Algorithmic game theory

Martin Balko

13th lecture

January 11th 2024

Multi-parameter mechanism design

Multi-parameter mechanism design

- In multi-parameter mechanism design, we have the following setting:
  - n strategic bidders,
  - a finite set Ω of outcomes,
  - each bidder i has a private valuation $v_i(\omega) \geq 0$ for every outcome $\omega \in \Omega$.
- Each bidder i submits his bids $b_i(\omega) \geq 0$ for each $\omega \in \Omega$ and our goal is to design a mechanism that selects an outcome $\omega \in \Omega$ so that it maximizes the social surplus $\sum_{i=1}^{n} v_i(\omega)$.
- The valuations now depend on possible outcomes, so, for example, if bidders compete for a single item, each bidder can have an opinion about each other bidder winning the item as well.
- Example (single-item auction): we set $\Omega = \{\omega_1, \ldots, \omega_n, \omega_\emptyset\}$ has size $n + 1$ and each outcome ω_i with $i \in \mathbb{N}$ corresponds to the winner i of the item. The last outcome $\omega_\emptyset$ corresponds to nobody getting the item. The valuations are $v_i(\omega_j) = 0$ for every $j \neq i$ and $v_i(\omega_i) = v_i$ otherwise.

The Vickrey–Clarke–Groves (VCG) mechanism

VCG mechanism (Theorem 3.18)

In every multi-parameter mechanism design environment, there is a DSIC social-surplus-maximizing mechanism.

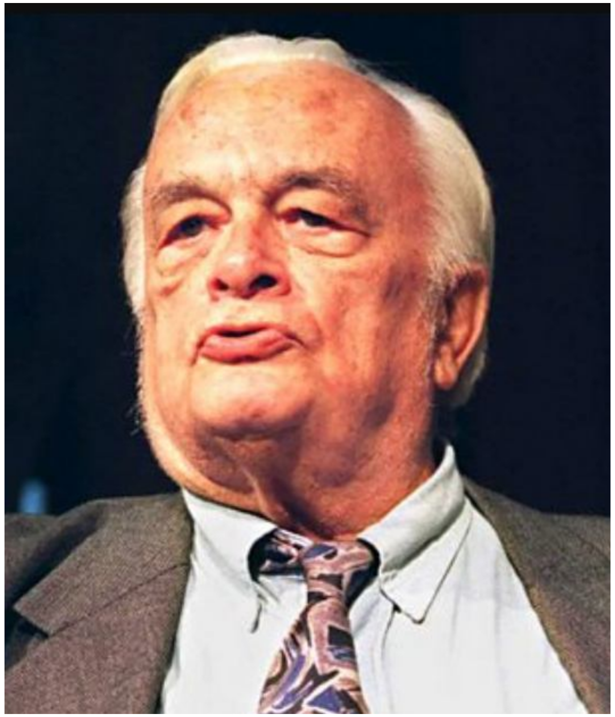

Figure: William Vickrey, Edward H. Clarke, and Theodore Groves.

Sources: : https://en.wikipedia.org, https://www.demandrevelation.com/, and https://www.researchate.net/

- We now present the proof.

VCG mechanism: proof idea

- The key idea is to consider the the loss of social surplus inflicted on the other $n-1$ bidders by the presence of bidder i. For example, in single-item auctions, the winning bidder inflicts a social surplus loss of the second-highest bid to the others.
- We define the payments to force each bidder to care about the others.
- We will see that the following allocation rule works

$$x(b) = \mathrm{argmax}_{\omega \in \Omega} \sum_{i=1}^{n} b_i(\omega)$$

together with this payment formula

$$p_i(b) = \max_{\omega \in \Omega} \left\{ \sum_{\substack{j=1 \\ j \neq i}}^{n} b_j(\omega) \right\} - \sum_{\substack{j=1 \\ j \neq i}}^{n} b_j(\omega^*),$$

where $\omega^* = x(b)$ is the outcome chosen by our allocation rule x for given bids b.

VCG auction example

- The government wants to construct roads connecting diverse cities, and he wants cities to pay for the roads.

