) = \int_{x \leq v_i} f_i(x) dx$.
- *quasi-linear* utility model:
  - $x_i \in \{0, 1\}$ indicates whether buyer i gets the item.
  - p_i is the price that buyer i pays to the mechanism.
  - The utility $u_i(\mathbf{x}, \mathbf{p})$ is then $x_i v_i - p_i$.
- Buyers behave rationally.

Mechanism Design

Q: How to maximize social welfare with Incentive Compatible mechanisms?

$$SW = \sum_{i \in [n]} x_i v_i$$

- Ensure that we sell the item to the k buyers with highest valuation!
- The *Vickrey auction* does it
- Buyers *submit their bids*: Direct Revelation Mechanism
- The Vickrey auction charges a price equal to the $k+1$-th highest bid.
- The Vickrey auction is Incentive Compatible (IC)

Revenue maximization

Bayesian setting is relevant:

- Known valuation distribution F_i of bidder
- Offer *monopoly price*:

$$r_i = argmax_p[p(1 - F_i(p))].$$

Second price auction with reserve price is optimal [Myerson, 1981]

- 1 item, 1 bidder $U[0,1]$, $r = 1/2$
- 1 item, 2 bidders $U[0,1]$:
  - second price auction with reserve price $1/2$ achieves revenue $5/12 > 1/3$
  - second price auction without reserve price achieves revenue $1/3$

One-sided vs Two-sided Markets

- In a one-sided market, the mechanism itself sells the item(s).
- In a two-sided market, the items are "sold" to the buyers by strategic agents called *sellers*.
- Mechanism is external entity and decides on the buyers and sellers who trade, and at which price.

A Standard Two-Sided Market Setting (1/2)

Double auctions

- There are k sellers, each with an identical copy of a single good for sale.
- There are n buyers, each interested only in receiving one copy of the good.
- w_j: the valuation of seller j, drawn from distribution G_j.
- v_i: the valuation of buyer i, drawn from F_i.

A Standard Two-Sided Market Setting (2/2)

An outcome consists of

- buyer allocation vector $\mathbf{x}^B \in \{0,1\}^n$
- seller allocation vector $\mathbf{x}^S \in \{0,1\}^k$
- buyer payment vector $\mathbf{p}^B \in \mathbb{R}^n$
- seller payment vector $\mathbf{p}^S \in \mathbb{R}^k$.

Negative payment means receiving money.

The utility model is symmetric for buyers and sellers:

- Buyer i's utility is $x_i^B v_i - p_i^B$.
- Seller j's utility is $x_j^S w_j - p_j^S$.

Ideal goals

- Maximize Social Welfare

$$SW = \sum_{i \in [n]} x_i^B v_i + \sum_{j \in [k]} x_j^S w_j$$

- Individual Rationality (IR), no agent gets negative utility
- Incentive Compatibility (IC) on the buyer and on the seller side
- We want our double auction to be *Budget Balanced (BB)*:

$$\sum_{i \in [n]} p_i^B + \sum_{j \in [k]} p_j^S = 0.$$

- *Weak Budget Balanced (BB)*: $\sum_{i \in [n]} p_i^B + \sum_{j \in [k]} p_j^S \geq 0.$
- The mechanism cannot subsidize the market (WBB) or make a surplus (BB)

Myerson and Satterthwaite impossibility results

Maximize Social Welfare is not possible with an (B)IC, IR, (W)BB mechanism
[Myerson and Satterthwaite, 1983]

- The results holds even for only one buyer and one seller with known distributions
- The *Second best BIC optimal mechanism provided in [MS83] is extremely complex and it does not have a closed form*
- *There is no guarantee on the Social Welfare that can be obtained by the mechanism*

Approximately optimal mechanisms

Seek for meaningful trade-offs between the IC, IR and BB requirements.

- Double auction mechanisms proposed in literature are either:
  - not IC
  - not BB
  - or do not have a good social welfare
- Many "large market" IR, IC, WBB results.

[McAfee 92]
[Dütting, Talgam-Cohen, Roughgarden, 2014]
[Blumrosen, Dobzinski, 2015]
[Segal-Halevi et al, 2016]

Trade-reduction Mechanism [McAfee 92]

- Order the buyers in decreasing order and the sellers in increasing order and find the breakeven index l.
- The first $l-1$ sellers give the item and receive w_l from the auctioneer;
- The first $l-1$ buyers receive the item and pay v_l to the auctioneer.

