Uniqueness of the VCG Mechanism

CSCI 1951k/2951z

2020-04-01

We prove that Groves mechanism is the *unique* DSIC mechanism, up to the choice of h_i. Further, the Vickrey-Clarke-Groves—the Groves mechanism where h_i is the Clarke pivot rule—is the unique DSIC and IR mechanism that does not pay to induce participation.

1 *Theorem and Assumptions*

In a single-parameter environment, a bidder's type can be described by a single value. In contrast, a **multi-parameter** environment is defined as follows:

- a finite set Ω of feasible allocations
- for each bidder $i \in N$, a private valuation, or type, $v_i : \Omega \to \mathbb{R}$

As Ω is assumed to be finite, we can likewise express bidder i's type T_i as $\mathbb{R}^{\Omega}$. In other words, v_i can be understood to be a vector that assigns a value to each feasible allocation.

Recall the Groves mechanism, a direct mechanism which collects reports $b_i : \Omega \to \mathbb{R}$ from each bidder $i \in N$, and is defined by the allocation rule:

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

and the payment rule:

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

where

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

for some function $h_i : T_{-i} \to \mathbb{R}$. Recall further that the Groves mechanism is DSIC.[1] Perhaps surprisingly, the Groves mechanism is the *unique* DSIC mechanism, up to the choice of h_i. A proof of this claim (Theorem 1.1) is the subject of these lecture notes.

The Clarke pivot rule, which defines the Vickrey-Clarke-Groves (VCG) mechanism, namely

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

was dictated by the IR constraint, together with the additional desideratum that participants should not be paid to play. It follows from Theorem 1.1 that the VCG mechanism is the unique DSIC mechanism that satisfies these additional design goals.

---

[1] The term DSIC generally refers only to direct mechanisms. Hence, we drop the qualifier "direct" when in reference to the mechanisms discussed in these notes; it is implied.

Theorem 1.1. *In a multi-parameter setting, if the set of possible types T_i is connected, for all bidders $i \in N$, then the Groves mechanism is the unique DSIC mechanism, up to an additive shift.*

Charging a bidder a payment that is independent of their own bid, as is done by the Groves mechanism, makes for a DSIC mechanism, since doing so aligns the bidders' incentives with the mechanism's (welfare maximization). In other words, charging bidders in this way is a sufficient condition for the DSIC property to hold.

Theorem 1.1 is a statement about the converse; it is a necessary condition. If a mechanism is DSIC, then it must be that each bidder pays an amount independent of their own bid.

Before proving Theorem 1.1, we define **connected**, and show by counterexample that it is a necessary assumption.

Example 1.2. Informally, we say that a type space $T_i \subseteq \mathbb{R}^{\Omega}$ is connected if, geometrically, it consists of no more than one solid object in $\mathbb{R}^{\Omega}$. We show by counterexample that if T_i is not connected, then DSIC payments may not be unique.

Consider, for example, a single-good auction in which $T_i = \mathbb{Z}^+$, for all bidders i. Define x to be the allocation rule in which the highest bidder wins the good. If p is the second-price payment rule, then (x, p) is a DSIC mechanism.

We argue that the alternative mechanism (x, p') is *also* DSIC, where p' is defined as follows:

- If there is a unique highest bidder, they pay the second highest bid plus 1/2.
- In the event that multiple bidders tie for first place, a winner is arbitrarily selected, and they pay their bid.

Assume $b^* \in \mathbb{Z}^+$ is the highest bid among all bidders other than bidder i. If bidder i's value for the good is less than b^*, then they prefer to lose the good. They can achieve this objective by bidding their value. Similarly, if their value is greater than $b^* + 1/2$, then they prefer to win the good, which again, they can achieve by bidding their value. Finally, consider the interval $[b^*, b^* + 1/2]$. Since the type space is restricted to integers, the only possible value on this interval is b^*. Bidding this true value leads to a tie, and utility 0; bidding less than b^* also yields utility 0; finally, bidding more than b^* yields negative utility. In sum, bidding truthfully maximizes utility in the (x, p') mechanism; hence, it is DSIC.

As the choice of $1/2 \in (0, 1)$ was arbitrary, payments are clearly[2] not unique. A unique winning bidder could just have well have paid the second-highest price plus 1/3 or 1/4, etc.