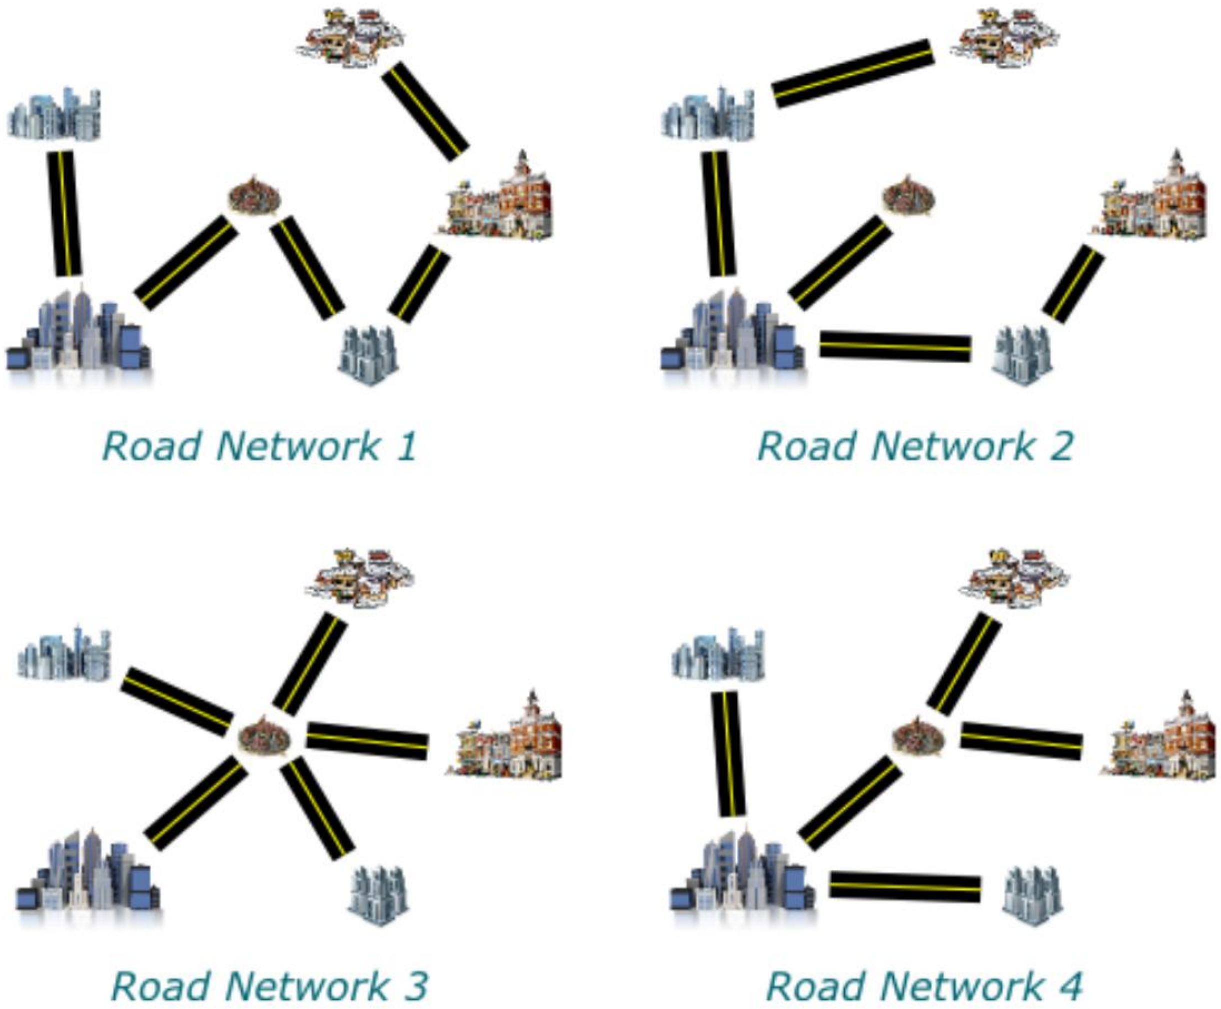

Sources: https://www.science4all.org/article/auction-design/

VCG auction example

	6 M$	14 M$	2 M$	**16 M$**
	5 M$	8 M$	4 M$	**12 M$**
	2 M$	1 M$	**20 M$**	4 M$
	4 M$	**6 M$**	3 M$	5 M$
	1 M$	1 M$	**6 M$**	2 M$
	1 M$	2 M$	2 M$	**3 M$**
Total (social welfare)	19 M$	32 M$	37 M$	**42 M$**

