Algorithmic game theory

Martin Balko

12th lecture

December 20th 2024

Revenue maximization: what we know

- For revenue maximization, we had to consider Bayesian model, where each bidder i draws his valuation according to some probability distribution F_i (with density f_i) that is known to the seller. We then try to maximize the expected revenue.
- In DSIC mechanisms, maximizing the expected revenue is then the same as maximizing the expected virtual social surplus $\sum_{i=1}^{n} \varphi_i(v_i)x_i(v)$. (Theorem 3.13) where

$$\varphi_i(x) = x - \frac{1 - F_i(x)}{f_i(x)}.$$

- If $F_1 = \cdots = F_n = F$ is regular, then Vickrey's auction with reserve $\varphi^{-1}(0)$ maximizes the expected revenue among all single-item auctions.
- What if the seller does not know the distributions $F_1, \ldots, F_n$?

The Bulow–Klemperer theorem

The Bulow–Klemperer theorem

The Bulow–Klemperer theorem (Theorem 3.15)

Let $F = F_1 = \cdots = F_n$ be a regular probability distribution. Then,

$$\mathbb{E}_{v_1,\ldots,v_{n+1}\sim F}\left[\mathrm{Rev}(VA_{n+1})\right] \geq \mathbb{E}_{v_1,\ldots,v_n\sim F}\left[\mathrm{Rev}(OPT_{F,n})\right],$$

where $\mathrm{Rev}(VA_{n+1})$ is the revenue of Vickrey auction VA_{n+1} with $n+1$ bidders (and no reserve) and $\mathrm{Rev}(OPT_{F,n})$ denotes the revenue of the optimal auction $OPT_{F,n}$ for F with n bidders.

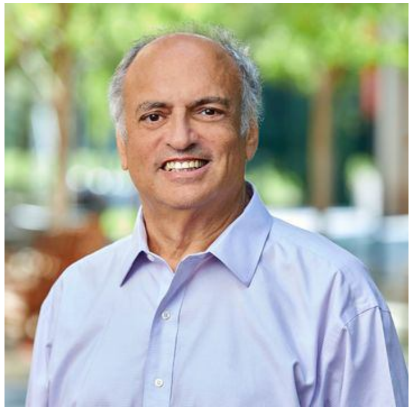

Figure: Jeremy Bulow and Paul Klemperer.

Sources: https://economics.stanford.edu/ and https://www.economics.ox.ac.uk/

- More competition is better than finding the right auction format.

The Bulow–Klemperer theorem: remarks

- If $F = F_1 = \cdots = F_n$ is a regular probability distribution, then

$$\mathbb{E}_{v_1,\ldots,v_{n+1}\sim F}[\mathrm{Rev}(VA_{n+1})] \geq \mathbb{E}_{v_1,\ldots,v_n\sim F}[\mathrm{Rev}(OPT_{F,n})].$$

- Recall that $OPT_{F,n}$ is the Vickrey auction with the reserve price $\varphi^{-1}(0)$, where φ is the virtual valuation function of F.
- The auction on the left side does not depend on F while the one on the right side does.
- The Bulow–Kemperer theorem implies that in all possible single-item auctions with identical regular probability distributions on valuations of $n \geq 2$ bidders, the expected revenue of Vickrey auction is at least $\frac{n-1}{n}$-fraction of the expected revenue of an optimal auction for the same number of bidders (Exercise).
- Informally: extra competition is more important than getting the auction format just right.

The Bulow–Klemperer theorem: proof

- We define an auxiliary auction $\mathcal{A}$ of $n+1$ bidders as follows:
  - Simulate the optimal auction $OPT_{F,n}$ on the bidders $1, \dots, n$,
  - If the item was not awarded, then give it to bidder $n+1$ for free.
- By definition, $\mathcal{A}$ always allocates the item and we have

$$\mathbb{E}_{v_1,\dots,v_{n+1}\sim F}\left[\mathrm{Rev}(\mathcal{A})\right] = \mathbb{E}_{v_1,\dots,v_n\sim F}\left[\mathrm{Rev}(OPT_{F,n})\right].$$

