Lecture 19

Markets, Mechanisms and Machines

David Evans and Denis Nekipelov

Mechanism design

- Mechanism design: theory for “rules of interaction” where *selfish behavior leads* to *good outcomes*.
- Selfish behavior: each agent maximizes her own utility (rational behavior)
- Leads: equilibrium (once actions are in equilibrium no one has incentive to deviate).
- Good outcomes: goals of the designer (social surplus or welfare, revenue of the auctioneer)

Mechanism design

Principles of mechanism design theory

- **Informative**: pinpoints salient features of environment and characteristics of good mechanisms
- **Prescriptive**: gives concrete suggestions for design of good mechanisms
- **Predictive**: theory predictions should match reality
- **Tractable**: theory should not assume super-natural ability for the agents or designer to optimize.

Mechanism design

- In many environments, optimal mechanisms do not agree with these principles
  - Complex product spaces and preferences (e.g. combinatorial auctions)
  - Complex information exchange requirements between agents
  - Complex structure of beliefs required to implement Bayes-Nash equilibrium

Mechanism design

- Complexity is one of the main obstacle of mechanism design with truly optimal mechanisms
- In fact, even for simple bimatrix games (games of complete information) computing Nash or mixed Nash equilibrium is difficult

Complexity of computing Nash equilibria

- Nash (1950): all games have a mixed Nash equilibrium
  - Exist distributions over players' actions such that each is best response to everyone else's actions
- **Theorem**. (Brouwer Fixpoint Theorem). If C is bounded, convex and closed, and $f : \mathrm{C} \longmapsto \mathrm{C}$ is continuous, there exists x s.t. $f(x) = x$.

Complexity of computing Nash equilibria

- n is number of players and S_i actions space of player i, and Δ_i be set of probability distributions over actions of player i, i.e.

  $\Delta_i = \{(p_s : s \in S_i) \mid p_s \geq 0\ 0 \text{ and } \Sigma_{s \in Si} p_s = 1\}$
- By C to denote the set of the mixed strategies of all the players, i.e. $C = \Delta_1 \times \ldots \times \Delta_n$
- C is convex, bounded and closed.
- We need a function $f : C \longmapsto C$ that the NE is fixpoint
- Natural answer is to use the best response.

Complexity of computing Nash equilibria

- *G*iven $\boldsymbol{p} = (p_1, p_2,\ldots, p_n) \in \mathrm{C}$, where $p_i \in \Delta_i$
- q_i is best response of player i
- Define function as $f(p) = (q_1,\ldots, q_n)$.
- Then can try to use fixed point argument
- Fundamental issue is that f may not be a function since the best response for the player might not be unique
- If we try to fix it somehow, resulting function may not be continuous

Complexity of computing Nash equilibria

- Consider $\max_q u_i(q,p_{-i})-||p_i-q||^2$
- p_i is the best response of player i
- For player i maximize penalized utility
- Suppose the maximizer for player i is q_i, define $f(p) = (q_1,\ldots, q_n)$
- **Lemma**. $\max_q u_i(q,p_{-i})-||p_i-q||^2$ has a unique maximum
- If a class of optimization problems has unique optimum then optimum is continuous function of coefficients in the objective function
- If $f(p)=p$ then p is Nash

Complexity of computing Nash equilibria

- Need “tractable” equilibrium concepts
- Desiderata:
  - Universality
  - Naturality and credibility
  - Efficiently computable
- Focus on last one: if computing equilibria is intractable, it is unlikely that mechanism designer can easily implement it

Complexity of computing Nash equilibria

- Is there an efficient algorithm for computing a mixed Nash equilibrium?
- For zero-sum games von Neumann work shows that Nash equilibrium can be characterize as solution of linear program
  - Using ellipsoid method we already showed that solutions to linear programs can be computed efficiently
- Non-zero-sum games do not reduce to linear programs
- Proposed algorithms are either of unknown complexity, or known to require exponential time.

Complexity of computing Nash equilibria

- Is there an efficient algorithm for computing a mixed Nash equilibrium?
- For zero-sum games von Neumann work shows that Nash equilibrium can be characterize as solution of linear program
  - Using ellipsoid method we already showed that solutions to linear programs can be computed efficiently
- Non-zero-sum games do not reduce to linear programs
- Proposed algorithms are either of unknown complexity, or known to require exponential time.

