

The VCG Mechanism

Lecture 8

For the last three lectures, we considered social choice functions with ordinal preferences: Those are rules to select an alternative given the ranking of every voter. We now generalize this and consider rules that select an alternative based on the voters' (now called players) cardinal utilities while being able to require the players to make payments. We will see that using payments will allow us to circumvent the issue of strategic misreporting.

1 Mechanisms

Definition 1 (Mechanism with Payments). Let $N = \{1, \dots, n\}$ be the set of players and let A be the set of alternatives. A *valuation profile* $\mathbf{v} = (v_1, \dots, v_n)$ gives for each player $i \in N$ a *valuation function* $v_i : A \rightarrow \mathbb{R}$. We say that the value of alternative x for player i is $v_i(x)$.

A (*direct revelation*) *mechanism* $M = (f, p)$ takes as input a valuation profile \mathbf{v} and returns an alternative $f(\mathbf{v}) \in A$ and payments $\mathbf{p}(\mathbf{v}) \in \mathbb{R}^n$, where $p_i(\mathbf{v})$ is the payment player i has to make.

Players have *quasi-linear* utilities: If alternative $x \in A$ is selected and player i 's payment is p_i , their utility from the outcome is

$$u_i(x, p_i) = v_i(x) - p_i.$$

Mechanisms and voting rules are different: The input of a voting rule is a preference profile (σ , where σ_i is the ranking of alternatives for voter i), while the input of a mechanism is a valuation profile (\mathbf{v} where v_i are the values, numbers in \mathbb{R} , of player i for the alternatives). Both voting rules and mechanisms output an alternative ($f(\sigma)$ and $f(\mathbf{v})$); mechanisms also output a payment $p_i(\mathbf{v})$ for each player i .

In that sense, social choice functions are the special case of mechanisms that don't use payments ($\mathbf{p} \equiv 0$) and that only make use of the order of the $v_i(x)$ for each player i , not their numerical value.

Mechanisms also generalize auctions:

Example 1 (Auctioning a single item). Let's consider an auction with a single item being up for sale and n players bidding for it. Each player has a value v_i^* for the item.

The possible outcomes, i.e., the alternatives, are $A = \{a_1, \dots, a_n\}$, where a_i denotes the scenario in which $i \in N$ gets the item. Thus, the valuation function of player i is

$$v_i(x) = \begin{cases} v_i^*, & x = a_i, \\ 0, & x \neq a_i. \end{cases}$$

Each player makes a bid $b_i \in \mathbb{R}$ for the item. A mechanism now is a rule that dictates, given bids b_1, \dots, b_n , which player gets the item and how much each player has to pay.

We give two examples, *first-price auctions* and *Vickrey auctions*¹ (often also called *second-price auctions*). In both, the highest bidder gets the item:

$$f(\mathbf{v}) = a_i \text{ where } i = \arg \max_j b_j.$$

In a first-price auction, the player that receives the item pays their bid, while all other players pay nothing:

$$p_i(\mathbf{v}) = b_i, \quad p_j(\mathbf{v}) = 0 \text{ for } j \neq i,$$

In a Vickrey auction, the player that receives the item now pays the *second highest* bid, while all other players still pay nothing:

$$p_i(\mathbf{v}) = \max_{j \neq i} b_j, \quad p_j(\mathbf{v}) = 0 \text{ for } j \neq i.$$

¹William Vickrey (1914 - 1996) was a professor of economics at Columbia University. He is one of only four individuals to receive the Nobel Prize posthumously.

In many settings, such as auctions, it is only natural to assume that players will act strategically to maximize their own utility. Thus, we may ask if we can expect players to report their valuation function v_i truthfully (for example, in the single-item auction setting, this corresponds to bidding their actual value, $b_i = v_i^*$). To do that, we need to extend our definition of strategyproofness to mechanisms.

Definition 2 (Strategyproofness). A mechanism $M = (f, p)$ is *strategyproof* if for any possible valuation profile \mathbf{v} , for all players $i \in N$ and for any possible deviations of this player, v'_i , it holds that

$$u_i(f(\mathbf{v}), p_i(\mathbf{v})) \geq u_i(f(v'_i, \mathbf{v}_{-i}), p_i(v'_i, \mathbf{v}_{-i})).$$

In other words, there exists no scenario in which a player can increase their utility under the mechanism by misreporting their valuation function: Truthfully reporting is a dominant strategy.