The mechanism is IC, WBB and achieve a $1-1/l$ approximation of the optimal social welfare.

Outline

1. Part I: Mechanism Design in Two-sided Markets
2. Part II: Bilateral Trade
3. Part III: Double Auctions
4. Conclusions

The Bilateral Trade Problem, $n = 1, k = 1$

- The double auction problem for one buyer with valuation v drawn from F and one seller with valuation w drawn from G.
- A trade is possible if $w \leq v$. Optimum social welfare:

$$OPT = \mathtt{E}_G[w] + \mathtt{E}_{F,G}[v - w | w \leq v]\mathtt{Pb}[w \leq v]$$
$$= \mathtt{E}[Seller\ \ value] + \mathtt{E}[Gain\ \ from\ \ trade]$$

- Every (DS)IC, BB mechanism is a posted price mechanism [Colini-Baldeschi, de Keijzer, Leonardi and Turchetta, 2016]
- How do we choose p in order to maximize

$$ALG = \mathtt{E}_G[w] + \mathtt{E}_{F,G}[v - w | w \leq p \leq v]\mathtt{Pb}[w \leq p \leq v]$$

- Set $p = m_G$, *median of the seller distribution* [McAfee 08]

The Bilateral Trade Problem

McAfee algorithm is a 2-apx of the social welfare [Blumrosen, Dobzinski, '15]:

- $$\begin{aligned} OPT &= \mathrm{E}_G[w] + \mathrm{E}_{F,G}[v-w|w \le p \le v]\mathrm{Pb}[w \le p \le v] \\ &+ \mathrm{E}_{F,G}[v-w|w \le v \le p]\mathrm{Pb}[w \le v \le p] \\ &+ \mathrm{E}_{F,G}[v-w|p \le w \le v]\mathrm{Pb}[p \le w \le v] \\ &\le 2 \times \mathrm{E}_G[w] + 2 \times \mathrm{E}_{F,G}[v-w|w \le p \le v]\mathrm{Pb}[w \le p \le v] \\ &= 2 \times ALG, \end{aligned}$$

  since $\mathrm{Pb}[w \le p] = \mathrm{Pb}[w \ge p] = 1/2$
- No deterministic algorithm which only depends on the seller distribution can improve
- A lower bound 1.33 and an upper bound 1.92 proved in [Colini-Baldeschi, de Keijzer, Leonardi and Turchetta, 2016]

The $e/e-1$-apx randomized mechanism for bilateral trade

Randomized $(e/e-1) = 1.58$-apx that depends only on the seller distribution [Blumrosen, Dobzinski, '16]

Random Quantile mechanism

Let $q(\cdot)$ be the quantile function of the seller, i.e., $G(q(x)) = x$. Post a price chosen randomly to both players as follows:

- Choose a number $x \in [1/e, 1]$ according to the cumulative distribution $D(x) = \ln(ex)$.
- Set the price to be $q(x)$.

- No quantile mechanism that uses only the seller distribution can achieve a better approximation
- A more involved mechanism achieves an $e/(e-1) - 0.0001$ approximation.

[Kang and Vondrak 2018]

Proof of the random quantile mechanism

- Assume the buyer has deterministic valuation b.
- The seller has value at least b with pb $1 - y$.
- Seller accepts price $q(x)$ with pb x.
- Density of price $q(x)$ is $d(x) = 1/x$.

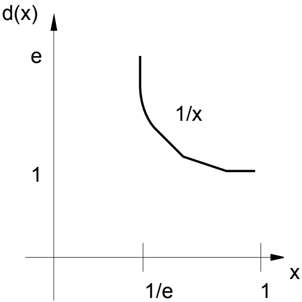

Proof of the random quantile mechanism

For a price $q(x)$, $x \in [1/e, y]$, trade occurs with probability x, and the realised efficiency is b:

$$QUANT(G, b) \geq \int_{1/e}^{y} x \cdot b \cdot \frac{1}{x} dx + b(1-y) \quad (1)$$

$$= b\left(y - \frac{1}{e}\right) + b(1-y) \quad (2)$$

$$= b\left(1 - \frac{1}{e}\right) \quad (3)$$

Algorithm OneSample

How many samples do we need if the distribution is unknown?

