24 Mechanisms with Payments

The Gibbard-Satterthwaite Theorem assumes that agents can have arbitrary preferences over the set of alternatives, and in particular does not apply in settings where the outcome selected by the mechanism includes monetary payments and the utility of each agent is *quasilinear*, i.e., a linear combination of a valuation for the alternative selected by the social choice function and the amount of money transfered to or from the agent. It is worth noting that this assumption makes utilities comparable across agents.

In cases where the outcome includes monetary payments, it will be instructive to separate these payments from the social choice and write a mechanism as a pair (f, p) of a social choice function $f : \Theta \to A$ and a payment function $p : \Theta \to \mathbb{R}^n$. The utility of agent i can then be written as $u_i(\theta', \theta_i) = \nu_i(f(\theta'), \theta_i) - p_i(\theta')$, where θ' is a profile of types revealed to the mechanism, θ_i is the true type of agent i, $\nu_i : A \times \Theta_i \to \mathbb{R}$ is a valuation function over alternatives, and $p_i(\theta') = (p(\theta'))_i$. The main result for the quasilinear setting is positive and provides a way to optimize the most natural social choice function, the one that maximizes social welfare. The *social welfare* of an alternative $a \in A$ is $\sum_{i \in N} \nu_i(a, \theta_i)$, i.e., the sum of all agents' valuations for this alternative.

24.1 Vickrey-Clark-Groves Mechanisms

The mechanisms implementing this social choice function are the so-called Vickrey-Clark-Groves (VCG) mechanisms. A mechanisms (f, p) is a *Vickrey-Clark-Groves mechanism* if

$$f(\theta) \in \arg\max_{a \in A} \sum_{i \in N} \nu_i(a, \theta_i) \qquad \text{and}$$

$$p_i(\theta) = h_i(\theta_{-i}) - \sum_{j \in N \setminus \{i\}} \nu_j(f(\theta), \theta_j) \qquad \text{for all } i \in N,$$

where $h_i : \Theta_{-i} \to \mathbb{R}$ is some function that depends on the types of all agents but i. The crucial component is the term $\sum_{j \in N \setminus \{i\}} \nu_j(f(\theta), \theta_j)$, which is equal to the social welfare for all agents but i. The utility of agent i adds its own valuation $\nu_i(f(\theta), \theta_i)$ and thus becomes equal to the social welfare of alternative $f(\theta)$ minus the term $h_i(\theta_{-i})$. The latter does not depend on θ_i and therefore has no strategic implications.

Theorem 24.1. *VCG mechanisms are strategyproof.*

Proof. Let $i \in N$, $\theta \in \Theta$, and $\theta'_i \in \Theta_i$. Then,

$$u_i(\theta, \theta_i) = \nu_i(f(\theta), \theta_i) - p_i(\theta)$$

$$\begin{aligned} &= \sum_{j\in N} v_j(f(\theta),\theta_j) - h_i(\theta_{-i}) \\ &\geqslant \sum_{j\in N} v_j(f(\theta_i',\theta_{-i}),\theta_j) - h_i(\theta_{-i}) \\ &= u_i((\theta_i',\theta_{-i}),\theta_i), \end{aligned}$$

where the inequality holds because $f(\theta)$ maximizes social welfare with respect to θ. □

Strategyproofness holds for any choice of the functions h_i, so it is natural to ask for a good way to define these functions. In many cases it makes sense that agents are charged rather than paid, but not more than their gain from participating in the mechanism. Formally, mechanism (f,p) makes *no positive transfers* if $p_i(\theta) \geqslant 0$ for all $i \in N$ and $\theta \in \Theta$, and is *ex-post individually rational* if it always yields non-negative utility for all agents, i.e., if $v_i(f(\theta)) - p_i(\theta) \geqslant 0$ for all $i \in N$ and $\theta \in \Theta$. It turns out that these two properties can indeed be achieved simultaneously. The so-called *Clark pivot rule* sets $h_i(\theta_{-i}) = \max_{a\in A} \sum_{j\in N\setminus\{i\}} v_j(a,\theta_j)$, such that the payment of agent i becomes $p_i(\theta) = \max_{a\in A} \sum_{j\in N\setminus\{i\}} v_j(a,\theta_j) - \sum_{j\in N\setminus\{i\}} v_j(f(\theta))$. Intuitively, this latter amount is equal to the externality agent i imposes on the other agents, i.e., the difference between their social welfare with and without i's participation. The payment makes the agent internalize this externality.

THEOREM 24.2. *A VCG mechanism with Clarke pivot rule makes no positive transfers. If $v_i(a,\theta_i) \geqslant 0$ for all $i \in N$, $\theta_i \in \Theta_i$, and $a \in A$, it also is individually rational.*

Proof. Fix $\theta \in \Theta$ and $i \in N$, and let $a = f(\theta)$ and $b \in \arg\max_{a'\in A} \sum_{j\in N\setminus\{i\}} v_j(a',\theta_j)$. Then, by choice of b, $p_i(\theta) = \sum_{j\in N\setminus\{i\}} v_j(b,\theta_j) - \sum_{j\in N\setminus\{i\}} v_j(a,\theta_j) \geqslant 0$, so the mechanism makes no positive transfers. Moreover,

$$\begin{aligned} u_i(\theta,\theta_i) &= v_i(a,\theta_i) + \sum_{j\in N\setminus\{i\}} v_j(a,\theta_j) - \sum_{j\in N\setminus\{i\}} v_j(b,\theta_j) \\ &\geqslant \sum_{j\in N} v_j(a,\theta_j) - \sum_{j\in N} v_j(b,\theta_j) \geqslant 0, \end{aligned}$$

where the two inequalities respectively hold because $v_i(b,\theta_i) \geqslant 0$ and by choice of a. □