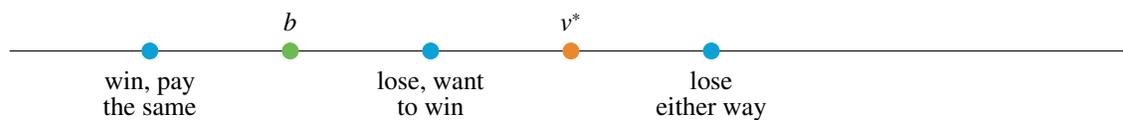
This is an analogous definition of strategyproofness for social choice functions, where v'_i corresponds to a false ranking and u_i corresponds to the voters' satisfaction with the outcome (how high this alternative is in their ranking).

We now return to [Example 1](#) to see if the mechanisms considered are strategyproof. First, we can see that the first-price auction is not strategyproof: In any scenario where there is a unique highest bid, the corresponding player can benefit from decreasing their bid to be just above the second highest bid. They still win the item, and their utility increases since their payment decreases. This is not possible in second-price auctions:

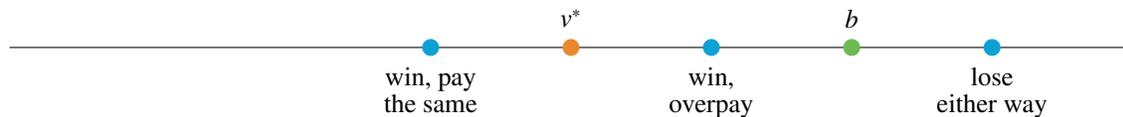
Theorem 1. *The Vickrey auction is strategyproof.*

Proof. Consider some player and let v^* be their valuation. Let b' be the highest bid (and valuation) that is not their bid. We want to show that bidding $b = v^*$ is an optimal strategy for the player; i.e., that it is never worse than any bid $b \neq v^*$.

First, let's assume that the player underbids, $b < v^*$. There are three possible outcomes based on b' : If $b' > v^*$, the player loses both if they bid b and when they bid v^* , so their utility is the same. If $b < b' \leq v^*$,² the agent loses if they bid b receiving utility 0, but wins if they bid v^* receiving utility $v^* - b' \geq 0$. If $b' \leq b$, the player wins and receives a utility of $v^* - b'$ both if they bid b or v^* . Thus, the utility of the player never increases due to bidding $b < v^*$ instead of v^* , so it is not beneficial to misreport in this way.



Now, let's assume that the player overbids, $b > v^*$. Again, there are three possible outcomes based on b' : If $b' > b$, the player loses both if they bid b and when they bid v^* , so their utility is the same. If $b \geq b' > v^*$, the agent wins if they bid b receiving utility $v^* - b' \leq 0$, but loses if they bid v^* receiving utility 0. If $b' \leq v^*$, the player wins and receives a utility of $v^* - b'$ both if they bid b or v^* . Thus, the utility of the player never increases due to bidding $b > v^*$ instead of v^* , so it is not beneficial to misreport in this way.



In conclusion, no player can ever benefit from making a bid different than their true valuation. Thus, playing $b = v^*$ is a dominant strategy; the Vickrey auction is strategyproof. □

2 The VCG mechanism

We have seen that Vickrey auctions are a very elegant mechanism for auctioning off a single item that incentivizes bidders to report their true valuation. The idea underlying Vickrey auctions can be generalized to apply to the general mechanism setting.

²We assume here that ties are broken in favor of the agent. The proof works analogously for any other reasonable tiebreaking rule.

Definition 3 (VCG Mechanism). The *Vickrey-Clarke-Groves (VCG) mechanism* is defined by

- the welfare-maximizing choice rule

$$f(v) \in \arg \max_{x \in A} \sum_{i \in N} v_i(x)$$

- and a payment rule

$$p_i(v) = \max_{x \in A^{-i}} \sum_{j \neq i} v_j(x) - \sum_{j \neq i} v_j(f(v)),$$

where A^{-i} is the set of alternatives that are available when player i is not present.