- We now argue that the Vickrey auction maximizes expected revenue over all auctions that are guaranteed to allocate the item.
- It follows from Theorem 3.13 that such an optimal auction awards the item to the bidder with the highest virtual valuation (even if this is negative). Since $F = F_1 = \cdots = F_n$ and F is regular, this bidder is always the bidder with the highest valuation and the Vickrey auction VA_{n+1} awards him the item.
- Thus, VA_{n+1} has expected revenue at least that of every auction that always allocates the item, including $\mathcal{A}$. Therefore

$$\mathbb{E}_{v_1,\dots,v_{n+1}\sim F}\left[\mathrm{Rev}(VA_{n+1})\right] \geq \mathbb{E}_{v_1,\dots,v_{n+1}\sim F}\left[\mathrm{Rev}(\mathcal{A})\right] = \mathbb{E}_{v_1,\dots,v_n\sim F}\left[\mathrm{Rev}(OPT_{F,n})\right]$$

□

Knapsack auctions

Going back to social surplus

- Let us go back to mechanisms that maximize social surplus.
- We know what awesome mechanisms are in single-item auctions. We also have Myerson's lemma for designing DSIC mechanisms.
- Can we design an awesome mechanism for every single-parameter environment?

- An awesome mechanism satisfies the following three properties:
  - DSIC: everybody has dominant strategy "bid truthfully" which guarantees non-negative utility,
  - strong performance: we maximize social surplus if everybody bids truthfully,
  - computational efficiency: the mechanism runs in polynomial time.

Auction example: scheduling TV commercials

- Bidders are companies such that each company has its own TV commercial of length w_i and is willing to pay v_i in order to have the commercial presented during a commercial break. The seller is a television station with a commercial break of length W.

Sources: https://mountain.com/ and https://www.eq-international.com/

- Can we design an awesome mechanism that assigns the slots?

Knapsack auction

- We formalize this auction as follows.
- In a knapsack auction of n bidders $1, \ldots, n$, each bidder i has publicly known size $w_i \geq 0$ and a private valuation $v_i \geq 0$. There is a single seller who has a capacity $W \geq 0$. The feasible set X consists of $\{0, 1\}$-vectors $(x_1, \ldots, x_n)$ such that $\sum_{i=1}^{n} x_i w_i \leq W$, where $x_i = 1$ indicates that bidder i is a winning bidder.

- We now try to design an awesome mechanism for knapsack auctions.
  - For bids $b = (b_1, \ldots, b_n)$, we choose $x(b)$ from X such that $\sum_{i=1}^{n} b_i x_i$ is maximized. Then, when bidders bid truthfully, the social surplus is maximized.
  - The allocation rule x is monotone (one-step function with breakingpoint at some z) and thus Myerson's lemma gives us a payment rule p such that (x, p) is DSIC. If $b_i < z$, then bidder i pays nothing, otherwise he pays $z \cdot (1 - 0) = z$.
  - So we have the first two conditions satisfied. However, the third one will be problematic since x solves the Knapsack problem.

Knapsack problem

- given a capacity W and n items of values $v_1, \ldots, v_n$ and sizes $w_1, \ldots, w_n$, find a subset of the items having a maximum total value such that the total size is at most W.

