CS4261 Algorithmic Mechanism Design

Game → Players $N=\{1,2,\dots,n\}$
→ Actions
→ Preferences over outcomes
→ general framework for strategic interaction.

$BR_i(\vec{p}_{-i}) = \{q \in \Delta(A_i) : u_i \dots\}$

NE

Socially optimal: maximize the total benefit

envy free: no one prefers another bundle more than his own

Normal form game → set of players $N=\{1,2,\dots,n\}$
→ set of actions for each player A_i
→ action profile $\vec{a} \in A_1 \times A_2 \times \dots \times A_n$
→ utility of each $\vec{a}$ for each player $u_i(\vec{a})$.

2,1	1,2
1,2	2,1

Socially optimal: maximise $\sum_i u_i(o)$ ← outcome

Pareto optimal: $\nexists o'$ s.t. $(\underbrace{u_i(o') \geq u_i(o)\ \forall i \in N}_{\text{everybody won't be worse-off with swapping}})$

↑ outcome

and $(\underbrace{u_i(o') > u_i(o)\ \exists i \in N}_{\text{someone gets strictly better-off with swapping}})$

$u_1(T, q) = \dots$

$BR_1(q) = \begin{cases} T\ (p=1) & \text{if } q > \frac{1}{2} \\ B\ (p=0) & \text{if } q < \frac{1}{2} \\ p \in [0,1] & \text{if } q = \frac{1}{2} \end{cases}$

player 1 ↑ if player 2 plays L with prob q

Dominant strategy: A best option for a player, regardless of what other players choose (may or may not exist)

(Weak) Domination: $\vec{p} \in \Delta(A_i)$ dominates $\vec{q} \in \Delta(A_i)$: $\forall \vec{p}_{-i} \in \Delta(A_{-i})$: $u_i(\vec{p}_{-i}, \vec{p}) \geq u_i(\vec{p}_{-i}, \vec{q})$

(Strong) Domination: strict ineq. for all.

Best response set

$BR_i(\vec{a}_{-i}) := \{ b \in A_i \mid b \in \arg\max u_i(\vec{a}_{-i}, b) \}$

↑ the set of options that yield the best outcome for me, given what everyone else is going to play.

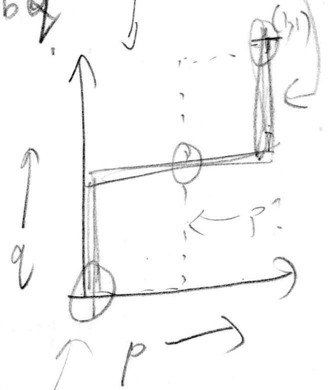

A	5,4	4,3
B	2,5	3,4

$A > B$ (dominates)

because $\{ 5 > 2,\ 4 > 3$.

Nash equilibrium: $\forall i \in N$, $a_i \in BR_i(\vec{a}_{-i})$

↑ i.e. it is in everyone's best response set.

intersections are Nash equilibriums.

Mixed strategies

- Player utility: expected value (we assume players are risk-neutral)
- Mixed Nash equilibrium: $\forall i \in N$: $\underbrace{\forall \vec{q}_i \in \Delta(A_i),\ u_i(\vec{p}) \geq u_i(\vec{p}_{-i}, \vec{q}_i)}$
  ↓ always exists!

  It is never better (in terms of utility/EV) to switch to a different mixed strategy.
- Best response → similarly defined.

Support of a distribution: set of options taken with nonzero probability.

5,4	4,4
5,5	4,5

↑ Nash equilibrium

Theorem: If action $a \in A_i$ is strictly dominated by some $\vec{p} \in \Delta(A_i)$ then a is never played (with any positive probability) in a Nash equi.

← so we can only remove strongly dominated strategies.

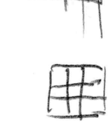 ← iterated removal of dominated strategies

A mixed strategy can dominate another strategy too!

$\}$ mixed > top row?