[2] It is terrible practice to write "clearly" in a research paper, and worse still, in lecture notes. Things are never crysyal clear to everyone, and will I even see this result so clearly myself when I return to it next year?

2 *Plan for the Proof*

We now set out to prove Theorem 1.1, the uniqueness of Groves payments. Let x be an allocation rule and let p be a payment rule such that (x, p) is DSIC. We will show that given x, this choice of p is unique, across *all* DSIC mechanisms, up to an additive shift.

Fix some bidder i and a value profile $\mathbf{v}_{-i}$. We abbreviate $x_i(\cdot, \mathbf{v}_{-i})$ and $p_i(\cdot, \mathbf{v}_{-i})$, by $x(\cdot)$ and $p(\cdot)$, respectively. In other words, these functions refer to the space of allocation and corresponding payment rules in terms of bidder i's report/type v_i.

Since (x, p) is DSIC, for any two types v and v' (for bidder i), $x(v) = x(v')$ implies $p(v) = p(v')$. Otherwise, if $p(v') < p(v)$, then i would be incentivized to report v', even if v were i's true type. So fixing an allocation rule, bidder i's payment rule as a function of its *type*, is (almost[3]) well defined. We thus write p_a hereafter to denote bidder i's corresponding payment, fixing allocation $a \in \Omega$.

The next bit of machinery we need is the notion of closeness.

Definition 2.1. Two allocations $a, b \in \Omega$ are called **close** if, for all $\epsilon > 0$, there exist types v and v' such that $x(v) = a$, $x(v') = b$, and $\|v' - v\|_\infty \leq \epsilon$, where $\|v\|_\infty = \max_{a \in \Omega} |v(a)|$.[4]

Consider the following graph:

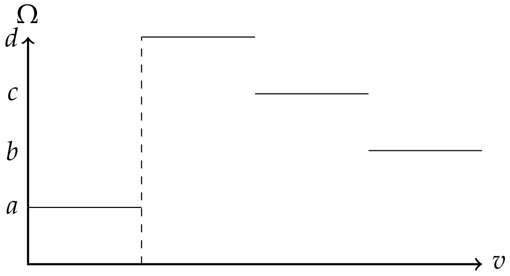

Figure 1: An allocation function $x(v)$.

Allocations a and d are close because we can find two types that are arbitrarily close to each other and that correspond to the two different allocations (i.e., choose types arbitrarily close to each side of the dotted line). Allocations a and b are not close because for small values of ϵ, we cannot find any two types within ϵ of each other that correspond to the two different allocations.

The proof of Theorem 1.1 follows from two key lemmas:

1. If the type space is connected, then any two allocations can be linked by a chain of pairwise-close allocations (Lemma 3.1).

   For example, in Figure 1, b is reachable from a via the chain a, d, c, b of pairwise-close allocations.

---

[3] up to the function h_i

[4] Recall that v, v' are vectors of dimension $|\Omega|$. The notation $v(a)$ refers to the value of the allocation a according to v.

2. If two allocations $a, b \in \Omega$ are close, then their corresponding payments p_a and p_b differ by an additive constant. Moreover, this constant is the *same* across all DSIC mechanisms (Lemma 3.3).

Putting these two lemmas together, we conclude that the payments p_a and p_b corresponding to any two allocations $a, b \in \Omega$ differ by the same additive constant in all DSIC mechanisms. In other words, all DSIC mechanisms differ from one another by an additive shift. This shift is determined by their respective choices of h_i.

For example, consider the VCG mechanism, which relies on the Clarke pivot rule h_i^c, and what we might call the VCG+1 mechanism, which relies on the alternative pivot rule $h_i^{+1}(\mathbf{b}_{-i}) = h_i^c(\mathbf{b}_{-i}) - 1$, for all $\mathbf{b}_{-i} \in T_{-i}$. The payments in VCG+1 and VCG differ by an additive shift. The former pays every bidder 1 to participate.

3 Key Lemmas

We prove these lemmas for one-dimensional type spaces only (i.e., in the single-parameter case, assuming $T_i \subseteq \mathbb{R}$). In the multi-dimensional case, a proof of Lemma 3.1 invokes concepts from topology, while the proof of Lemma 3.3 is more notation heavy, but conceptually similar.

