The Vickrey-Clarke-Groves Mechanism

CS 1951k/2951z

2020-02-5

We introduce the Vickrey-Clarke-Groves Mechanism mechanism, a direct auction for multiple goods, and argue that it is DSIC.

1 *Combinatorial Auctions*

In single-parameter environments, there is sometimes a single good, and sometimes more than one good (e.g., sponsored search; forthcoming), but the bidders themselves are always characterized by but one parameter. In multi-parameter environments, there are multiple (indivisible) goods, and the bidders' values for **bundles** (i.e., subsets) of those goods cannot be characterized by a single parameter. Auctions designed for multi-parameter environments are called **combinatorial auctions**, as bidders value combinations of goods.

Analogous to the single-parameter setting, we begin our study of combinatorial auctions by focusing on direct mechanisms. As usual, we are driven by three design goals:

1. incentives: we seek auctions in which reporting bids truthfully is a dominant strategy (i.e., DSIC mechanisms)
2. economic efficiency (i.e., welfare maximization)
3. computational efficiency

As in single-parameter environments, we again adopt the IPV model in the multi-parameter/combinatorial setting. In this case, a bidder i's valuation cannot be expressed as a single number, but is instead a function of the auction's outcome, which in general describes not only i's allocation and payments, but all bidders'. Without other-regarding preferences, however, each bidder i's valuation depends only on i's allocation, among the set of possible allocations $\mathcal{A}$:[1]

$$v_i : \mathcal{A} \to \mathbb{R}, \quad \forall i \in N.$$

[1] We assume the set $\mathcal{A}$ of possible allocations is the same for all bidders, namely 2^G, or the power set of G, meaning the set of all subsets of G. The notation for the power set derives from its cardinality. If the cardinality of a set X is m, then the cardinality of the power set of X is 2^m.

In this lecture, we study a direct mechanism (meaning, a direct combinatorial auction) in which all bidders' reports take the form of valuations, meaning values for all possible bundles of goods. It should be apparent immediately that designing such a mechanism in a computationally efficient manner is a non-trivial endeavor. Even describing a single bidder's valuation—and hence, communicating bids to the auctioneer—takes exponential space and time.

2 *The Vickrey-Clarke-Groves Mechanism*

The standard procedure for designing a DSIC auction is as follows:

1. Find an economically efficient allocation, given bids. In this lecture, we take as our objective welfare maximization: i.e., given bids $\mathbf{b}$, find an allocation ω^* s.t.

$$\omega^* \in \arg\max_{\omega \in \Omega} \sum_{i \in N} b_i(\omega).$$

   This problem is called the **winner determination problem**.

2. Charge each winner an appropriate payment, so as to ensure the DSIC property.

The complexity of the winner determination problem can make it impossible to deploy auctions designed via this procedure in practice. We return to this issue later. For the moment, we seek to explore the theoretically interesting question, what payment formula ensures the DSIC property? The answer is the following:

$$p_i(\omega^*) = h_i(\mathbf{b}_{-i}) - \sum_{j \neq i \in N} b_j(\omega^*)$$

where, as above, ω^* denotes an optimal allocation. A mechanism that uses this payment formula is called a **Groves** mechanism. Before we unpack this formula, let's prove that this mechanism is dominant-strategy incentive compatible.

Theorem 2.1. *The Groves mechanism is DSIC.*

Proof. Fix a bidder i. Given a bid profile $\mathbf{b}_{-i}$, let $f(b_i, \mathbf{b}_{-i})$ compute a welfare-maximizing allocation, assuming i bids b_i: i.e.,

$$f(b_i, \mathbf{b}_{-i}) \in \arg\max_{\omega \in \Omega} \sum_{i \in N} b_i(\omega).$$

The utility of bidder i is

$$\begin{aligned} u_i(b_i, \mathbf{b}_{-i}) &= v_i(f(b_i, \mathbf{b}_{-i})) - \left(h_i(\mathbf{b}_{-i}) - \sum_{j \neq i \in N} b_j(f(b_i, \mathbf{b}_{-i})) \right) \\ &= v_i(f(b_i, \mathbf{b}_{-i})) + \sum_{j \neq i \in N} b_j(f(b_i, \mathbf{b}_{-i})) - h_i(\mathbf{b}_{-i}). \end{aligned}$$

The VCG mechanism selects an allocation $f(b_i, \mathbf{b}_{-i})$ that maximizes

$$b_i(f(b_i, \mathbf{b}_{-i})) + \sum_{j \neq i \in N} b_j(f(b_i, \mathbf{b}_{-i})).$$

Bidder i wants the mechanism to select an allocation $f(b_i, \mathbf{b}_{-i})$ that maximizes

$$v_i(f(b_i, \mathbf{b}_{-i})) + \sum_{j \neq i \in N} b_j(f(b_i, \mathbf{b}_{-i}))$$

In order to ensure the mechanism maximizes this quantity, bidder i must report $b_i = v_i$. □

We will choose

$$h_i(\mathbf{b}_{-i}) = \max_{\omega \in \Omega} \sum_{j \neq i \in N} b_j(\omega)$$

This choice is called the **Clarke pivot rule**.