Zero-sum game:
→ 2-player game
→ $\forall a_i \in A_1,\ \forall b_j \in A_2 : \underbrace{u_1(a_i, b_j)}_{A_{ij}} = -u_2(a_i, b_j)$

→ Mixed Nash eqm is poly-time computible via simplex.

Duality of linear optimization (simplex) problem:

Primal	Dual
minimize $\vec{c}^T\vec{x}$	maximize $\vec{b}^T\vec{y}$
s.t. $A\vec{x} \geq \vec{b}$	s.t. $A^T\vec{y} \leq \vec{c}$
and $x_i \geq 0\ \forall i \in N$	$y_j \geq 0\ \forall j \in M$

Given A (matrix), $\vec{b}$, $\vec{c}$, the optimal for both problems are the same.

Theorems: If $\vec{x}^*$ and $\vec{y}^*$ are the optimal solutions, then:

1. Optimality: $\vec{c}^T\vec{x}^* = \vec{b}^T\vec{y}^*$
2. Complementary Slackness:
   If $x_i^* > 0$ then $(A^T\vec{y}^*)_i = c_i$
   If $y_j^* > 0$ then $(A\vec{x}^*)_j = b_j$

von Neumann Minimax thm: For any payoff A ← matrix,

$$\underbrace{\max_{\vec{p}\in\Delta(A_1)} \min_{\vec{q}\in\Delta(A_2)}}_{\text{maximin}} \underbrace{\vec{p}A\vec{q}^T}_{\text{get expected payoff for strategy}} =: V_+ = V_- := \underbrace{\min_{\vec{q}\in\Delta(A_2)} \max_{\vec{p}\in\Delta(A_1)}}_{\text{minimax}} \vec{p}A\vec{q}^T$$

To prove: 1. show that it is (at least) as good if player 2 uses a pure strategy, i.e.

$$\max_{\vec{p}\in\Delta(A_1)} \underbrace{\min_{\vec{q}\in\Delta(A_2)} \vec{p}A\vec{q}^T}_{\text{min over all mixed strategies}} = \max_{\vec{p}\in\Delta(A_1)} \underbrace{\min_j \left((\vec{p}A)_j\right)}_{\text{min over all pure strategies}}$$

2. Use LP duality:

$$\max V_+ \quad \text{s.t. } V_+ \leq \sum_{i=1}^{n} a_{ij} x_i \text{ for all } j \in M \text{ and } \sum_{i=1}^{n} x_i = 1$$

← all pure strategies

$$= \min V_- \quad \text{s.t. } V_- \geq \sum_{j=1}^{m} a_{ij} y_j \text{ for all } i \in N \text{ and } \sum_{j=1}^{m} y_j = 1$$

Support of a Nash eqm: support of $\vec{p}$: $\{a : p(a) > 0\}$ ← all those options that have +ve contribution to $\vec{p}$.

Thm: If $(\vec{p}, \vec{q})$ is a Nash eqm and $a \in supp(\vec{p})$ then:
$u_1(a, \vec{q}) \geq u_1(a', \vec{q})\ \forall a' \in A_1$

Solving Nash eqm for two players directly cannot be done with LP, but with thm it can be done:

Directly: Find $\vec{p}, \vec{q}$ s.t. $\sum_{a\in A_1} p(a) = 1$, $\sum_{b\in A_2} q(b) = 1$
$\forall \vec{p}' \in \Delta(A_1) : u_1(\vec{p}, \vec{q}) \geq u_1(\vec{p}', \vec{q})$
$\forall \vec{q}' \in \Delta(A_2) : u_2(\vec{p}, \vec{q}) \geq u_2(\vec{p}, \vec{q}')$
(there are infinitely many constraints)

With thm: $\forall B_1 \subseteq A_1, \forall B_2 \subseteq A_2$: Find $\vec{p}, \vec{q}$ s.t. $\sum_{a\in A_1} p(a) = 1$; $\sum_{b\in A_2} q(b) = 1$
$\forall a \notin B_1, p(a) = 0$; $\forall a \in B_1, p(a) > 0$
$\forall b \notin B_2, q(b) = 0$; $\forall b \in B_2, q(b) > 0$
$\forall b \in B_2, \forall b' \in A_2, u_2(\vec{p}, b) \geq u_2(\vec{p}, b')$
$\forall a \in B_1, \forall a' \in A_1, u_1(a, \vec{q}) \geq u_1(a', \vec{q})$

exponential time alg. by trying all subsets.