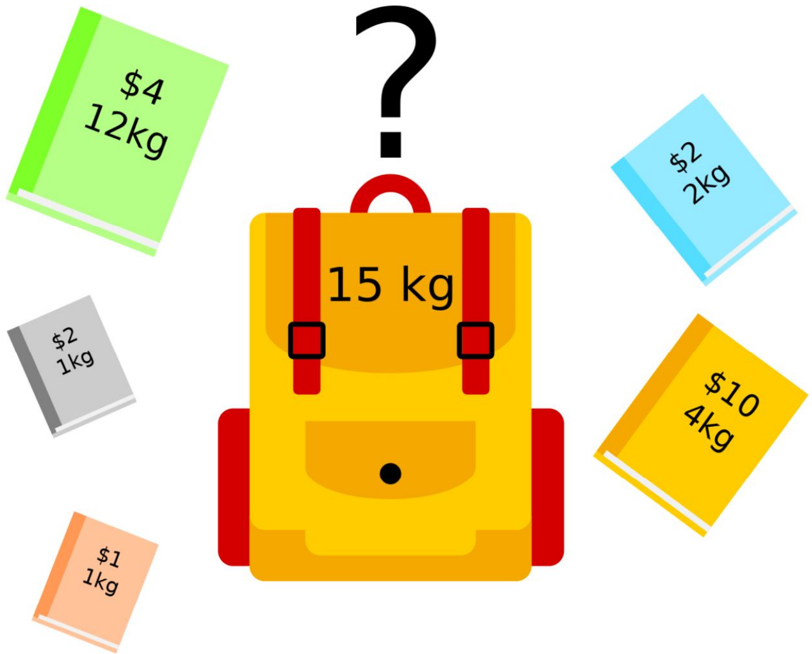

Sources: https://en.wikipedia.org and https://twitter.com/

- This problem is NP-hard.
- There is a pseudo-polynomial time algorithm using dynamic programming and a fully polynomial-time approximation scheme.

We are not always awesome

- Assuming P $\neq$ NP, we cannot satisfy the third condition (polynomial time), since the Knapsack problem is NP-hard, and choosing the allocation rule so that the social surplus is maximized solves it. Thus, knapsack auctions are not awesome.
- Relaxing the first condition (DSIC property) does not help, since it is the last two conditions that collide. We might relax the third condition, say to pseudopolynomial time. This is helpful if our instances are small or structured enough and we have enough time and computing power to implement optimal surplus-maximization.
- The dominant paradigm is to relax the second constraint (optimal surplus) as little as possible, subject to the first (DSIC) and the third (polynomial-time) constraints.
- Myerson's Lemma implies that the following goal is equivalent: design a polynomial-time and monotone allocation rule that comes as close as possible to maximizing the social surplus.

Approximation with monotonicity

- This resembles the primary goal in approximation: design algorithms for NP-hard problems that are as close to optimal as possible, subject to a polynomial-time constraint.
- For us, the algorithms must additionally obey a monotonicity constraint.
- The "holy grail" in algorithmic mechanism design: for as many NP-hard problems as possible, match the best-known approximation guarantee for (not necessarily monotone) approximate surplus maximization algorithms, subject to $P \neq NP$. That is, we would like the DSIC/monotone constraint to cause no additional surplus loss.

- We now illustrate this approach by designing an allocation rule that gives at least half of the optimum social surplus in knapsack auctions.

Greedy allocation for knapsack auctions

- We assume without loss of generality that no bidder i has $w_i > W$. We also assume that the bidders $1, \ldots, n$ are sorted in the order $<$ so that

$$\frac{b_1}{w_1} \geq \cdots \geq \frac{b_n}{w_n}.$$

- Consider the following greedy allocation rule $x^G = (x_1^G, \ldots, x_n^G) \in X$, which for given bids $b = (b_1, \ldots, b_n)$ selects a subset of bidders so that $\sum_{i=1}^n x_i^G w_i \leq W$ using the following procedure.
  - Pick winners in the order $<$ until one does not fit and then halt.
  - Return either the solution from the first step or the highest bidder, whichever creates more social surplus.
- The reason for the second step is that the solution in the first step might be highly suboptimal if there is a very valuable and very large bidder. Consider, for example, $n = 2$ with $b_1 = 2$, $w_1 = 1$, $b_2 = W$, and $w_2 = W$ for a very large W.
- The rule x_G is monotone (Exercise).

2-approximation for knapsack auctions I

Theorem 3.10

Assuming truthful bids, the social surplus of the greedy allocation rule x^G is at least one half of the maximum possible social surplus.

