

Mechanism Design II: more VCG and the Revelation Principle

1 Mechanism design, general setup (recap)

We assume we have n players, and a set of “alternatives” A (we will also call them outcomes or allocations) such as who gets all the various items. Each player i has a valuation function $v_i : A \rightarrow \mathbb{R}$. These can be arbitrary (e.g., you can have items worth more together than separately like a printer and ink, and you can even have higher value on allocations that give more items to your friends and lower value on allocations that give items to your enemies). The one assumption we will make is that players’ utilities are *quasilinear*: The utility for player i of allocation a , paying p_i is $u_i = v_i(a) - p_i$.

2 VCG with Clarke pivot rule

We ended last time describing VCG with the Clarke pivot rule, the standard version of VCG.

Given a vector v of valuation functions,

- Let $f(v) = \operatorname{argmax}_{a \in A} \sum_j v_j(a)$ be the allocation that maximizes social welfare.
- Let $p_i(v) = \max_a \left[\sum_{j \neq i} v_j(a) \right] - \sum_{j \neq i} v_j(f(v)) = \sum_{j \neq i} v_j(f(v_{-i})) - \sum_{j \neq i} v_j(f(v))$.

In other words, you charge each player i an amount equal to how much less happy they make everyone else (how much they reduce everyone else’s total social welfare) by causing $f(v)$ to be chosen rather than $f(v_{-i})$. This is often called *charging them their externality*.

In addition to incentive-compatibility and maximizing social welfare, this mechanism has the following two additional nice properties:

1. The auctioneer never pays the bidder
2. Assuming the v_i ’s themselves are non-negative, no player ever gets negative utility. This is called *ex-post individual rationality*. For example, in an auction of goods where people’s valuations depend only on what they get, then among other things this implies that people who don’t get anything don’t have to pay anything.

Why does this satisfy $p_i(v) \geq 0$? This is because the first term is the max.

Why does this satisfy individual rationality assuming the v_i functions are non-negative? That's because:

$$\begin{aligned} u_i(f(v), p) &= v_i(f(v)) + \sum_{j \neq i} v_j(f(v)) - \max_a [v_{-i}(a)] \\ &= \max_a [v(a)] - \max_a [v_{-i}(a)]. \end{aligned}$$

This can't be negative because one option for the first “ a ” is to use the second “ a ”, and v_i itself is non-negative.

3 More examples

3.1 Combinatorial auctions

In a combinatorial auction, we have m items, and each bidder has a valuation function over *subsets* of items. Valuation functions are assumed to be normalized so that for any bidder i , $v_i(\emptyset) = 0$, and one typically also assumes “free disposal”, namely that if $S \subseteq T$ then $v_i(S) \leq v_i(T)$. Two sets S and T are said to be “complements” for bidder i if $v_i(S \cup T) > v_i(S) + v_i(T)$ and are said to be “substitutes” for bidder i if $v_i(S \cup T) < v_i(S) + v_i(T)$. Implicit is an assumption that bidders only care about what they get, and not about what anyone else gets.

For a combinatorial auction, what VCG would do is find the social-welfare-maximizing allocation of items, namely a partition of the items into subsets S_1, S_2, \dots, S_n maximizing $\sum_i v_i(S_i)$, and then will charge each bidder their externality, namely the difference between $\max_{\{S'_1, \dots, S'_n\}} \sum_{j \neq i} v_j(S'_j)$ and $\sum_{j \neq i} v_j(S_j)$. One of the questions on the homework is to work out a specific example with two items and three bidders.

One problem with combinatorial auctions is that implementing VCG can be computationally hard. In fact, there is even the problem of how a general bidder should describe their valuation function to the mechanism. To address this, we'll often assume that bidders can implement certain oracles. For example, if you propose a bundle of items to a bidder, they can say how much it's worth to them (a value query), or if you assign prices to individual items, a bidder can figure out what subset of items they would buy at those prices (a demand query). We will talk about these later. In fact, we'll see that a simple mechanism that involves setting prices to items and having bidders come in one at a time and take what they want at those prices (which is trivially IC) has some nice guarantees in terms of approximately maximizing social welfare.

3.2 Public projects

In the public projects problem, we imagine the government is considering undertaking some public project like building a bridge, or a park, etc. This project has some known cost C , and each person i has some value $v_i \geq 0$ on the project (you can look in the chapter for an extension to the case that people might have negative value on the project). We will think of the government itself as having value $-C$ on building the project because the government has to pay for it (that money goes outside the system). So, to maximize social welfare, the government should build the project if $\sum_i v_i > C$ and should not build the project if $\sum_i v_i < C$.