Sperner's lemma

Given any n-simplex (e.g. triangle, tetrahedron, ...), with corners coloured in distinct colours:

- Triangulate it in any way, where colours of vertices on each face must come from a colour at any of its corners.
- Then there exists an odd number of simplices whose vertices use all $(n+1)$ colours

$\uparrow$ dimensions.

Proof for $n=2$:

colours of vertices

$Q :=$ num. of triangles with (G,B,B) or (G,G,B)

$R :=$ num. of triangles with (R,G,B)

$X :=$ num of edges with (G,B) on boundary

$Y :=$ num of edges with (G,B) on interior

$\therefore\ 2Q + R =$ number of directed GB or BG edges $= 2Y + X$

$\uparrow$ can be shown to be odd.

$\because X$ is odd, $\therefore R$ is odd.

Brouwer fixed point thm: Given $f: A \to B$ continuous and K is compact convex set then:

$$\exists x \in K \text{ s.t. } f(x) = x$$

Nash theorem: A Nash equilibrium always exists.

Regret Minimisation:

$n :=$ number of actions

$\ell_\pi^t :=$ expected loss of algorithm π at time t

$\uparrow$ might play mixed strategies

$L_\pi^t = \sum_{\tau=1}^t \ell_\pi^\tau :=$ total loss of algorithm π up to time t.

Regret: $L_\pi^t - L_{best}^t$

$\uparrow$ best action, not best algorithm!

Greedy algorithm: At time t, choose the action that minimises L_i^t.

(tie-break by lower index)

Regret of $O(n)$

i.e. $\dfrac{L_{greedy}^T}{L_{best}^T} \in O(n)$

average is a Nash equilibrium.

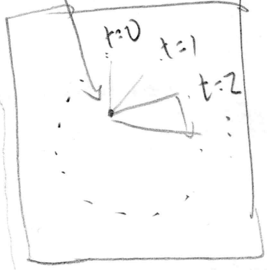

Randomised greedy alg.:

if there are ties for the best action, then choose uniform probability.

Multiplicative weight updates (MWU)

Initially: $\begin{cases} w_i^1 = 1 \\ p_i^1 = \frac{1}{n} \end{cases}$

At t, $\begin{cases} w_i^t = w_i^{t-1} e^{-\varepsilon \ell_i^{t-1}} \\ p_i^t = \frac{w_i^t}{W^t}, \quad W^t := \sum w_i^t \end{cases}$

when $\varepsilon = 0 \Rightarrow$ pure random play

$\varepsilon \to +\infty \Rightarrow$ randomised greedy.

Multiplicative weight updates

$$W^{t+1} = \sum_{i=1}^n w_i^{t+1} \geq w_{best}^{t+1}$$

$$w_i^{t+1} = e^{-\varepsilon L_i^t}$$

Formula: $\forall x \in [-1, 1]: e^x \le 1 + x + x^2$

$$\therefore W^{t+1} \le \sum_{i=1}^n w_i^t \left(1 + (-\varepsilon \ell_i^t) + (-\varepsilon \ell_i^t)^2\right)$$

(since $-1 \le \ell_i \le 1$) $\le (1+\varepsilon^2)\left(\sum w_i^t\right) - \varepsilon\left(\sum w_i^t \ell_i^t\right)$

$$= W_t \Big(1 + \varepsilon^2 - \varepsilon \underbrace{\sum_{i=1}^n \ell_i^t p_i^t}_{= \ell_\pi^t}\Big)$$