Lemma 3.1. *If T_i is connected, then for any two allocations $a, b \in \Omega$, there exists a chain of allocations $a = a_1, a_2, a_3, \ldots, a_m = b$ such that every consecutive pair a_k, a_{k+1} in this chain is close.*

Proof. Let $a, b \in \Omega$ and let $x(v) = a$ and $x(v') = b$, for some $v, v' \in T_i \subseteq \mathbb{R}$. WLOG, assume $v < v'$. Since T_i is connected, $[v, v'] \subseteq T_i$, and so there is a corresponding allocation for every $v_0 \in [v, v']$, each one close to the last. Since Ω is finite, these allocations form a finite sequence beginning with a and ending with b. $\square$

Remark 3.2. We already argued that p_a is well-defined [5]: i.e., the payment p_a associated with allocation $a \in \Omega$ is such that for all v satisfying $x(v) = a$, $p(v) = p_a$. Why does this result not automatically prove Lemma 3.3? The reason is, well-definedness only implies that in a single DSIC payment scheme, say p^1, that p_a^1 and p_b^1 differ by an additive constant. But it is conceivable that there is another DSIC payment scheme, say p^2, in which p_a^2 and p_b^2 differ by some other additive constant. Lemma 3.3 establishes that p_a and p_b differ by the same additive constant in *all* DSIC payment schemes.

Lemma 3.3. *If $a, b \in \Omega$ are close, then $|p_a - p_b|$ is the same additive constant in all DSIC mechanisms.*

[5] up to h_i

Proof. Assume $a, b \in \Omega$ are close, so that for all $\epsilon > 0$, there exist types v_a and v_b with $x(v_a) = a$ and $x(v_b) = b$ s.t. $\|v_b - v_a\|_\infty < \epsilon$. We invoke the DSIC property to characterize $|p_b - p_a|$.

Let $u_i(z; v)$ denote bidder i's utility from reporting z when its type is v. By the DSIC property, it must be the case that $u_i(v_a; v_a) \geq u_i(v_b; v_a)$ and $u_i(v_b; v_b) \geq u_i(v_a; v_b)$. Equivalently,

$$\begin{aligned} v_a(a) - p_a &\geq v_a(b) - p_b \\ v_b(b) - p_b &\geq v_b(a) - p_a. \end{aligned}$$

We rearrange the terms in these inequalities to form lower and upper bounds, respectively, on $p_b - p_a$:

$$p_b - p_a \geq v_a(b) - v_a(a) \tag{1}$$

$$v_b(b) - v_b(a) \geq p_b - p_a \tag{2}$$

In other words,

$$p_b - p_a \in [v_a(b) - v_a(a), v_b(b) - v_b(a)] \,.$$

By the triangle inequality, and because a and b are close, for all $\epsilon > 0$,

$$\begin{aligned} (v_b(b) - v_b(a)) - (v_a(b) - v_a(a)) &= (v_b(b) - v_a(b)) + (v_a(a) - v_b(a)) \\ &\leq |v_b(b) - v_a(b)| + |v_a(a) - v_b(a)| \\ &\leq 2\epsilon. \end{aligned}$$

Letting $\epsilon \to 0$, it follows that $|p_b - p_a|$ is an additive constant. As we assumed nothing about the mechanism other than that it is DSIC, this property holds of all DSIC mechanisms. □

Example 3.4. Imagine two allocations l and w, where in the former, the bidder loses, and in the latter, they win. Imagine further that these two allocations are close, so that, for example, the losing valuation $v_l = (0, 10 - \epsilon)$ and the winning valuation $v_w = (0, 10)$. Then:

$$\begin{aligned} p_w - p_l &\in [v_l(w) - v_l(l), v_w(w) - v_w(l)] \\ &= [(10 - \epsilon) - 0, 10 - 0] \\ &= [10 - \epsilon, 10] \end{aligned}$$

In the limit as $\epsilon \to 0$, $p_b - p_a = 10$.

Theorem 1.1 states that incentive compatibility dictates only the difference between the payments associated with one allocation or another; it does not pin them down. Fixing the payment for any one type in a DSIC mechanism, however, fully determines the payments for all others (i.e., pins them down). After fixing the payments in two DSIC mechanisms, the difference between all corresponding payments is a "shift." All DSIC mechanisms are equivalent up to this shift, which is dictated by the choice of h_i.