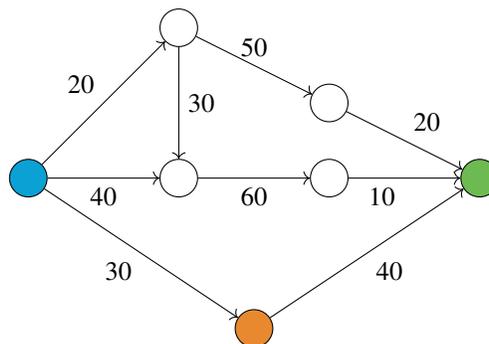
The payment rule states that each player has to pay the ‘harm’ they cause to all other agents. In particular, the first term is the social welfare the other players could achieve if player i doesn’t exist; the second term is the social welfare the other bidders achieve under the outcome $f(v)$ (i.e., if i does exist).

Thus, this is a generalization of the second-price auction described before: For all players but the highest bidder, they pay nothing since their (non)existence doesn’t change the social welfare among the other players. For the highest bidder, the first term is the second highest bid (the social welfare if they didn’t exist), and the second term is 0 (the social welfare of all players but them if they get the item), so they pay exactly the second highest bid. Let’s look at an example of this:

Example 2 (VCG mechanism for auctions). Consider an auction with a single item and two players. Player 1 has value $v_1^* = 7$ for the item, and player 2 has value $v_2^* = 3$. According to the VCG mechanism, player 1 gets the item since they have the highest valuation. If player 1 didn’t exist, player 2 would receive the item, getting utility of 3, while currently their utility is 0 since player 1 receives it. Thus, $p_1(v) = 3 - 0 = 3$. If player 2 didn’t exist, player 1 would still get the item, so $p_2(v) = 7 - 7 = 0$.

Indeed, the VCG mechanism does not only apply to single-item auctions but to a much larger class of problems:

Example 3 (VCG mechanism for routing). In the following graph, we need to construct a path from the blue node to the green node. Each edge is a player, and if the path uses their edge, the player has a cost corresponding to the number next to the edge.



Thus, any path from the blue node to the green node is an alternative for the mechanism. The valuation of a player for an alternative is the negative cost of the player’s edge if it is part of the path, else 0. Here A^{-i} are all paths that don’t include the edge associated with player i .

The VCG mechanism will choose the alternative that maximizes the social welfare, which in this case corresponds to the lowest-cost path, from the blue to the orange to the green node for a total cost of 70.

For the payments, let’s first consider the player corresponding to the edge from the blue to the orange node. Without this edge, the minimum weighted path from blue to green (the alternative that incurs the least social cost) has cost $20 + 50 + 20 = 90$. With the blue-orange edge, the cost incurred by *other* agents is just 40 (the edge from the orange to the green node). Thus, the payment made by the blue-orange player is $(-90) - (-40) = -50$. This player gets paid 50. Similarly, we can see that the orange-green player pays $(-90) - (-30) = -60$, i.e., receives 60. All other players pay 0, since the social welfare (negative total cost of the chosen alternative) both with and without them is -70 .

Theorem 2. *The VCG mechanism is strategyproof.*

Proof. Recall that:

$$f(\mathbf{v}) \in \arg \max_{x \in A} \sum_{i \in N} v_i(x),$$

and the utility of player i is:

$$\begin{aligned} u_i(f(\mathbf{v}), p_i(\mathbf{v})) &= v_i(f(\mathbf{v})) - p_i(\mathbf{v}) \\ &= v_i(f(\mathbf{v})) - \left[\max_{x' \in A^{-i}} \sum_{j \neq i} v_j(x') - \sum_{j \neq i} v_j(f(\mathbf{v})) \right] \\ &= \sum_{j \in N} v_j(f(\mathbf{v})) - \max_{x' \in A^{-i}} \sum_{j \neq i} v_j(x') \end{aligned}$$

Now, consider a misreport to v'_i . The second term above will be unchanged since it is independent of player i ; it is the optimal social welfare that the players can achieve without i . The first term, however, changes to

$$\sum_{j \in N} v_j(f(v'_i, \mathbf{v}_{-i})) \leq \max_{x \in A} \sum_{j \in N} v_j(x) = \sum_{j \in N} v_j(f(\mathbf{v})),$$

implying that the utility of player i cannot increase from misreporting. Thus, reporting the true valuation is a dominant strategy; the VCG mechanism is strategyproof. \square

The VCG mechanism has a second important property: It is individually rational, meaning no player is better off not participating in the mechanism at all.