Auctions that use the aforementioned design process together with the Clarke pivot rule in the payment formula are called **Vickrey-Clarke-Grove (VCG)** auctions.[2] With the Clarke pivot rule and non-negative valuations, the VCG mechanism is individually rational, and never pays bidders to participate (i.e., all payments are non-negative).

[2] VCG auctions are a special case of VCG mechanisms, which apply beyond auctions: e.g., in the realm of public goods provisioning.

While the VCG mechanism is DSIC and IR, it still exhibits some bizarre behavior. You will explore the following anomalies in this week's homework exercises:

1. The VCG mechanism may allocate goods to bidders with strictly positive valuations, and still generate zero revenue.
2. The VCG mechanism may generate less revenue for the auctioneer when an additional bidder participates.
3. The VCG mechanism may generate more utility for bidders who collude to submit untruthful bids. Indeed, bidders can collude with themselves by submitting what are called **false-name** bids!

3 *Interpreting the VCG Payment Formula*

We now describe two ways to interpret the VCG payment formula.

Suppose bidder i were not present. The VCG mechanism would generate welfare

$$\max_{\omega \in \Omega} \sum_{j \neq i \in N} b_j(\omega).$$

When bidder i is present, the set of bidders $N \setminus i$ may not achieve (collectively) as much welfare as when i was not present. With i present, the welfare generated is

$$\omega^* \in \arg\max_{\omega \in \Omega} \sum_{j \in N} b_j(\omega),$$

and the net difference in welfare for the set of bidders in $N \setminus i$ is

$$\max_{\omega \in \Omega} \sum_{j \neq i \in N} b_j(\omega) - \sum_{j \neq i \in N} b_j(\omega^*).$$

This quantity is exactly the payment the VCG mechanism charges bidder i. Thus, we can view bidder i's payment as the **externality** they impose on all the other bidders, collectively.

Example 3.1. Assume a single-good auction, in which bidder i has the highest valuation. At the optimal allocation ω^*, $\sum_{j \neq i \in N} b_j(\omega^*) = 0$, so the VCG mechanism charges the winner $\max_{\omega \in \Omega} \sum_{j \neq i \in N} b_j(\omega)$, which (not coincidentally!) is the second-highest bid. Thus, the second-price, sealed-bid auction (a.k.a., the Vickrey auction) is the VCG mechanism assuming only one good.

Another way of interpreting the payment formula is in terms of rebates. Expand the payment formula as follows:

$$\begin{aligned} p_i(\omega^*) &= \max_{\omega \in \Omega} \sum_{j \neq i \in N} b_j(\omega) - \sum_{j \neq i \in N} b_j(\omega^*) \\ &= \max_{\omega \in \Omega} \sum_{j \neq i \in N} b_j(\omega) - \sum_{j \neq i \in N} b_j(\omega^*) + b_i(\omega^*) - b_i(\omega^*) \\ &= b_i(\omega^*) - \left[\sum_{j \in N} b_j(\omega^*) - \max_{\omega \in \Omega} \sum_{j \neq i \in N} b_j(\omega) \right]. \end{aligned}$$

We see that bidder i's payments are precisely their bid less a non-negative quantity. Moreover, this quantity can be understood as the amount of additional welfare that can be attributed to i's presence: i.e., the value of the welfare-maximizing allocation with i less the value of the welfare-maximizing allocation without i.

Example 3.2. In a Vickrey auction for a single good, suppose, without loss of generality, bidder 1 wins, and bidder 2's bid is the second highest. Bidder 1 pays

$$p_1(\omega^*) = b_1 - (b_1 - b_2) = b_2.$$

We now present an example of how the VCG mechanism operates with three bidders and two goods.

Example 3.3. Suppose there are three bidders and two goods, A and B, with valuations as described in Table 1.

Bundle	v_1	v_2	v_3
$\emptyset$	0	0	0
$\{A\}$	2	1	1
$\{B\}$	2	1	1
$\{A, B\}$	2	1	4

Table 1: Bidder valuations for winning subset of goods in $G = \{A, B\}$

The welfare-maximizing allocation ω^* gives bundle $\{A, B\}$ to bidder 3, and results in total welfare of 4.