$1 + x \le e^x \rightarrow$

$$\le W_t e^{\varepsilon^2 - \varepsilon \ell_\pi^t}$$

continue

$$\le W_1 e^{t\varepsilon^2 - \varepsilon\left(\sum_{i=1}^t \ell_\pi^i\right)}$$

$$= n \cdot e^{t\varepsilon^2 - \varepsilon L_\pi^t}$$

$$\therefore \frac{e^{-\varepsilon L_{best}^t}}{} \le W_{t+1} \le n \cdot e^{t\varepsilon^2 - \varepsilon L_\pi^t}$$

$$\Rightarrow e^{\varepsilon L_\pi^t} \le e^{\varepsilon L_{best}^t} \cdot n \cdot e^{\varepsilon^2 t}$$

$$L_\pi^t \le L_{best}^t + \frac{\log n}{\varepsilon} + \varepsilon t.$$

when $t > \log n$, average loss per round for alg. is $O(1)$ to average loss per round for best.

Multiplicative weight update & minimax thm

If player 2 uses an online algorithm A with regret R, then average loss is at most $v_- + \frac{R}{T}$ after T rounds:

Pf: $L_A^T \leq L_{best}^T + R \leq T v_- + R \Rightarrow \frac{L_A^T}{T} \leq v_- + \frac{R}{T}$.

there is some pure strategy that guarantees P2 a loss of at most v_- against any $\vec{p}$ at every round.

$\Rightarrow$ If both players use MWU to pick a strategy at every round, then their average strategies are a Nash eqm of the underlying minimax game.

Routing games

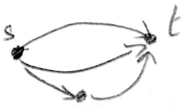

$\ell_e : \mathbb{R} \to \mathbb{R}^+$
(congestion function of edge e)

- atomic when edges must have integer flow

Equilibrium: switching to a different path never yields a better result.

Braess' Paradox: Adding an edge increases the social cost.

Price of Anarchy: $\frac{\text{WorstNash}(G)}{\text{OPT}(G)}$

In non-atomic version, pure Nash eqm always exists!

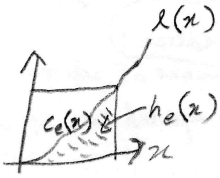

Let $c_e(x) := x \cdot \ell_e(x)$ (total social cost of congestion at e)

Let $c'_e(x) := \frac{d}{dx} c_e(x)$

Lemma: $(\forall P, P' \in \mathcal{P}_i,\ f_P^* > 0 \Rightarrow c'_P(f^*) \leq c'_{P'}(f^*))$ iff a flow is optimal. (i.e. marginal social cost is no more than any other path's)

$\mathcal{P}_i$: set of paths

$\ell_e^*(x) := c'_e(x) = \ell_e(x) + x \cdot \ell'_e(x)$; $\ell_P^*(x) = \sum_{e \in P} \ell_e^*(x)$

Corollary: a flow is optimal for $\langle G, r, \ell \rangle \iff$ it is a NE for $\langle G, r, \ell^* \rangle$

Thm: If there exists $\alpha > 1$ s.t. $c_e(x) \leq \alpha h_e(x)$ $\forall e \in E$, then $PoA \leq \alpha$.

$h_e(x) := \int_0^x \ell_e(y)\,dy$

Cooperative games

examples:
- Induced subgraph (undired, weighted): value of a coalition is the sum of edge weights in the coalition.
- Network flow (directed, weighted): value of a coalition is the maxflow using those edges only.
- Weighted voting games: $(w_1, \ldots, w_n ; q)$ — q: a cutoff; each player (n players total) has a weight.
  coalition value $= 1$ if sum of member weights $\geq q$, 0 otherwise
- Bankruptcy problem: split debt amongst creditors in some way.