Definition 4 (Individually Rational Mechanisms). For a mechanism $M = (f, p)$, denote by $f(\mathbf{v}_{-i}) \in A^{-i}$ the outcome of the mechanism when i isn't present. The mechanism M is *individually rational* if

$$u_i(f(\mathbf{v}), p_i(\mathbf{v})) \geq u_i(f(\mathbf{v}_{-i}), 0).$$

Here, $u_i(f(\mathbf{v}), p_i(\mathbf{v}))$ corresponds to the utility of player i when participating in the mechanism and $u_i(f(\mathbf{v}_{-i}), 0)$ corresponds to the utility of player i when not participating in the mechanism.

Theorem 3. *The VCG mechanism is individually rational.*

Proof. The difference $u_i(f(\mathbf{v}), p_i(\mathbf{v})) - u_i(f(\mathbf{v}_{-i}), 0)$ is

$$\begin{aligned} &v_i(f(\mathbf{v})) - \left[\sum_{j \neq i} v_j(f(\mathbf{v}_{-i})) - \sum_{j \neq i} v_j(f(\mathbf{v})) \right] - v_i(f(\mathbf{v}_{-i})) \\ &= \sum_{j \in N} v_j(f(\mathbf{v})) - \sum_{j \in N} v_j(f(\mathbf{v}_{-i})) \\ &= \max_{x \in A} \sum_{j \in N} v_j(x) - \sum_{j \in N} v_j(f(\mathbf{v}_{-i})) \\ &\geq 0, \end{aligned}$$

where the inequality follows from the fact that the maximum social welfare of $f(\mathbf{v})$ is, by definition, at least as high as the social welfare of any other alternative. \square

3 Sponsored search

A real-world application of the VCG mechanism are the ranking of sponsored websites in online search. When you search Google, for example to buy something, often the first few results are sponsored. How far up a companies' website will be placed in this sponsored ad section is determined by an auction. We give a model of these position auctions and explore how the VCG mechanism can be applied to them.

We assume that in these auctions, each company submits a bid-per-click b_i , corresponding to their (estimated) value for a single person clicking on the link to visit their homepage. Early mechanisms for sponsored searches ranked ads based on their bid-per-click. They then determined prices in a first-price auction manner, meaning that company i would pay exactly b_i for every click received on their ad. One big issue with this setup, though, is that it does not take into account how frequently a link will get clicked. In particular, if a company with a low-quality product (receiving little clicks even if ranked in first place) submits a slightly higher bid than another company with a high-quality product (receiving, say, 10 times as many clicks when ranked first), the low-quality company will get displayed first even though the total value (clicks time bid-per-click) would be a lot higher if the high-quality company was ranked first.

Because of this, we turn to a more general model of sponsored search that also considers the predicted number of clicks for every ad.

Definition 5 (Position Auctions). A *position auction* consists of

- a set of positions $M = \{1, \dots, n\}$, ordered from top (1) to bottom (n), with positions $m + 1, \dots, n$ considered “unassigned” and $m + 1 \leq n$,
- a value-per-click w_i for each player i (which is reported, potentially untruthfully), and
- a predicted click-through rate CTR_{ij} corresponding to the probability in which the ad from player i will be clicked if it is in position j .

The value of player i from their ad being displayed in position j is $v_{ij} = CTR_{ij} \cdot w_i$, their expected value from clicks if they are in position j .

Here, we assume that click-through rates are separable into a position score and a quality score, meaning that $CTR_{ij} = \text{pos}_j \cdot Q_i$, where $\text{pos}_j \in [0, 1]$ is the probability that a user sees position j and $Q_i \in [0, 1]$ is the probability that the user clicks on the ad i if they see it (the quality of the ad). We assume that $\text{pos}_1 \geq \dots \geq \text{pos}_m$ and $\text{pos}_{m+1} = \dots = \text{pos}_n = 0$.