Sources: https://www.science4all.org/article/auction-design/

- Cities pay their negative externalities on the collectivity. Other cities would be happier without the biggest city (NYC, say). How much happier they would be is exactly what NYC must pay.

VCG auction example

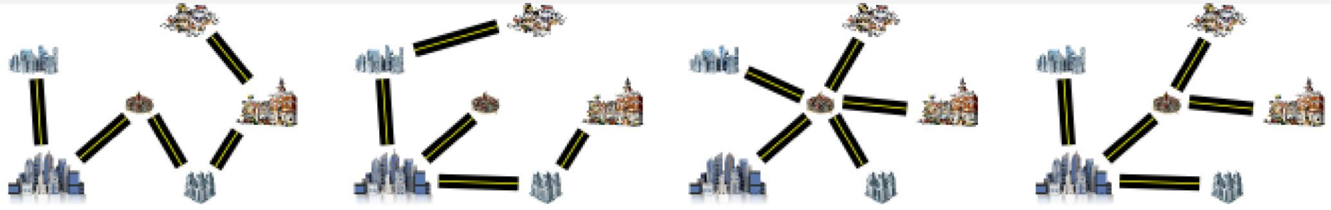

	Road Network 1	Road Network 2	Road Network 3	Road Network 4
	6 M$	14 M$	2 M$	**16 M$**
	5 M$	8 M$	4 M$	**12 M$**
	2 M$	1 M$	**20 M$**	4 M$
	4 M$	**6 M$**	3 M$	5 M$
	1 M$	1 M$	**6 M$**	2 M$
	1 M$	2 M$	2 M$	**3 M$**
Total (social welfare)	19 M$	32 M$	37 M$	**42 M$**

Sources: https://www.science4all.org/article/auction-design/

- If NYC was not there, then road network number 3 (RN3) would have been chosen, as opposed to RN4. The value of RN3 for the other cities would be 35 M$, as opposed to the 26 M$ of RN4. Therefore, the negative externality of NYC is $35 - 26 = 9$ M$.

VCG mechanism: proof I

- We proceed in two steps. First, we assume, without justification, that bidders truthfully reveal their private information, and figure out which outcome from Ω to pick.
- We maximize the social surplus, so our allocation rule needs to pick an outcome that maximizes the social surplus. So, given bids $b = ((b_1(\omega))_{\omega\in\Omega}, \ldots, (b_n(\omega))_{\omega\in\Omega})$, we define the allocation rule by

$$x(b) = \operatorname{argmax}_{\omega\in\Omega} \sum_{i=1}^{n} b_i(\omega).$$

- Second, we need to choose payment rule $p = (p_1, \ldots, p_n)$ so that the multi-parameter mechanism (x, p) is DSIC.
- We choose p so that our assumption about truthful bidders is justified.
- The key idea turns is considering the the loss of social surplus inflicted on the other $n-1$ bidders by the presence of bidder i.

VCG mechanism: proof II

- Formally, we choose the payment rule as

$$p_i(b) = \max_{\omega \in \Omega} \left\{ \sum_{\substack{j=1 \\ j \neq i}}^{n} b_j(\omega) \right\} - \sum_{\substack{j=1 \\ j \neq i}}^{n} b_j(\omega^*)$$

  for every bidder i, where $\omega^* = x(b)$ is the outcome chosen by our allocation rule x for given bids b.
- The first term is the surplus of the remaining $n-1$ bidders if we omit bidder i. The second term is the social surplus if we consider bidder i.
- By definition, the mechanism (x, p) is maximizing social surplus, assuming truthful bids. It remains to prove that it is DSIC.
- That is, we need to show that each bidder i maximizes his utility $v_i(x(b)) - p_i(b)$ by setting $b_i(\omega) = v_i(\omega)$ for every $\omega \in \Omega$.
- One can show that we have $0 \leq p_i(b) \leq b_i(\omega^*)$ (Exercise), hence truthtelling agents are guaranteed non-negative utility.