Algorithm OneSample

1. Sample p from seller's distribution;
2. Post price p and allow the agents to trade.

Theorem

The algorithm OneSample *provides a 2 approximation of the expected maximal welfare.*

[Dütting, Fusco, Lazos, Leonardi 2019]

Algorithm SampleQuantile

The SampleQuantile Algorithm has parameters $n \geq 0, 1/e > \delta > 0$:

1. Sample $z \in [1/e, 1]$ with CDF $\ln(e \cdot x)$.
2. Draw n samples from G.
3. Sort the samples in increasing order and choose the $(z - \frac{\delta}{2e}) \cdot n$-th one. Call that sample p.
4. Post price p and allow the agents to trade.

Theorem

For every $\varepsilon \in (0, \frac{4}{e})$, given $n = \frac{16e^2}{\varepsilon^2}\log(\frac{4}{\varepsilon})$ samples, SampleQuantile *provides an $\left(1 - \frac{1}{e} - \varepsilon\right)$ approximation of the optimal expected social welfare*

[Dütting, Fusco, Lazos, Leonardi 2019]

Outline

Sequential Posted Price Mechanisms

Definition

Sequential posted price (SPP) mechanisms offer *one* take-it-or-leave-it price to each buyer according to some order until all the items are sold.

Why do we study SPP mechanisms?

- Very popular mechanisms in practice
- Conceptually simple.
- Not direct revelation mechanisms
- Buyers have obvious dominant strategies
- They are easy to analyze
- Seemingly needed for DSIC, BB double auction.

Drawback: Require prior information about buyer and seller valuations

One-sided SPP mechanisms

- There is an auctioneer with k identical items to sell.
- There are n buyers. They want no more than 1 item.
- For buyer i, valuation v_i is drawn from a finite distribution $F_i \in \mathbb{R}_{\geq 0}$.

How well can SPP mechanisms approximate SW and revenue?

For social welfare the optimal mechanism is VCG

For revenue the optimal mechanism is Myerson

SPP Mechanism [Chawla et al. (2010)]

- For buyer i, let $q_i := \mathbf{Pr}[\text{Optimal mechanism gives item to buyer } i]$.
- Let $\bar{p}_i$ be such that $\mathbf{Pr}_{v_i \sim F_i}[v_i > \bar{p}_i] = q_i$.
- The SPP with prices $\bar{p} = (\bar{p}_1, \ldots, \bar{p}_n)$, offered in non-increasing order, 2-approximates revenue or social welfare of optimal mechanism.

Adapting SPP Mechanisms for Two-Sided Markets (1/2)

SPP mechanisms are adapted to two-sided markets:

1. Decide on an order σ of the buyers.
2. Decide on an order λ of the sellers.
3. Decide on prices p_{ij} for all $i \in [n]$, $j \in [k]$.
4. Iteratively offer the price p_{ij} to the next buyer-seller pair (i, j) according to σ and λ.
   - If both accept, let them trade at price p_{ij}. Allocate an item to i. Deallocate an item from j. Charge p_{ij} to i and $-p_{ij}$ to j. Move to the next seller of λ. Move to the next buyer of σ.
   - If seller rejects, move to the next seller in λ.
   - If buyer rejects, move to the next buyer in σ.

Adapting SPP Mechanisms for Two-Sided Markets (2/2)

Things to note about two-sided SPP mechanisms:

- Inherently BB.
- Behaving "truthfully" is not always a dominant strategy. However:

Lemma

If prices only depend on the buyer, and not on the seller (i.e., $p_{ij} = p_{ij'}$ for all $i \in [n], j, j' \in [k]$) and are posted in a non-increasing order, then "truthfulness" is a dominant strategy.

Approximation result for double auctions

Theorem

There exists a BB double auction with a dominant strategy that 6-approximates the expected optimal social welfare (even with an additional matroid constraint on the set of buyers that trade).

[Colini-Baldeschi, de Keijzer, Goldberg, Leonardi, Roughgarden, and Turchetta, 2016]

Outline of a simpler mechanism

How this mechanism works:

- For $i \in [n]$, let $\bar{p}_i$ denote the price from the single-sided mechanism.
- Let σ denote the order of the buyers by decreasing $\bar{p}_i$ (also according to Chawla et al. (2010)).
- Let λ be a uniform random ordering of the sellers.
- Set $p_{ij} = p_i = \max\{\bar{p}_i, m_{(k/2)}\}$ where $m_{(k/2)}$ is the median of the sellers' median valuations.

Analysis of the mechanism (1/2)

- Let at most $k/4$ pairs trade.
- This leaves $3k/4$ sellers with their item.
- The sellers prepared to trade are the sellers with the lowest valuations.
- So: $(4/3)\mathbf{ALG}_s \geq \mathbf{OPT}_s$.