Complexity of computing Nash equilibria

- Recall from last lecture that NP-complete problems are those that cannot be efficiently solved unless P=NP
- Example: Boolean satisfiability problem (SAT)
  - Determine if there exists an interpretation (assignment to TRUE or FALSE) that satisfies a given Boolean formula
  - For given Boolean formula replace inputs with TRUE or FALSE that entire formula = TRUE
  - If this replacement is possible, the problem is called satisfiable

Complexity of computing Nash equilibria

- NP-complete problems like SAT are typically shown to be intractable from possibility that solution might not exist
- This argument is at the core of NP-completeness proof
- Unlike any known NP-complete problem, solution to problem of computing Nash equilibrium always exists (Nash's theorem)
- This indicates that while computing mixed Nash is not P, it is also not NP-complete

Complexity of computing Nash equilibria

- Suppose there is reduction from SAT to Nash: efficient algorithm that takes as input instance of SAT and outputs instance of Nash
- Then if we provide solution to instance of Nash, we could tell if SAT has solution
- We turn this into nondeterministic algorithm to verify if instance of SAT has solution
  - Guess solution of Nash instance, and check that it indeed implies that SAT instance has no solution
- Existence of such on-deterministic algorithm for SAT would be similar to establishing P=NP

Complexity of computing Nash equilibria

- Papadimitriou (1994) considers a class of (seemingly) unrelated "search" problems
- Given an input, find a solution (which then can be easily checked) or report that none exists
- Note asymmetry between these outcomes: "none exists" is not required to be easy to verify.
- Search problem is total if the solution always exists
- Can describe a specific subset of total search problems

Complexity of computing Nash equilibria

- Consider directed graph
- Vertex in directed graph is “unbalanced” if number of its incoming edges differs from number of its outgoing edges
- For each directed graph and unbalanced vertex there must exist at least one other unbalanced vertex
- Problem:
  - **Input**: directed graph G and a specified unbalanced vertex of G.
  - **Output**: Some other unbalanced vertex.

Complexity of computing Nash equilibria

- Such problems are called PPAD (polynomial parity argument for directed graphs)

Theorem (Daskalakis, Goldberg, Papadimitriou, 2008). The problem of computing mixed Nash equilibria is PPAD-complete

- Nash is hard if P≠NP
- Existing algorithms seem to confirm it

Approximation

- When optimal mechanism is not easily available, good mechanisms can be generated by approximations

Definition. For an environment given implicitly, denote an approximation mechanism and its performance by APX, and a reference mechanism and its performance by REF.

(i)For any environment, APX is a β approximation to REF if

$\text{APX} \geq \text{REF}/\beta$

(ii) For any class of environments, a class of mechanisms is a β approximation to REF if for any environment in the class there is a mechanism APX in the class that is a β approximation to REF.

(iii) For any class of environments, a mechanism APX is a β approximation to REF if for any environment in the class APX is a β approximation to REF.

Example: posted price mechanism

For a given price p uniform pricing mechanism serves a single item to the first agent willing to pay p

Theorem. If the values of agents are independently drawn from the distribution F, uniform pricing mechanism with price $p = F^{-1}(1 - 1/n)$ is the $e/(e - 1)$ approximation to the optimal social surplus

Recall: optimal social surplus corresponds to allocating item to the highest-value agent

Example: posted price mechanism

- Take second-price auction as REF and uniform pricing as APX
- REF optimizes surplus subject to ex post supply constraint (only 1 item is available), i.e. allocates each agent with ex ante probability $1/n$
- Consider mechanism UB that maximizes social surplus subject to constraint that each agent is allocated with ex ante probability $1/n$ but does not have ex post supply constraint
- Note that UB $\geq$ REF

Example: posted price mechanism

- Since UB has no supply constraint, we can optimize it for each agent separately
- Socially optimal way of allocating to a single agent with ex ante probability $1/n$ is to offer price
  $p = F^{-1}(1 - 1/n)$
- In this case $\Pr(v > p) = 1/n$ (i.e. agent is allocated)
- Social surplus of UB: $\text{UB} = n\, \mathrm{E}[v \mid v > p] \Pr(v > p)$
  $= \mathrm{E}[v \mid v > p]$
- Now we relate UB to the social surplus of REF

Example: posted price mechanism

- REF can allocate to only one agent
- The agent is allocated only if her value exceeds the uniform price $p = F^{-1}(1 - 1/n)$. The probability of this is $1/n$
- The item is not allocated if all values are below p
- This can happen with probability $(1-1/n)^n < 1/e$, i.e. the probability that the item is allocated is $>1 - 1/e$
- Expected surplus of APX is then

$$\text{APX} \geq (1 - 1/e)\ \text{E}[v \mid v > p] = (1 - 1/e)\text{UB} \geq (1 - 1/e)\text{REF}$$