$N = \{1, 2, \ldots, n\}$

characteristic fn: $v : 2^N \to \mathbb{R}$ (value function mapping from power set to $\mathbb{R}$)

coalition structure (CS) = some partition of N; $OPT(G) = \max_{CS} \sum_{S \in CS} v(S)$

imputation: $\forall S \in CS, \sum_{i \in S} x_i = v(S)$ (the amount to pay each player)

simple game: $v(S) \in \{0, 1\}$

monotone: $\forall S \subseteq T \subseteq N \implies v(S) \leq v(T)$ (i.e. adding people to a group is not worse.)

super additive game: $\forall S, T \subseteq N$: $v(S) + v(T) \leq v(S \cup T)$ where S, T are disjoint

convex game: $S \subseteq T \subseteq N$ and $i \in N \setminus T \Rightarrow v(S \cup \{i\}) - v(S) \leq v(T \cup \{i\}) - v(T)$

i.e. joining the big group is always better.

Core of a cooperative game:

some S elements in the coalition

making a new coalition with those elements only.

An imputation $\vec{x}$ is in the core if $\sum_{i\in S} x_i =: \underbrace{x(S) \geq v(S)}$, $\forall S \subseteq N$

(i.e. no subset of the coalition will want to break off.)

* $\Rightarrow \underbrace{v(S)} \leq \sum_{i\in S} x_i \leq \underbrace{v(N) - v(N \setminus S)}$

amount that S can get on their own.

marginal utility of the coalition by adding S to them

Proof that core is empty is NP-hard.

Simple game: $\forall S \subseteq N$, $v(S) \in \{0,1\}$

winning coalition: those with value 1

losing coalition: otherwise

veto player: a player that is a member of every winning coalition (can't win without them)

Thm: Let $G = \langle N, v\rangle$ be a simple game, and $v(N) = 1$, then $\left(\text{Core}(G) \neq \phi \Leftrightarrow G \text{ has veto players}\right)$.

Lemma: Core of induced subgraph game is not empty $\Leftrightarrow$ graph has no negative cut.

Shapley value: $\frac{1}{n!} \sum_{\sigma \in \Pi(N)} m_i(\sigma)$

marginal contribution of i in σ.

Satisfies:

- Efficiency $\left(\sum_{i\in N} \phi_i = v(N)\right)$: all money are distributed
- Symmetry $\left(\forall S \subseteq N \setminus \{i,j\} : v(S\cup\{i\}) = v(S \cup \{j\}) \Rightarrow \phi_i = \phi_j\right)$: equal players are paid equally
- Dummy/Null player $\left(\forall S \subseteq N \setminus \{i\} : v(S\cup\{i\}) = v(S) \Rightarrow \phi_i = 0\right)$: those who don't contribute are not paid anything
- Additivity/Linearity $\left(\phi_i(v_1) + \phi_i(v_2) = \phi_i(v_1+v_2)\right)$: combined game combines the payment

value function 1   value function 2

Thm: Shapley value is the only division that satisfies all four properties

Nash bargaining solution

$= \max_{(v_1,v_2)\in S} (v_1 - d_1)(v_2 - d_2)$

assuming they are in first quadrant of (v_1, v_2).

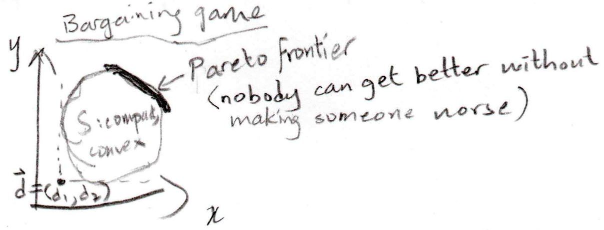

Satisfies:

- Efficiency: no outcome Pareto-dominates $(f_1(S,\vec{d}), f_2(S,\vec{d}))$
- Symmetry: reflecting S and $\vec{d}$ on $y = x$ should reflect the chosen point
- Independence of Irrelevant Alternatives (IIA): Removing portions of S that do not contain the chosen point should not change the chosen point.
- Invariance under Equivalent Representations (IER): Translation and scaling on either axis should translate and scale the chosen point in the same way.