- Proof (sketch): Let $w_1, \ldots, w_n$ be the given sizes, $v_1, \ldots, v_n$ the valuations (and also the bids), and W be the capacity.
- First, we consider a relaxation of the problem, where we can choose each bidder i with fraction $\alpha_i \in [0, 1]$ so that i contributes with $\alpha_i \cdot v_i$ to the solution. The greedy algorithm to solve this fractional version: pick winners in the order $<$ until the capacity W is fully used with the possibility to pick the last winner fractionally, if needed.
- We show that this algorithm maximizes the surplus over all feasible solutions to the fractional knapsack problem.
  - Let $1, \ldots, k$ be the winners selected by the greedy algorithm and suppose for contradiction that there is another feasible solution that gives higher social surplus.

2-approximation for knapsack auctions II

- ◦ Then our solution can be improved by changing one of the constraints α_i to some larger β_i. Since $\alpha_1 = \cdots = \alpha_{k-1} = 1$, we have $i \geq k$. Since $\sum_{l=1}^{k} \alpha_l w_l = W$, there is a $j \in \{1, \ldots, k\}$ with $j < i$ and $\beta_j < \alpha_j$. We can assume that these are the only changed coefficients (Exercise).
- ◦ Then $(\beta_i - \alpha_i)w_i \leq (\alpha_j - \beta_j)w_j$, as $\sum_{l=1}^{k} \alpha_l w_l = W$ and we add $(\beta_i - \alpha_i)w_i$ to the size, while removing $(\alpha_j - \beta_j)w_j$.
- ◦ On the other hand, since the social surplus is now larger, we have $(\beta_i - \alpha_i)v_i > (\alpha_j - \beta_j)v_j$. By dividing the left side of the second inequality with the left side of the first inequality and doing the same for the right sides, we obtain $v_i/w_i > v_j/w_j$, which, since $j < i$, contradicts our choice of the order $<$.

- Now, assume that in the fractional setting the first $k-1$ winners i have $\alpha_i = 1$ while $\alpha_k < 1$. Then the social surplus achieved by step 1 is exactly $\sum_{i=1}^{k-1} \alpha_i v_i = \sum_{i=1}^{k-1} v_i$. The social surplus in step 2 is at least v_k.
- Thus, we have social surplus at least $\max\{v_k, \sum_{i=1}^{k-1} v_i\}$, at least half of the optimal fractional solution, which is at least the non-fractional optimum. □

Multi-parameter mechanism design

Multi-parameter mechanism design

- In single-parameter mechanism design, each bidder has only single piece of private information—his valuation v_i. We consider more general model where bidders have different private valuations for different items.