VCG mechanism: proof III

- We fix bidder i and the bids of other bidders b_{-i}.
- If $x(b) = \omega^*$, then, by the choice of p, the utility of bidder i equals

$$v_i(\omega^*) - p_i(b) = \left(v_i(\omega^*) + \sum_{j \neq i} b_j(\omega^*) \right) - \left(\max_{\omega \in \Omega} \left\{ \sum_{j \neq i} b_j(\omega) \right\} \right).$$

- The second term is independent of b_i, thus bidder i needs to maximize the first term of the expansion in order to maximize his utility. However, bidder i cannot influence ω^* directly, the mechanism (x, p) gets to choose ω^* so that sums of bids are maximized.
- Best case for bidder i is when the mechanism picks ω^* that maximizes the first term of the expansion, that is, bidder i wants to select

$$\operatorname{argmax}_{\omega \in \Omega} \left\{ v_i(\omega) + \sum_{j \neq i} b_j(\omega) \right\}.$$

- If i bids truthfully, then this agrees with our choice of x. Thus, bidding truthfully results in maximizing i's utility. $\square$

Proof of Myerson's lemma

Myerson's lemma

Myerson's lemma (Theorem 3.8)

In a single-parameter environment, the following three claims hold.

(a) An allocation rule is implementable if and only if it is monotone.

(b) If an allocation rule x is monotone, then there exists a unique payment rule p such that the mechanism (x, p) is DSIC (assuming that $b_i = 0$ implies $p_i(b) = 0$).

(c) The payment rule p is given by the following explicit formula

$$p_i(b_i; b_{-i}) = \int_0^{b_i} z \cdot \frac{\mathrm{d}}{\mathrm{d}z} x_i(z; b_{-i}) \, \mathrm{d}z$$

for every $i \in \{1, \ldots, n\}$.

- We have applied this result many times, but we have not seen its proof yet. Let's fix that.

Proof of Myerson's lemma I

- Let x be an allocation rule and p be a payment rule such that (x, p) is DSIC. We prove all three claims at once use a clever swapping trick.
- The DSIC property says that, for every z, $u_i(z; b_{-i}) = v_i \cdot x_i(z; b_{-i}) - p_i(z; b_{-i}) \leq v_i \cdot x_i(v_i; b_{-i}) - p_i(v_i; b_{-i})$.
- For two possible bids y and z with $0 \leq y < z$, bidder i might as well have private valuation z and can submit the false bid y if he wants, thus the DSIC condition gives

$$u_i(y; b_{-i}) = z \cdot x_i(y; b_{-i}) - p_i(y; b_{-i}) \leq z \cdot x_i(z; b_{-i}) - p_i(z; b_{-i}) = u_i(z; b_{-i}).$$

- Analogously, we can have $v_i = y$ and $b_i = z$ and thus (x, p) satisfies

$$u_i(z; b_{-i}) = y \cdot x_i(z; b_{-i}) - p_i(z; b_{-i}) \leq y \cdot x_i(y; b_{-i}) - p_i(y; b_{-i}) = u_i(y; b_{-i}).$$

- By putting these inequalities together, we obtain the following payment difference sandwich:

$$z(x_i(y; b_{-i}) - x_i(z; b_{-i})) \leq p_i(y; b_{-i}) - p_i(z; b_{-i}) \leq y(x_i(y; b_{-i}) - x_i(z; b_{-i})).$$

- Since $0 \leq y < z$, we obtain $x_i(y; b_{-i}) \leq x_i(z; b_{-i})$. Thus, if (x, p) is DSIC, then x is monotone.