1. If bidder 1 did not exist, the maximum possible welfare would be 4. Bidder 1's contribution to welfare is zero. Payments for bidder 1 are

$$p_1(\omega^*) = 4 - 4 = 0 - (4 - 4) = 0.$$

2. If bidder 2 did not exist, the maximum possible welfare would be 4. Bidder 2's contribution to welfare is zero. Payments for bidder 2 are
$$p_2(\omega^*) = 4 - 4 = 0 - (4 - 4) = 0.$$
3. If bidder 3 did not exist, the maximum possible welfare would be 3. We could allocate A to bidder 1 and B to bidder 2 (or vice versa). Payments for bidder 3 are
$$p_3(\omega^*) = 3 - 0 = 4 - (4 - 3) = 3.$$

4 *The Winner Determination Problem*

Implementing the VCG mechanism requires solving the winner determination problem (i.e., solving for an optimal allocation) $n + 1$ times. This is a serious impediment to its widespread use, because solving the winner determination problem turns out to be an NP-hard problem. Nonetheless, it is a problem of great practical importance; as such, it has been studied empirically,[3] like other NP-hard problems: e.g., satisfiability and the travelling salesperson problem.

The solution to the winner determination problem can be formulated as a constrained optimization problem:

1. Objective function: Maximize welfare.
2. Constraints:
   (a) Do not allocate any good to more than one bidder.
   (b) Do not allocate to any bidder more than one bundle.[4]

Let $x_{i,S}$ indicate whether bidder i is allocated bundle S or not:
$$x_{i,S} = \begin{cases} 1, & \text{if } i \text{ is allocated bundle } S \\ 0, & \text{otherwise.} \end{cases}$$
Finding an optimal allocation can then be expressed as an integer linear program. Given truthful reports $\mathbf{v}$, the primal form is:
$$\begin{aligned} \max_{\mathbf{x}} \quad & \sum_{i \in N} \sum_{S \subseteq G} x_{i,S} v_i(S) \\ \text{subject to} \quad & \sum_{i \in N} \sum_{S \subseteq G : j \in S} x_{i,S} \leq 1, && \forall j \in G \\ & \sum_{S \subseteq G} x_{i,S} \leq 1, && \forall i \in N \\ & x_{i,S} \in \{0, 1\}, && \forall i \in N, \forall S \subseteq G \end{aligned}$$
This problem is computationally expensive because there are an exponential number of variables, but it would be difficult to solve regardless, because integer linear programming is NP-hard.

[3] Kevin Leyton-Brown, Mark Pearson, and Yoav Shoham. Towards a universal test suite for combinatorial auction algorithms. In *Proceedings of the 2nd ACM conference on Electronic commerce*, pages 66–76. ACM, 2000

[4] This set of constraints is necessary because the value of bundle $S_1 \cup S_2$ may not equal the value of S_1 plus the value of S_2, and maximizing the sum total of these values is the objective.

Remark 4.1. Not all instances of the winner determination problem are NP-hard, even in the multi-parameter case. For example, when bidders' valuations are additive, so that the value of a set of goods is the sum total of their individual values: i.e., for all $i \in N$,

$$v_i(S) = \sum_{j \in S} v_i(j)$$

the winner determination problem can be solved in polynomial time.

Now, suppose we relax the integrality constraints, so that allocations are now in $[0, 1]$ instead of $\{0, 1\}$. With a bit of work, we can show that the dual of the relaxed LP is:

$$\begin{aligned} \min_{\mathbf{u}, \mathbf{p}} \quad & \sum_{i \in N} u_i + \sum_{j \in G} p_j \\ \text{subject to} \quad & u_i \geq v_i(S) - \sum_{j \in S} p_j, && \forall i \in N, \forall S \subseteq G \\ & u_i \geq 0, && \forall i \in N \\ & p_j \geq 0, && \forall j \in G. \end{aligned}$$

We have intentionally used the variable names u and p, as the dual linear program can be interpreted as solving for utility and payments, respectively. While the number of variables in the dual is only polynomial in size, the number of constraints is exponential.

A The Primal and the Dual

In matrix form, given $\mathbf{c} \in \mathbb{R}^n$, $\mathbf{A} \in \mathbb{R}^{m \times n}$, $\mathbf{b} \in \mathbb{R}^m$, $\mathbf{x} \in \mathbb{R}^n$, and $\mathbf{y} \in \mathbb{R}^m$, the relaxed primal can be written as

$$\begin{aligned} \max_{\mathbf{x}} \quad & \mathbf{c}^T \mathbf{x} \\ \text{subject to} \quad & A\mathbf{x} \leq \mathbf{b} \\ & \mathbf{x} \geq 0 \end{aligned}$$

In matrix form, the dual can be written as

$$\begin{aligned} \min_{\mathbf{y}} \quad & \mathbf{b}^T \mathbf{y} \\ \text{subject to} \quad & A^T \mathbf{y} \geq \mathbf{c} \\ & \mathbf{y} \geq 0 \end{aligned}$$

References

[1] Kevin Leyton-Brown, Mark Pearson, and Yoav Shoham. Towards a universal test suite for combinatorial auction algorithms. In *Proceedings of the 2nd ACM conference on Electronic commerce*, pages 66–76. ACM, 2000.