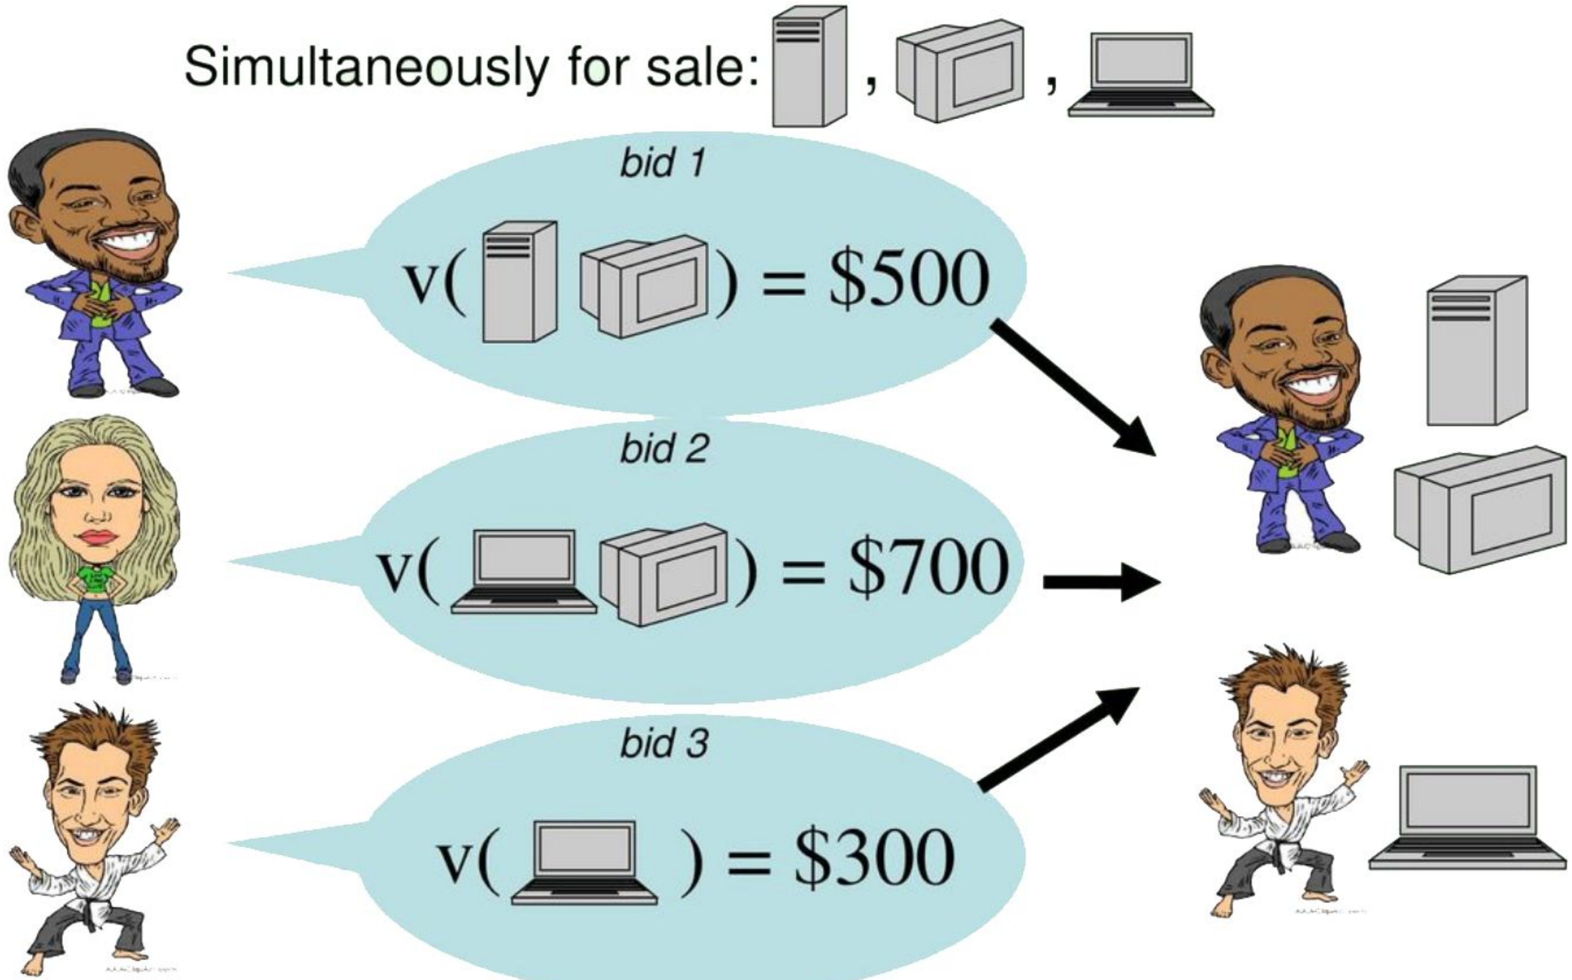

Sources: Auctions & Combinatorial auctions (Vincent Conitzer)

Multi-parameter mechanism design formally

- 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.