Analysis of the mechanism (2/2)

Now the buyers' side.

- By charging at least $m_{(k/2)}$, we expect at least half of the sellers are prepared to trade.
- This implies: with probability at least $1/2$, at least $k/4$ sellers are prepared to trade.
- In case $p_i = \bar{p}_i$ for all buyers in σ. We get $\mathbf{ALG}_b \geq (1/2)(1/4)(1/2)\mathbf{OPT}_b$.
- In the case $p_i = m_{(k/2)}$ for a subset of the buyers, some social welfare on the buyers' side may be lost.
- In that case we show that there are corresponding sellers with a higher valuation.
- $(4/3)\mathbf{ALG}_s + 16\mathbf{ALG}_b \geq \mathbf{OPT}_b$

Together:

$16\mathbf{ALG} \geq (4/3)\mathbf{ALG}_s + (4/3)\mathbf{ALG}_s + 16\mathbf{ALG}_b \geq \mathbf{OPT}_b + \mathbf{OPT}_s = \mathbf{OPT}$

Outline

1. Part I: Mechanism Design in Two-sided Markets
2. Part II: Bilateral Trade
3. Part III: Two-sided Auctions
   - Two-sided Double Auctions
   - Two-sided Combinatorial Auctions
4. Conclusions

Two-sided combinatorial auctions

- Every v_i and every w_j map from $2^{[k]}$ to $\mathbb{R}_{\geq 0}$
- We will consider probability distributions over the following classes of valuation functions:
  - v is *additive* iff $v(S) = \sum_{j \in S} \alpha_j v(\{j\})$ for all $S \subseteq [k]$ for some real numbers α_j.
  - v is *fractionally subadditive (or XOS)* if and only if there exists a collection of additive functions $a_1, \ldots, a_d$ such that for every bundle $S \subseteq [k]$ it holds that $v(S) = \max_{i \in [d]} a_i(S)$
  - Fractionally subadditive (or XOS) generalizes submodular functions

The mechanism for two-sided combinatorial auctions

- For each item $j \in [k]$, let $SW_j^B(\boldsymbol{v})$ its expected contribution to the social welfare.
- Set $p_j := \frac{1}{2}\mathbb{E}_{\boldsymbol{v}}\left[SW_j^B(\boldsymbol{v})\right]$.
- For all $j \in [k]$:
  1. Set $q_j := 1/(2Pr[w_j \leq p_j])$.
  2. With probability q_j, offer payment p_j in exchange for her item. Otherwise, skip this seller.
  3. If j accepts the offer, set $\Lambda_1 := \Lambda_1 \cup \{j\}$.
- For all $i \in [n]$:
  1. Let $D(v_i, \boldsymbol{p}, \Lambda_i)$ be the demand set of buyer i at price p_j.
  2. Buyer i chooses a bundle $B_i \in D(v_i, \boldsymbol{p}, \Lambda_i)$.
  3. Allocate the accepted items to buyer i
  4. $\Lambda_{i+1} := \Lambda_i \setminus B_i$.

Results

- A 6-approximate DSIC mechanism for buyers with XOS-valuations and sellers with one item at their disposal (i.e., *unit-supply sellers*);
- a 6-approximate BIC mechanism for buyers with XOS-valuations and non-unit supply sellers with additive valuations;
- a 6-approximate DSIC mechanism for buyers with additive valuations and sellers with additive valuations.

[Colini-Baldeschi, de Keijzer, Goldberg, Leonardi, Roughgarden, Turchetta, 2017]

Outline

Conclusions on Two-sided Market Design

- Algorithmic mechanism design in two-sided markets finds many relevant applications to digital markets
- Simple mechanisms achieve good efficiency while obeying the IR, IC, BB requirements
- Many open problems and applications to digital markets

Conclusions of the first part.

- Economic decisions are taken more and more often by algorithms
- There are several barriers to the reach of good equilibria between agents:
  - computational complexity
  - coordination between agents
  - selfish behaviour
- In the last two decades Economics and Computer Science have made huge progresses in modelling and quantifying these phenomena

Conclusions

Many topics have not been touched in this talk:

- Repeated games
- Mechanism design for social good
- Social choice and voting
- Behavioural cues, e.g., altruistic or myopic behaviour
- Complex market structures
- Many applications to the modelling of social systems and biological evolution

Coming next

- II. Algorithms for Online Labour marketplaces