Proof of Myerson's lemma II

- In the rest of the proof, we assume that the allocation x is monotone.
- Let i and b_{-i} be fixed, so we consider x_i and p_i as functions of z.
- First, we also assume that the function x_i is piecewise constant. Thus, the graph of x_i consists of a finite number of intervals with "jumps" between consecutive intervals:

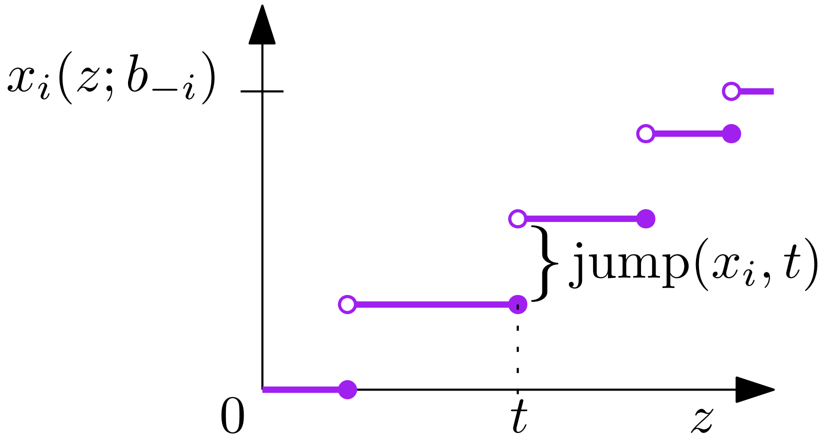

- For a piecewise constant function f, we use $\mathrm{jump}(f, t)$ to denote the magnitude of the jump of f at point t.
- If we fix z in the payment difference sandwich and let y approach z from below, then both sides become 0 if there is no jump of x_i at z. If $\mathrm{jump}(x_i, z) = h > 0$, then both sides tend to $z \cdot h$.

Proof of Myerson's lemma III

- Thus, if the mechanism (x, p) is supposed to be DSIC, then the following constraint on p must hold for every z:

$$\mathrm{jump}(p_i, z) = z \cdot \mathrm{jump}(x_i, z).$$

- If we combine this constraint with the initial condition $p_i(0; b_{-i}) = 0$, we obtain a formula for the payment function p for every bidder i and bids b_{-i} of other bidders,

$$p_i(b_i; b_{-i}) = \sum_{j=1}^{\ell} z_j \cdot \mathrm{jump}(x_i(\,\cdot\,; b_{-i}), z_j),$$

  where $z_1, \ldots, z_\ell$ are the breakpoints of the allocation function $x_i(\,\cdot\,; b_{-i})$ in the interval $[0, b_i]$.
- With some additional facts from calculus, this argument can be generalized to general monotone functions x_i. We omit the details.

Proof of Myerson's lemma IV

- It remains to show that if x is monotone, then the mechanism (x, p) is indeed DSIC.
- This argument works in general, but we present it only for piecewise constant functions.
- We recall that the utility $u_i(b_i; b_{-i}) = v_i \cdot x_i(b_i; b_{-i}) - p_i(b_i; b_{-i})$.
- Using the expression of the payment, we see that the payment $p_i(b_i; b_{-i})$ of bidder i corresponds to the part of $[0, b_i] \times [0, x_i(b_i; b_{-i})]$ lying to the left of the curve $x_i(\cdot ; b_{-i})$.
- It will follow from a picture that it is optimal for bidder i to bid $b_i = v_i$.