Good markets

- Thick (lots of buyers & sellers, everyone is aware of their options)
- Timely (not too fast or too slow)
- Safe (people cannot be hurt by revealing preferences, fair outcomes, people are better off by participating)

Unravelling

- Matching is done earlier & earlier (bad)

Matching scenario

- Set of students $S = \{s_1, \dots, s_n\}$
- Set of hospitals $H = \{h_1, \dots, h_m\}$
- Each student ranks hospitals with a total ordering : $\succ_s$
- Each hospital ranks students with a total ordering : $\succ_h$

Outcome: A one-to-one matching $M : S \to H$

Blocking pair of a matching M : A pair (s,h) where $h \succ_s M(s)$ and $s \succ_h M^{-1}(h)$ (i.e. both of them prefer each other to their assigned counterpart)

Gale-Shapley Algorithm

→ polytime & returns a stable matching

must always exist.

- Start with all students unassigned
- While there are unassigned students :
  - Each unassigned student proposes to their favourite not-yet-proposed-to hospital
  - Each hospital looks at the list of students that proposed to it at this round and whoever is assigned to it already (if exists), and picks the most prefered one. All others remain/become unassigned.

Matching with couples → stable matching might not exist (but in practice usually exists)

Stable if : No (s,h) prefers each other

No $((c_1, c_2), (h_1, h_2))$ prefers each other

Thm: Gale-Shapley assigns each student to their most preferred hospital in which some stable matching exists.

⇒ It is better for students than hospitals

⇒ Students cannot game the system, but hospitals can (to get a more preferred student)

Allocation of Indivisible Goods

- $\pi(i)$: set of goods allocated to player i under π

$v_i(S) := \sum_{g \in S} v_{ig}$ (additive valuations) — value of good g to player i.

Assume that $\forall i,j \in N,\ v_i(N) = v_j(N)$.

Desirable properties

- Optimality : π is optimal $:= \pi \in \arg\max_{\pi'} \sum_i v_i(\pi'(i))$
- Pareto optimality : no allocation dominates π. (somebody will get strictly worse off in any other outcome)
- Envy-freeness : no player wants another's bundle : $\forall i,j \in N,\ v_i(\pi(i)) \geq v_i(\pi(j))$
- Maximin share : if I get to partition the items, what is the maximum value (to me) of the worst bundle? (it is independent of the goods given to other players). i.e. $MMS_i := \max_{\pi} \min_{S \in \pi} v_i(S)$
- Maximin share requirement: Each player gets at least their maximin share, i.e. $\forall i \in N,\ v_i(\pi(i)) \geq MMS_i$

Approximate solutions:

- EF-1 : no player wants another's bundle with the best good removed : $\forall i,j \in N,\ \exists g \in \pi(j)$ s.t. $v_i(\pi(i)) \geq v_j(\pi(j) \setminus \{g\})$
- α-EF (for some $0 < \alpha < 1$) : $\forall i,j \in N,\ v_i(\pi(i)) \geq \alpha \cdot v_i(\pi(j))$
- α-MMS (for some $0 < \alpha < 1$) : $\forall i \in N,\ v_i(\pi(i)) \geq \alpha \cdot MMS_i$

Thms for indivisible goods with additive valuations

- EF allocation does not always exist (consider a single good divided amongst two players)
- EF-1 allocation always exists
- MMS allocation does not always exist
- $\frac{2}{3}$-MMS allocation always exists
- Deciding if an EF allocation exists is NP-complete (by reduction from PARTITION problem, where every item is valued the same by all players)

Algorithm for finding an EF-1 allocation in $O(\underbrace{m}_{\text{number of items}}\underbrace{n^3}_{\text{number of players}})$ time:

- While there is an unallocated good g:
  - Give g to a player that nobody envies (which must always exist because the envy graph is a DAG)
  - Decycle the envy graph
- End.

For each cycle, swap the bundles amongst them. The cycle disappears, and other people that envy player j will now envy the player that j's bundle was given to. So the number of edges decreases. Continue until there are no cycles left.