Multi-parameter mechanism: example

- Consider two bidders 1 and 2 in an auction in which the seller sells two items t_1 and t_2 and each bidder is allowed to obtain only one item.
- The set of possible outcomes is

$$\Omega = \{(0,0),(0,1),(0,2),(1,0),(2,0),(1,2),(2,1)\},$$

  where an outcome (i,j) corresponds to bidder 1 receiving the item t_i (or nothing if $i=0$) and bidder 2 receiving the item t_j (or nothing if $j=0$).
- The valuations of bidder 1 are $v_1(i,j)=0$ if $i=0$, $v_1(i,j)=10$ if $i=1$, and $v_1(i,j)=5$ if $i=2$. Similarly, the valuations of bidder 2 are $v_2(i,j)=0$ if $j=0$, $v_2(i,j)=5$ if $j=1$, and $v_2(i,j)=3$ if $j=2$.
- Thus, we see that both bidders prefer item t_1 and also want to buy at least something.
- The optimal allocation is to sell t_1 to bidder 1 and t_2 to bidder 2, which leads to the social surplus $10+3=13$.
- We now state a cornerstone of mechanism design, the VCG mechanism. It states that we can still do DSIC social surplus maximization.

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.

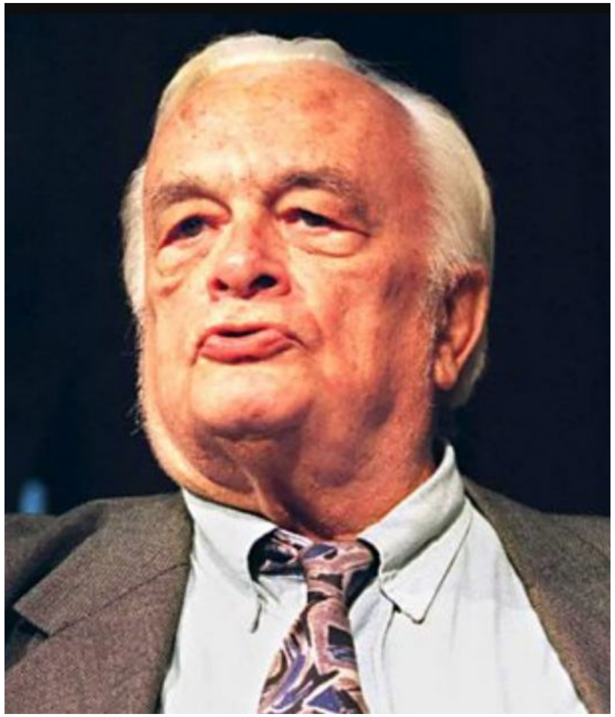

Figure: William Vickrey, Edward H. Clarke, and Theodore Groves.

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

- We postpone the proof of this result to the last lecture.

VCG mechanism: remarks

- 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.