Proof of Myerson's lemma by picture

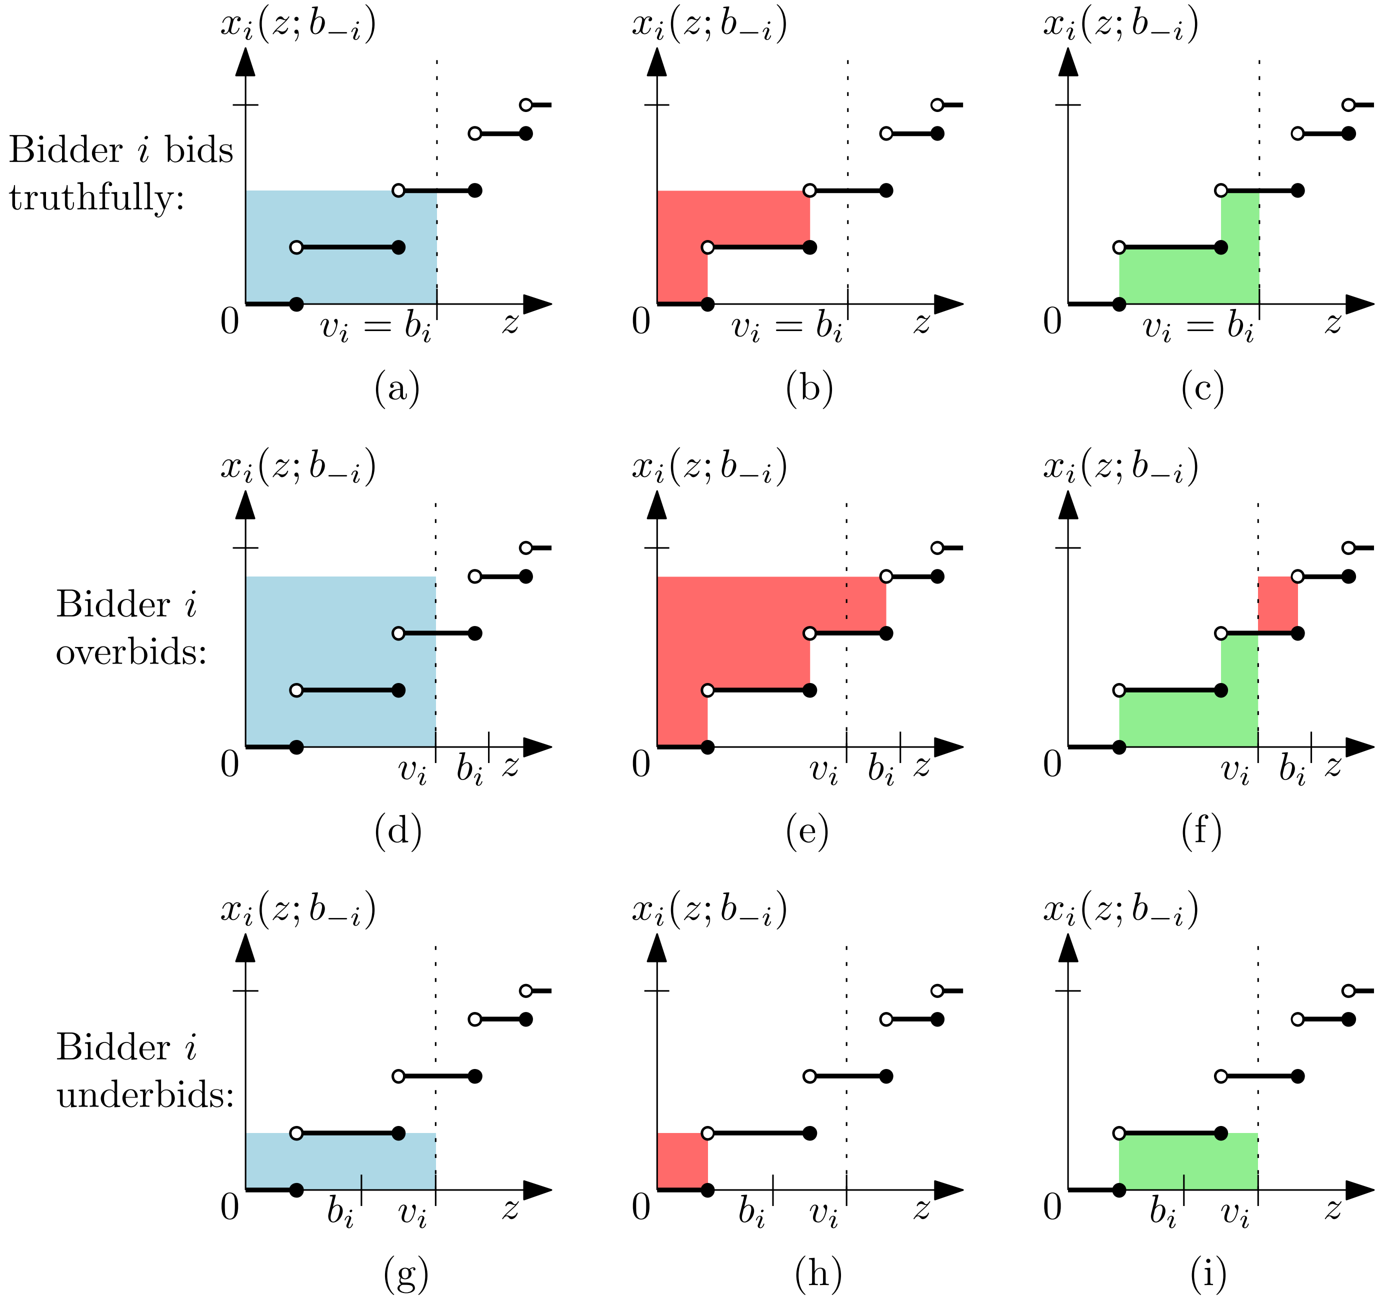

Revelation principle

Going beyond DSIC?

- So far we have considered only DSIC mechanisms, as they are easy to play and predict. A natural question is whether we lose anything by restricting ourselves to these mechanisms.
- We split the DSIC property into the following two parts:
  - Each bidder has a dominant strategy, no matter what his private valuation is.
  - This dominant strategy is direct revelation (that is, truthtelling).
- Example (auction that satisfies only the first property):
  - Consider a single-item auction, where the seller on given bids $(b_1, \dots, b_n)$ runs a Vickrey's auction on bids $(2b_1, \dots, 2b_n)$.
  - Then each bidder has a dominant strategy and thus the first property is satisfied.
  - However, this dominant strategy is not direct revelation, but to bid half of your value.
- So is DSIC too restrictive?

Revelation principle

- No! We will see that only the first condition matters while the second one then comes for free.
- Clearly, the first condition is not always satisfied (first-price auctions). Without it, it is hard to predict the behavior of the bidders. Although it sometimes makes sense to relax the first condition.
- The Revelation principle says that there is no need to relax the second condition.

Revelation principle (Theorem 3.19)

For every multi-parameter mechanism M in which every bidder has a dominant strategy, no matter what his private valuation is, there is an equivalent mechanism M' in which each bidder has a dominant strategy that is a direct revelation.