Rent Division

$N := \{1, \dots, n\}$ players

$G := \{g_1, \dots, g_n\}$ rooms

$v_{ij} :=$ player i's valuation of room j.

Assume that $\forall i,j \in N$, $\sum_k v_{ik} = \sum_k v_{jk} =: R$ ← total rent that needs to be paid.

Output: a room allocation $\sigma : N \to N$ and a rent division $\vec{p}$ where $\sum_j p_j = R$ (p_j: payment for room j)

EF outcome: $\langle \sigma, \vec{p} \rangle$ s.t. $v_{i\sigma(i)} - p_{\sigma(i)} \geq v_{ij} - p_j \ \forall i,j \in N$.

i.e. "I don't prefer your room for the price you're being charged"

Thms for rent division:

- EF outcome always exists, and can be computed efficiently
- In any EF outcome, room allocation (i.e. σ) is optimal (1st welfare thm)
- If outcome $\langle \sigma, \vec{p} \rangle$ is EF, then so is $\langle \sigma', \vec{p} \rangle$ for any optimal room allocation σ'. (2nd welfare thm) — same price vector

General algoritmic framework

- Compute a socially optimal allocation (max weighted matching) → Find an EF price vector (linear programming) → This price vector will work with any optimal allocation (2nd welfare thm)

Ways to choose amongst EF outcomes

- Equitability: minimise the disparity between the highest and lowest utilities: $\arg\min_{\vec{p} \in EF} \max_{i,j \in N} (u_i(\vec{p}) - u_j(\vec{p}))$
- Maximin: maximise the minimum utility: $\arg\max_{\vec{p} \in EF} \min_{i \in N} u_i(\vec{p})$

Thms:

- There is a unique maximin EF price vector, and this vector is also equitable (however, there might exist equitable price vectors that are not maximin)

Single item auctions

- English auction: auctioneer sets a starting price, bidders take turns raising their bids, last bidder wins and pays his bid.
  - It is rational to bid iff $p + \delta \leq v$ (p: current price, v: my valuation), and I should bid $p + \delta$. Winner will eventually pay second highest value or second highest value $+ \delta$.
- Japanese auction: auctioneer sets a starting price then keeps raising it until all but one bidder drops out, last bidder pays the current price.

- Dutch auction: auctioneer sets a high starting price and then starts lowering it until a bidder accepts the price.
- Sealed-bid auction: all bidders simultaneously submit their bids, the higher bidder gets the item and pays
  - his bid (first-price auction), or
  - 2nd highest bid (second-price, or Vickrey, auction)

[Bidders must trust the auctioneer in a sealed-bid auction.

Thm: In a Vickrey auction, truthful bidding is a dominant strategy (i.e. truthful bidding will never make a person worse off regardless of the actions of everyone else.)

Note: There are many NEs in a Vickrey auction.
e.g. $\vec{v} = (50, 30, 70)$ (valuations)
possible NE: $(50, 30, 70)$, $(0, 0, 70)$, $(70, 0, 0)$

$(0,0,70)$: third player wins and pays nothing
$(70,0,0)$: first player wins and pays nothing.

Multi unit auctions (identical items)
- There are multiple, identical items
- Each bidder wants one item, and values it at v_i
- How should the items be distributed, and how much should they pay?

Multi unit auctions (different items)
- There are multiple, different items.
- Each player wants one item, values item j at v_{ij}.
- How should the items be distributed, and how much should they pay?

(General) Mechanism Design
- Players $N = \{1, \dots, n\}$
- Outcomes $O = \{o_1, \dots, o_m\}$
- Each player has a utility function $u_i : O \to \mathbb{R}$ (takes into consideration payments if any)
- Centre chooses an outcome to maximise some function (e.g. $\sum_i u_i(o^*)$).