An outcome of the election is an assignment $\sigma : N \rightarrow M$ of the players to positions. Thus, the possible alternatives of the mechanism are all possible assignments. The goal is to maximize the social welfare, i.e. the total value created,

$$\max_{\sigma \in A} \sum_{i \in N} CTR_{i, \sigma(i)} \cdot w_i$$

Theorem 4. *The assignment that ranks the ads according to $w_i \cdot Q_i$ is optimal.*

Proof. Let’s assume towards a contradiction that this is not the case, i.e., that there exists players i and j such that $Q_i \cdot w_i > Q_j \cdot w_j$ but player i is placed in position k and player j is placed in position k' where $k > k'$. We wish to show that the assignment with these two players swapped has at least the total value of the current assignment. First, note that the total value of the other bids outside of i or j is the same across the two assignments because their positions don’t change. Thus, it suffices to show that the increase in value for bidder i from moving from k to k' is at least the decrease in value for bidder j from moving from k' to k . That is

$$Q_i \cdot w_i \cdot (\text{pos}_k - \text{pos}_{k'}) + Q_j \cdot w_j \cdot (\text{pos}_{k'} - \text{pos}_k) \geq 0,$$

which follows since $Q_i \cdot w_i > Q_j \cdot w_j$ and $\text{pos}_k \geq \text{pos}_{k'}$. □

This allows us to state what the VCG mechanism applied to position auctions looks like.

Definition 6 (The VCG Position Auction). The *VCG position auction* returns the assignment σ that ranks players by $w_i Q_i$ (thus maximizing the social welfare). Player i pays

$$\sum_{k=\sigma(i)+1}^{m+1} (\text{pos}_{k-1} - \text{pos}_k) \cdot Q_{\sigma^{-1}(k)} \cdot b_{\sigma^{-1}(k)},$$

the total gain in value all players ranked below i would make by moving one position up (thus, the loss in social welfare among all other players due to i existing).

Example 4 (VCG Position Auction). Consider $m = 3$ with position effects

$$\text{pos}_1 = 0.2,$$

$$\text{pos}_2 = 0.18,$$

$$\text{pos}_3 = 0.1.$$

Further, $n = 4$, all ads have $Q_i = 1$, and the (reported) values-per-click are

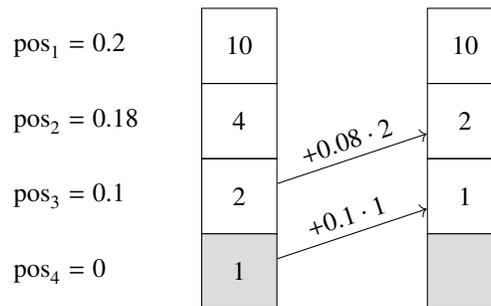
$$w_1 = 10,$$

$$w_2 = 4,$$

$$w_3 = 2,$$

$$w_4 = 1.$$

First note that since the optimal ranking is assigning position in order of $Q_i \cdot v_i = v_i$, player 1 will get position 1, player 2 will get position 2, etc. To illustrate the payments, let's calculate the payment of player 2. By definition, it will be the marginal harm they cause to the other players with their bid. Note that player 1 is unaffected by player 2's bid, so we can disregard them. We see that without player 2, player 3 with value-per-click $w_3 = 2$ will move up to position 2 and so their value will increase from $2 \cdot 0.1$ to $2 \cdot 0.18$. Player 4 with value-per-click $w_4 = 1$ will move up to position 3 so their value will increase from $1 \cdot 0$ to $1 \cdot 0.1$.



Thus, player 2 has to make a payment of

$$(0.18 - 0.1) \cdot 2 + (0.1 - 0) \cdot 1 = 0.26.$$

Facebook, X (then Twitter) and Google Ads (then Adwords) have all used VCG position auctions at some point for their sponsored search. However, a practical shortcoming of the VCG position auction is transparency. Under VCG, especially if the click-through rates are predicted with complex methods, it is difficult for the advertising companies to predict their payments. These platforms still adapt their mechanism frequently, with their choices being informed by ongoing research on mechanism design.