Bayesian approximation

- Focus on auction environment with asymmetric value distributions
- I.e. the second price auction with a reserve is no longer optimal
- We will only consider distributions with monotone hazard rate (recall the terminology from Myerson's optimal auction)
- Our goal is to find simple compelling approximations for optimal auction in the asymmetric environment

Bayesian approximation

- In symmetric settings Myerson's optimal auction maximizes the revenue of the auctioneer by maximizing the virtual surplus
- Each agent's virtual value is characterized by
  $$V(v) = v - (1 - F(v)) / f(v)$$
- Reserve prices are set by discarding all agents with negative virtual values
- We want to determine if a version of the optimal auction remains approximately optimal in the asymmetric settings

Bayesian approximation

- Consider generalization of the second price environment:

The second-price auction with (discriminatory) reserves $\boldsymbol{p} = (p_1,\ldots,p_n)$ is:

I. reject each agent i with $v_i < p_i$,
II. allocate the item to the highest valued agent remaining or none if none exists), and
III. charge the winner her critical price.

- Question: Can this (simple) design provide a good approximation for optimal auction?

Bayesian approximation

Theorem. For single-item environments and agents with values drawn independently from (non-identical) regular distributions, the second-price auction with (asymmetric) monopoly reserve prices obtains at least half the revenue of the (asymmetric) optimal auction.

Bayesian approximation

Proof:

- We proved that for regular distributions (MHR) expected revenue is equal to expected virtual surplus

Lemma. For any virtual value function, the virtual values corresponding to values that exceed the monopoly price are non-negative.

Lemma. For any distribution, the value of an agent is at least her virtual value for revenue.

- Both results follow from MHR and the structure of virtual value $V(v) = v - (1 - F(v)) / f(v)$

Bayesian approximation

Proof:

•Let REF denote the optimal auction and its expected revenue and APX denote the second price auction with monopoly reserves and its expected revenue

•Let I be the winner of the optimal auction and J be the winner of the monopoly reserves auction

•Then REF=E[$V_I(v_I)$] and APX=E[$V_J(v_J)$]

•By law of total probability

$$\begin{aligned} \text{REF} = \text{E}[V_I(v_I)] &= \text{E}[V_I(v_I) \mid I = J]\Pr(I = J) \quad \text{(a)} \\ &+ \text{E}[V_I(v_I) \mid I \neq J]\Pr(I \neq J) \quad \text{(b)} \end{aligned}$$

Bayesian approximation

Proof:

- Part (a):

$$\mathrm{E}[V_I(v_I) \mid I = J]\Pr(I = J) = \mathrm{E}[V_J(v_J) \mid I = J]\Pr(I = J)$$

$$\leq \mathrm{E}[V_J(v_J) \mid I = J]\Pr(I = J) + \mathrm{E}[V_J(v_J) \mid I \neq J]\Pr(I \neq J)$$

$$= \mathrm{E}[V_J(v_J)] = \mathrm{APX}$$

Bayesian approximation

Proof:

•Part (b):

$E[V_I(v_I) \mid I \neq J]\Pr(I \neq J) \leq E[\, v_I \mid I \neq J]\Pr(I \neq J)$

by the property of virtual values;

•Given that J is the winner of APX (second price auction), her payment is at least v_I

Thus

$E[\, v_I \mid I \neq J]\Pr(I \neq J)$

$$\leq E[\text{ Payment}_J(v_J) \mid I \neq J]\Pr(I \neq J)$$

Bayesian approximation

Proof:

•Part (b):

Given that payments are non-negative

$$E[\text{ Payment}_J(v_J) | I \neq J]\Pr(I \neq J)$$

$$\leq E[\text{ Payment}_J(v_J) | I \neq J]\Pr(I \neq J)$$

$$+ E[\text{ Payment}_J(v_J) | I = J]\Pr(I = J) = \text{APX}$$

Therefore

$$E[V_I(v_I) \mid I \neq J]\Pr(I \neq J) \leq \text{APX}$$

Bayesian approximation

Proof:

- Collect terms:
  - Part (a): $E[V_I(v_I) \mid I = J]\Pr(I = J) \leq \text{APX}$
  - Part (b): $E[V_I(v_I) \mid I \neq J]\Pr(I \neq J) \leq \text{APX}$
- Therefore

$$\text{REF} = E[V_I(v_I) \mid I = J]\Pr(I = J) + E[V_I(v_I) \mid I \neq J]\Pr(I \neq J) \leq 2\ \text{APX}$$