Incentive Compatibility: If everyone else reports their true valuations, I should report truthfully too. (i.e. reporting true valuations is a NE)

Dominant Strategy Incentive Compatibility (i.e. strategyproofness): I should report truthfully regardless of everyone else. (i.e. reporting true valuations is a (weakly) dominant strategy.)

Revelation Principle: Given any mechanism M, there exists a mechanism M' whose inputs are users' valuations, and whose outputs are exactly like those of M. (i.e. there is always an incentive-compatible mechanism)

Problems:
- computational burden is pushed to the center
- the direct mechanism might have additional bad equilibria.

Vickrey-Clarke-Groves Mechanisms
1. Choose an outcome that maximises $\sum_i v_i(o^*)$ (i.e. socially optimal outcome)
2. To determine the payment that player j must make:
   - pretend j does not exist, and choose o^*_{-j} that maximises $\sum_{i \neq j} v_i(o^*_{-j})$
   - j pays $\sum_{i\neq j} v_i(o^*_{-j}) - \sum_{i \neq j} v_i(o^*)$ (i.e. the negative externality that player j imposes on others)

Thm: In a VCG mechanism, truthful reporting is a dominant strategy.

Proof: $u_i(o^*) = v_i(o^*) - \left(\sum_{i\neq j} v_i(o^*_{-j}) - \sum_{i \neq j} v_i(o^*)\right)$

$$= \underbrace{\sum_i v_i(o^*)}_{\text{total social welfare}} - \underbrace{\sum_{i\neq j} v_i(o^*_{-j})}_{\text{does not depend on player } j\text{'s reporting}}$$

(any other outcome o' must have lower total social welfare than o^* (which is socially optimal))

Assumptions:
- Choice set monotonicity: $O_{-i} \subseteq O$
- No negative externalities: $\forall o_{-i} \in O_{-i}, v_i(o_{-i}) \geq 0$

Facility Location

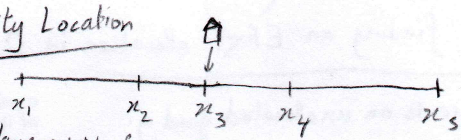

- Players: $N = \{1, \dots, n\}$
- Each player has a location $x_i \in \mathbb{R}$ (WLOG $x_1 \leq x_2 \leq \dots \leq x_n$).
- Centre should find an optimal location to build the facility. $f : \mathbb{R}^n \to \mathbb{R}$

Measures: (To minimise...)
- Total cost: $\sum_i |f(\vec{x}) - x_i|$
- Max cost: $\max_i |f(\vec{x}) - x_i|$

Thm: Choosing the median position (tiebreak towards larger player?) has the minimum total cost. It is also strategyproof and group-strategyproof.

strategyproof: regardless of what others do, it is best to play truthfully

group-strategyproof: no subset of players can collude to make everybody in the subset strictly better off.

Thm: Any deterministic truthful mechanism has worst-case approximation of at least 2 to the maximum cost. (i.e. we are using the max cost measure.)

Proof by contradiction, with two players.

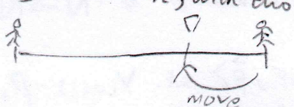

player to optimal position, then reason backwards

Thm: Any randomised truthful mechanism has worst-case expected approximation of at least $\frac{3}{2}$ to the maximum cost.

Proof:

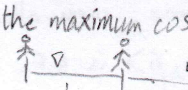

WLOG let $E[f] \leq \frac{1}{2}$. Then the person at 1 has expected cost of $\geq \frac{1}{2}$. Then if he lies and claims that he is at 2 instead, the mechanism has to pick a location either $\leq \frac{1}{2}$ or $\geq \frac{3}{2}$.

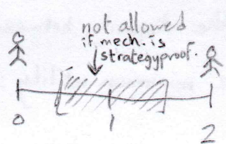

Now, if instead we have one person at 0 and one person at 2 playing truthfully, due to the above reasoning, the mechanism has to pick a location either $\leq \frac{1}{2}$ or $\geq \frac{3}{2}$, hence the maximum cost is $\geq \frac{3}{2}$.