Proof of the Revelation principle

- The proof proceeds by a simulation argument.
- For each bidder i and his valuations $v_i(\omega)_{\omega\in\Omega}$, let $s_i(v_i(\omega)_{\omega\in\Omega})$ be the dominant strategy of i in the mechanism M.

- We now construct the mechanism M' that additionally satisfies the second condition:
  - The mechanism M' accepts sealed bids $b_1(\omega)_{\omega\in\Omega}, \dots, b_n(\omega)_{\omega\in\Omega}$ from the bidders.
  - Then, M' submits the bids $s_1(b_1(\omega)_{\omega\in\Omega}), \dots, s_n(b_n(\omega)_{\omega\in\Omega})$ to M and M' outputs the same outcome as M.

- The direct revelation is a dominant strategy in M', as if a bidder i has valuations $v_i(\omega)_{\omega\in\Omega}$, then submitting any other bid than $v_i(\omega)_{\omega\in\Omega}$ can only result in playing a different strategy than $s_i(v_i(\omega)_{\omega\in\Omega})$ in M'. This, however, can only decrease the utility of i. $\square$

Exams

Exams info

- Exam format:
  - Oral exam with preparation, 3 hours max.
  - I will ask about three topics, one survey question, one question with proofs, and one exercise.
  - The more points you have, the easier the exam is. If you have at least 25 points, you do not need to solve the exercise part.
- Dates (so far):
  - 14.1. – 9:00–19:00, capacity 30
  - 20.1. – 9:00–19:00, capacity 30
  - 28.1. – 9:00–16:00, capacity 20
  - 3.2. – 9:00–16:00, capacity 20
  - 11.2. – 9:00–16:00, capacity 20
- What you should know: everything that we covered (everything is included in the lecture notes).