- This means that APX (second price auction with monopoly reserves) produces at least half of the optimal revenue

Posted prices

- Disadvantages of auctions
  - require multiple rounds of communication (can be slow)
  - require all agents to be present at the time of the auction
- In many environments these features are prohibitive: routing, online and offline retail
- Posted pricing does not have these disadvantages and provides strong revenue guarantees
  - No room for collusion
  - Can be used to set starting prices if auctions are subsequently used

Posted prices

- Consider oblivious posted prices (agents arrive and face their prices in arbitrary order)
- Theory is based on *prophet inequality* from optimal stopping theory
  - Gambler faces sequence of n games
  - Game i has prize v_i as independent draw from F_i
  - Order of the games and price distributions known to gambler
  - In game i gambler observes prize $v_i \sim F_i$ and must decide whether to keep the prize and stop or return the prize and continue
  - Only allowed to keep one prize

Posted prices

- What is the optimal stopping rule for the gambler?
  - Use backwards induction: in game n gambler stops regardless of prize realization
  - Expected value from stopping in game n is $E[v_n]$
  - Then in game $n - 1$ the gambler stops if v_{n-1} is greater than $p_{n-1}=E[v_n]$
  - Expected value from stopping in game $n - 1$ is

    $$p_{n-2}=E[v_{n-1} \mid v_{n-1}>p_{n-1}](1 - F_{n-1}(p_{n-1}))+p_{n-1}F_{n-1}(p_{n-1})$$
  - By the same principle, expected value from stopping in game $n - 2$ is

    $$p_{n-3}=E[v_{n-2} \mid v_{n-2}>p_{n-2}](1 - F_{n-2}(p_{n-2}))+p_{n-2}F_{n-2}(p_{n-2})$$

Posted prices

- This leads to sequence of thresholds $p_1,\ldots,p_n$ defining optimal stopping rule for gambler
- Has typical drawbacks of optimal strategies
  - Complicated (takes n numbers to describe it)
  - sensitive to small changes of gamble (e.g. order of games)
  - Little room for intuitive understanding of properties of good strategies.
  - Does not generalize well to give solutions to similar gambles
- May be attractive to look at simple approximations instead

Posted prices

- May be attractive to look at simple approximations instead
- Uniform threshold strategy is given by single threshold p and requires gambler to accept first prize i with $v_i \geq p$
- Threshold strategies are suboptimal
  - E.g. prescribes not to stop at game n if $v_n < p$
- Call prize selection procedure when multiple prizes are above p *tie-breaking rule*
- For gambler's gambler's game it is lexicographic (smallest i)

Prophet inequality

Theorem. For any product distribution on prize values $\boldsymbol{F} = F_1 \times \ldots \times F_n$, there exists a uniform threshold strategy such that the expected prize of the gambler is at least half the expected value of the maximum prize; moreover, the bound is invariant with respect to the tie-breaking rule; moreover, for continuous distributions with non-negative support one such threshold strategy is the one where the probability that the gambler receives no prize is exactly 1/2.

Prophet inequality

- Discussion
  - Even though gambler does not know realizations of the prizes in advance she can still do half as well as a prophet who does.
  - This result implies that optimal (backwards induction) strategy has this performance guarantee
  - However, such guarantee was not obvious from original formulation of optimal strategy
  - Unlike backwards induction, it is very simple
  - Result driven by trading off probability of not stopping and receiving no prize with the probability of stopping early with a suboptimal prize

Prophet inequality

Proof

- Let REF denote prophet and her expected prize (E[$\max_i v_i$]) and APX denote gambler with strategy p and her expected price
- Define q_i=1-$F_i(p)$=Pr($v_i \geq p$) the probability that prize i is above threshold p and $\chi=\Pi_i(1-q_i)$ is the probability that gambler rejects all prices
- We allow prophet not to pick any prizes is all their values are negative
- Use notation $(x)^+=\max\{x,0\}$

Prophet inequality

Proof

- Bound expected prize of the prophet from above
- REF=E[$\max_i v_i$]=p+E[$\max_i (v_i-p)$]

  $\leq p$+E[$\max_i (v_i-p)^+$]

  $\leq p$+Σ_iE[$(v_i-p)^+$]