The VCG mechanism with the Clarke pivot rule will collect the values v_1, \dots, v_n , build the project if their sum is greater than C , and then charge each citizen their externality. Specifically:

1. If $\sum_j v_j < C$ then the project is not built, and nobody is charged anything.
2. If $\sum_j v_j \geq C$ then the project is built and citizens are charged as follows. For each i ,
 - (a) If $\sum_{j \neq i} v_j \geq C$ then citizen i is charged 0.
 - (b) If $\sum_{j \neq i} v_j < C$ then citizen i is charged $C - \sum_{j \neq i} v_j$.

We can see that this charges each citizen their externality. In particular, in case 2(b), the social-welfare-optimizing alternative if citizen i were not there would have been not to build the project.

Notice that that in general, the government does not collect enough money to pay for the project (it might even collect no money at all). It turns out this is unavoidable (see Section 9.5.5).

3.3 Social choice revisited with money

Consider a social choice problem where we want to make a choice among k alternatives, say picking one of k restaurants to order lunch from. What VCG would do is ask each person to list how much they value each of the k alternatives, and then it would sum these numbers up and select the alternative of highest total value. It would then charge each person their externality. For example, suppose a restaurant A was chosen such that the total sum of everyone else's value for that restaurant is \$480, but the restaurant B that maximizes the sum of everyone else's value has a sum of everyone else's value of \$500. Then you would be charged \$20. It must be the case that you valued restaurant A at least \$20 since otherwise B would have been chosen (assuming no negative valuations), so participating is Individually Rational. Again, the total dollar amount charged to everyone might not be enough to pay for the catering order (it could even be that everyone is charged \$0) but money is just being used to make the system incentive compatible.

3.4 Bilateral trade

Suppose you want to facilitate trades between a buyer and a seller. (Think of yourself as helping to move resources around to people who need them the most). E.g., a seller has a car, which has value v_s to them. There is a buyer who might want the car, and it has value v_b to them. To maximize social welfare, what we want is that if $v_b > v_s$, then the car goes to the buyer and the seller gets some money, and if $v_b < v_s$ then the car stays with the seller. Here, the Clarke pivot rule is a little funny because it corresponds to you stealing the car and the auctioning it off to the highest bidder at the second highest bid. (E.g., if $v_s > v_b$ then the seller's externality is v_b so they have to pay v_b). Instead, here the most natural version of VCG is that if $v_b \leq v_s$ then there is no trade and no payments, whereas if $v_b > v_s$ then the buyer gets the car, the buyer pays v_s and the seller gets v_b . E.g., if the buyer values the car at \$1000 and the seller values it at \$700, then the buyer pays \$700 and the seller gets \$1000. (Can you see why this is IC for the seller and why it is IC for the buyer?) So the mechanism is subsidizing the trade. It turns out that you can't avoid that if you want to be IC and maximize social welfare (see Section 9.5.5).

4 The Revelation Principle

VCG is a “direct revelation mechanism” where everyone submits their valuations and the mechanism computes the outcome and payments. You can also have mechanisms where there is a process, like an ascending auction. But if such a mechanism has dominant strategies, then there is also an incentive-compatible direct-revelation mechanism. Why? Just take in the valuations and then act on behalf of the players (which the mechanism knows how to do since it can just play the dominant strategy).

This allows us to focus on single-round mechanisms when understanding what IC mechanisms look like.

5 Characterization of IC direct-revelation mechanisms

Theorem 1 *A direct-revelation mechanism is IC iff it has the following two properties:*

1. *If $f(v_i, v_{-i}) = f(v'_i, v_{-i})$ then $p_i(v_i, v_{-i}) = p_i(v'_i, v_{-i})$. In other words, for any given v_{-i} , you can view it as assigning prices $p_i(a, v_{-i})$ to each alternative $a \in f(., v_{-i})$.*
2. *$f(v_i, v_{-i}) = \arg \max_{a \in f(., v_{-i})} [v_i(a) - p_i(a, v_{-i})]$. In other words, it then chooses the best alternative for you at these prices.*

Proof: The easier direction is: if the mechanism satisfies these two conditions, then it is incentive-compatible. That is because the price for each alternative doesn’t depend on your own reported valuation function, and then the mechanism is optimizing utility on your behalf.

In the other direction, we need to show that if either (1) or (2) is violated, then the mechanism is *not* incentive-compatible. Let’s begin with (1). If (1) is not satisfied, then for some i , some v_i , some v_{-i} and some v'_i , we have $f(v_i, v_{-i}) = f(v'_i, v_{-i})$ but $p_i(v_i, v_{-i}) \neq p_i(v'_i, v_{-i})$. This means that player i with true valuation function equal to whichever of these is larger would prefer to misreport as the valuation for whichever of these is smaller, since that produces the same allocation but at a lower price.

Now, suppose that (1) is satisfied but (2) is not. So, for some i , some v_i , and some v_{-i} , the allocation chosen is different from the argmax allocation a^* . But, since a^* is in the range of $f(., v_{-i})$, player i would prefer to misreport their valuation to be whichever v'_i produces allocation a^* , since that allocation and payment have higher utility for player i . ■