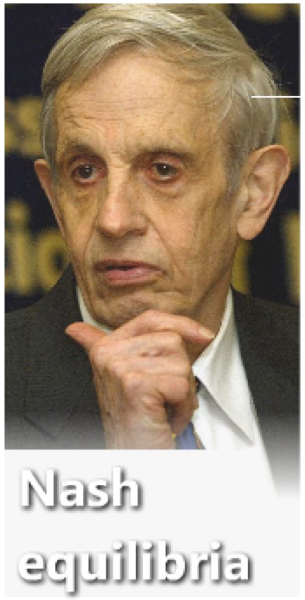

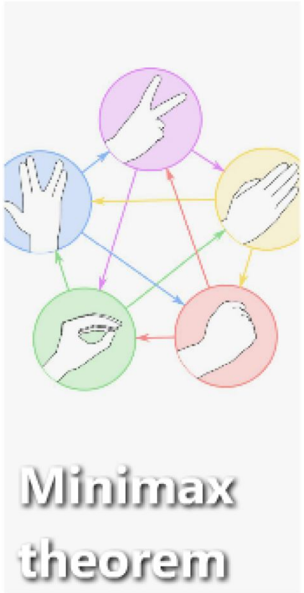

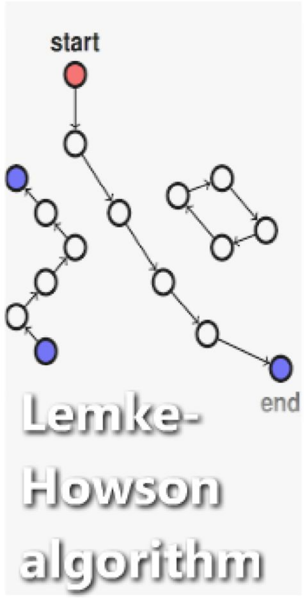

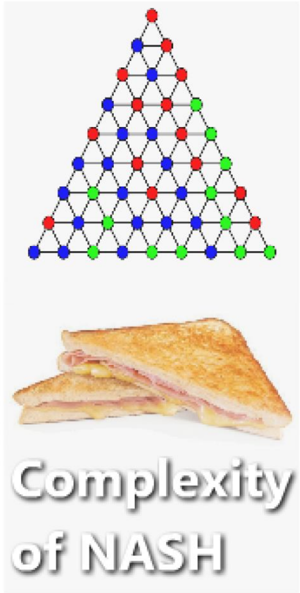

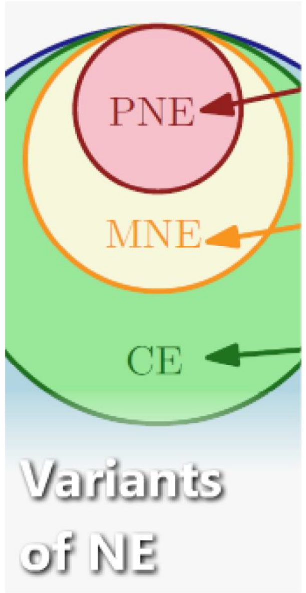

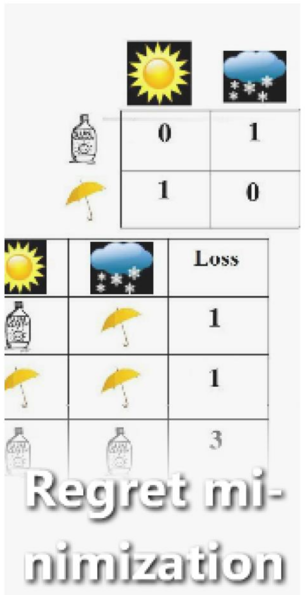

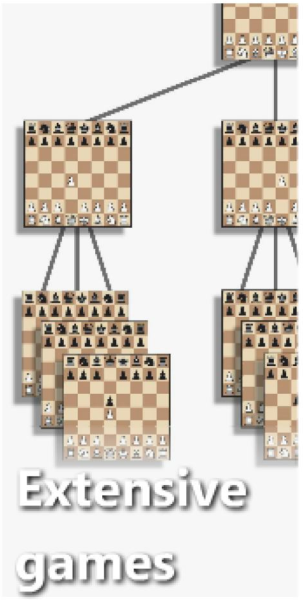

Thank you for your attention.