Prophet inequality

Proof

- Bound expected prize of the gambler from below
- Suppose that gambler receives prize g
- Split value of the prize into p and g - p (guaranteed part and "surplus")
- Expected value of the prize splits into APX_1 and APX_2
- Then $APX_1 = p$ Pr(gambler gets a prize) = $(1-\chi)p$
- To evaluate APX_2, denote by E_i the event that all prizes excluding i^{th} are below p

Prophet inequality

Proof

- Bound expected prize of the gambler from below
- Then $\mathrm{APX}_2 \geq \Sigma_i \mathrm{E}[(v_i\text{-}p)^+ | \, E_i]\Pr(E_i) \geq \chi \, \Sigma_i \mathrm{E}[(v_i\text{-}p)^+]$
  - $\Pr(E_i) = \Pi_{i\neq j}(1\text{-}q_j) \geq (1\text{-}q_i) \, \Pi_{i\neq j}(1\text{-}q_j) = \chi$
  - And v_i is independent from E_i (can drop conditioning)
- $\mathrm{APX}=\mathrm{APX}_1+\mathrm{APX}_2 \geq (1\text{-}\,\chi)p + \chi \, \Sigma_i \mathrm{E}[(v_i\text{-}p)^+]$
- Plug χ =1/2: this corresponds to threshold p such that

  $\chi = \Pi_i(1\text{-}F_i(p))=1/2$
- Combining two inequalities produces

  $2\mathrm{APX} \geq p + \Sigma_i \mathrm{E}[(v_i\text{-}p)^+] \geq \mathrm{REF}$
- This proves the prophet inequality

Prophet inequality

- Prophet inequality is tight: better approximation bound cannot generally by obtained by uniform threshold strategy
- Invariance to the tie-breaking rule implies that prophet inequality also approximates settings similar to gambler's game
- In oblivious posted pricing agents arrive in worst-case order and the first agent who desires to buy the item at her offered price does so.
- Thus, can use prophet inequality to show that there exist oblivious posted pricings that guarantee half the optimal surplus

Posted prices

- Second price auction obtains optimal social surplus $\max_i v_i$
- Uniform posted price corresponds to uniform threshold for values
- In worst case arrival order agent with lowest value above posted price buys.
- This is just like the gambler's game with tie-breaking by smallest value v_i.
- Recall that prophet inequality is invariant w.r.t. to tie-breaking rules

Posted prices

Theorem. In single-item environments there is an anonymous pricing whose expected social surplus under any order of agent arrival is at least half of that of the optimal social surplus.

Posted prices

- Now consider revenue from posted price
- Revenue-optimal single-item auction selects winner with highest (positive) virtual value
- Note that the gambler’s problem (maximizing prize) is similar to the auctioneer’s problem (but maximizing virtual value)
- Uniform threshold for the gambler’s prize corresponds to uniform threshold for virtual values maximized by the auctioneer
- Note: uniform threshold for virtual values corresponds to non-uniform (a.k.a., discriminatory) prices.

Posted prices

- Definition: uniform virtual price π corresponds to uniform virtual pricing $\boldsymbol{p}=(p_1,\ldots,p_n)$ such that $V_i(p_i)=\pi$
- Compare uniform virtual pricing to gambler's game
  - Both use uniform threshold to select maximum
  - Uniform virtual pricing obtains worst revenue when agents arrive in order of increasing price (in value space).
  - Implicitly breaks ties by smallest posted price p_i.
  - Gambler's threshold strategy breaks ties by ordering of games (i.e., by smallest i).
- Irrespective of tie-breaking rule bound of prophet inequality holds

Posted prices

Theorem. In single-item environments there is a uniform virtual pricing (for virtual values equal to marginal revenues) whose expected revenue under any order of agent arrival is at least half of that of the optimal auction.

Posted prices

Proof

- Uniform virtual price π corresponds to uniform virtual pricing $\boldsymbol{p}=(p_1,\ldots,p_n)$.
- Worst case outcome of such posted pricing:
  - When there is only one agent i with value $v_i \geq p_i$ revenue is
  - When there are multiple agents S (values exceed offered prices) lowest price arrives first and pays $\min_{i \in S} p_i$.
- This is version of gambler's game with tie-breaking rule by smallest p_i
- Bound revenue of uniform virtual pricing with worst-case arrival order, relate its revenue to its virtual surplus.

Posted prices

Proof

- By Myerson's theorem can express optimal revenue as expected maximum virtual value
- Expected revenue of a uniform virtual pricing is equal to its expected virtual surplus.
- By prophet inequality, there is uniform virtual price that obtains a virtual surplus of at least ½ of maximum virtual value
- Thus, the revenue of the corresponding price posting is at least half the optimal revenue.