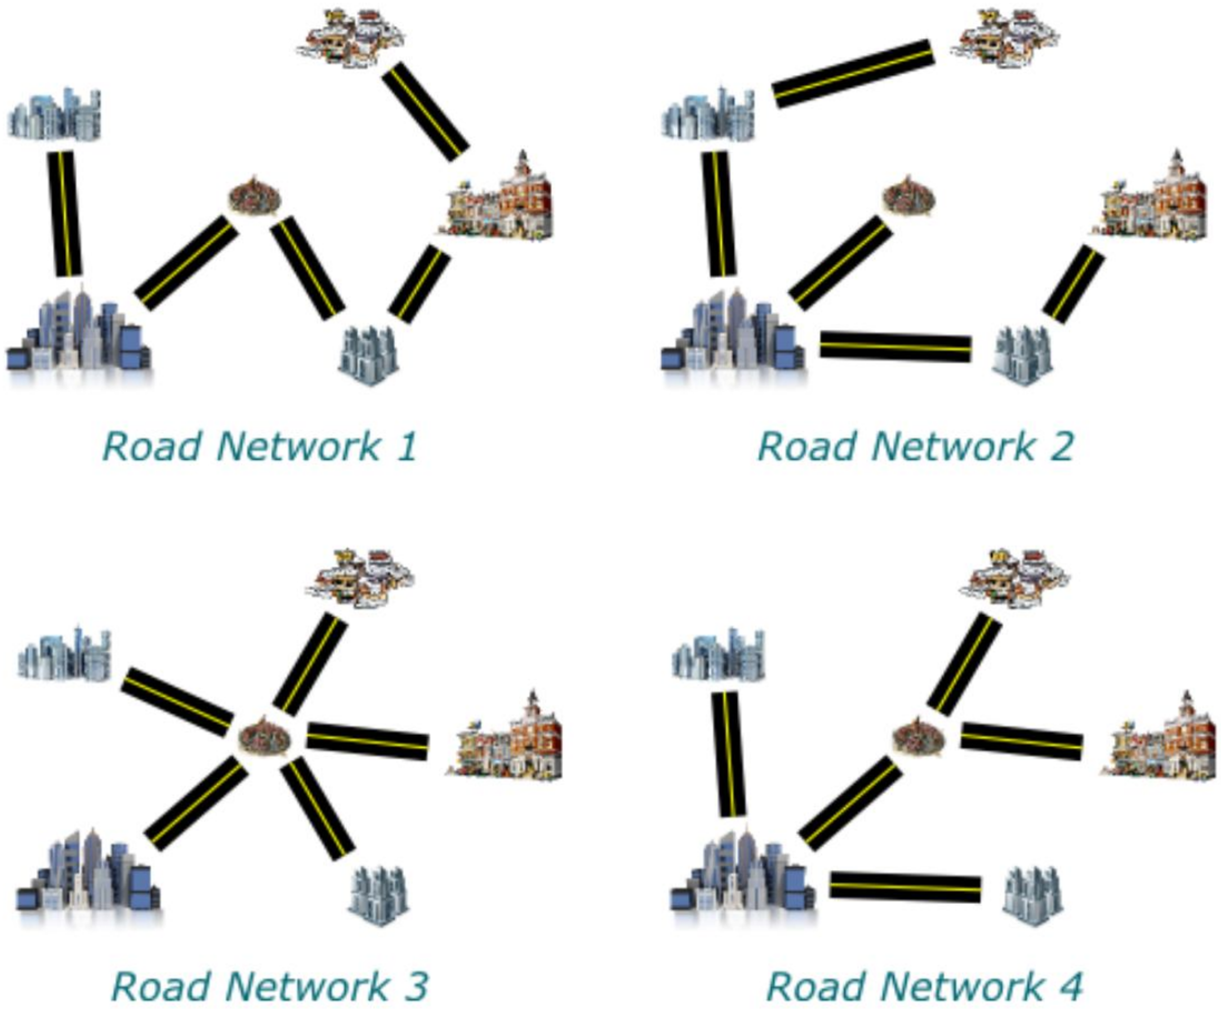

Sources: https://www.science4all.org/article/auction-design/

VCG auction example

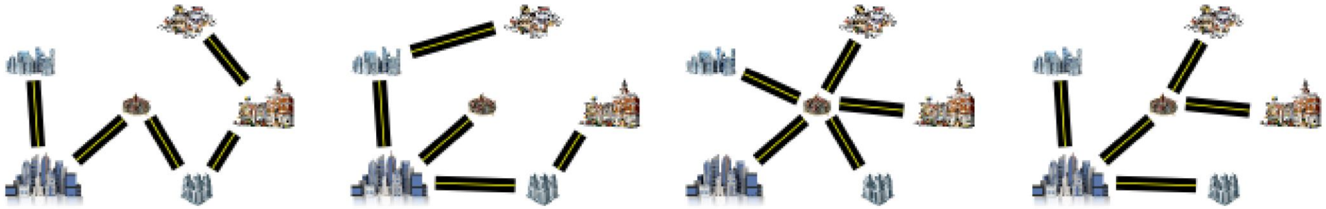

	Road Network 1	Road Network 2	Road Network 3	Road Network 4
	6 M$	14 M$	2 M$	**16 M$**
	5 M$	8 M$	4 M$	**12 M$**
	2 M$	1 M$	**20 M$**	4 M$
	4 M$	**6 M$**	3 M$	5 M$
	1 M$	1 M$	**6 M$**	2 M$
	1 M$	2 M$	2 M$	**3 M$**
Total (social welfare)	19 M$	32 M$	37 M$	**42 M$**

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

- 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$.

Source: https://www.kindpng.com/

Thank you for your attention and merry Christmas!