Consider for example the application of the VCG mechanism with Clarke pivot rule to an auction of a single good. In this case $A = N$, and the valuation function can be written as $v_i : A \to \mathbb{R}$, such that $v_i(a)$ is equal to agent i's valuation for the good if $a = i$ and zero otherwise. Since only a single agent can receive the good, $\max_{a\in A} \sum_{i\in N} v_i(a) = \max_{i\in N} v_i(i)$, and thus $f(\theta) \in \arg\max_{i\in N} v_i(i)$. Moreover, $p_i(\theta) = \max_{a\in A} \sum_{j\in N\setminus\{i\}} v_j(a) - \sum_{j\in N\setminus\{i\}} v_j(f(\theta))$. The first term is equal to $\max_{j\in N\setminus\{i\}} v_j(j)$ if $a = i$, the second term is zero if $a = i$ and equal to the first term otherwise, and thus $p_i(\theta) = \max_{j\in N\setminus\{i\}} v_j(j)$ if $f(\theta) = i$ and $p_i(\theta) = 0$ otherwise. We thus obtain the well-know Vickrey (or second-price) auction, which assigns the good to the agent with the highest bid and charges this agent a payment equal to the second-highest bid.

24.2 Characterizations of Strategyproof Mechanisms

One might wonder whether other objectives can be implemented in the quasilinear setting besides maximization of social welfare. Two characterizations of strategyproof mechanisms (f, p) exist in the literature. The first characterization states that a mechanism is strategyproof if and only if the payment of an agent is independent of its reported type and the chosen outcome simultaneously maximizes the utility of all agents.

THEOREM 24.3. *A mechanism* (f, p) *is strategyproof if and only if for every* $i \in N$ *and* $\theta \in \Theta$*,*

$$p_i(\theta) = t_i(\theta_{-i}, f(\theta)) \qquad \textit{and}$$
$$f(\theta) \in \arg\max_{a \in A(\theta_{-i})} (v_i(\theta_i, a) - t_i(\theta_{-i}, a)),$$

where $t_i : \Theta_{-i} \times A \to \mathbb{R}$ *is a* price function *and* $A(\theta_{-i}) = \{f(\theta_i, \theta_{-i}) : \theta_i \in \Theta_i\}$ *is the range of* f *given that the reported types of all agents but* i *are fixed to* θ_{-i}*.*

Alternatively, strategyproof mechanisms can be characterized purely in terms of their social choice function. SCF f satisfies *weak monotonicity* if for all $\theta \in \Theta$, $i \in N$, and $\theta'_i \in \Theta_i$, $f(\theta) = a \neq b = f(\theta_i, \theta_{-i})$ implies that $v_i(a, \theta_i) - v_i(b, \theta_i) \geqslant v_i(a, \theta'_i) - v_i(b, \theta'_i)$. Intuitively, an SCF is weakly monotone if a change in the social choice due to a change of type of a single agent means that the agent's value for the new choice must have increased relative to its value for the old choice.

THEOREM 24.4. *If mechanism* (f, p) *is strategyproof, then* f *satisfies weak monotonicity. If SCF* f *satisfies weak monotonicity and for each* $i \in N$*,* $\{(v_i(a, \theta_i))_{a \in A} : \theta_i \in \Theta_i\} \subseteq \mathbb{R}^{|A|}$ *is a convex set, then there exists a payment function* $p : \Theta \to \mathbb{R}^n$ *such that* (f, p) *is strategyproof.*

This result reduces the characterization of strategyproof mechanisms to one of weakly monotone social choice function. The problem with the latter is that weak monotonicity is a local condition that is hard to check, and existence of a global condition depends on the domain of possible preferences. Good global conditions are known to exist for two extreme cases: domains that are unrestricted in the sense that the utilities an agent assigns to the alternatives in A can be arbitrary vectors in $\mathbb{R}^{|A|}$, and domains that are essentially one-dimensional.

A closer look at the unrestricted case reveals that the only strategyproof mechanisms are simple variations of VCG mechanisms, which allow for the assignment of weights to agents and alternatives and for restrictions of the range. SCF f is called an *affine* *maximizer* if there exist $A' \subseteq A$, $w_i \in \mathbb{R}_{>0}$ for $i \in N$, and $c_a \in \mathbb{R}$ for $a \in A'$ such that for every $\theta \in \Theta$, $f(\theta) \in \arg\max_{a \in A'}(c_a + \sum_{i \in N} w_i v_i(a, \theta_i))$. It is easy to see that VCG mechanisms can be generalized to affine maximizers.

THEOREM 24.5. *Let* f *be an affine maximizer, and for each* $i \in N$ *and* $\theta \in \Theta$*, let* $p_i(\theta) = h_i(\theta_{-i}) - \sum_{j \in N \setminus \{i\}} (w_j/w_i) v_j(f(\theta), \theta_j) - c_{f(\theta)}/w_i$*, where* $h_i : \Theta_{-i} \to \mathbb{R}$*. Then* (f, p) *is strategyproof.*

Proof. The utility of agent $i \in N$ is

$$u_i(\theta, \theta_i') = v_i(f(\theta), \theta_i') - h_i(\theta_{-i}) + \sum_{j \in N \setminus \{i\}} (w_j/w_i) v_j(f(\theta), \theta_j) + c_{f(\theta)}/w_i.$$

By adding $h_i(\theta_{-i})$, which does not depend on θ_i, and multiplying by w_i, we see that $u_i(\theta, \theta_i')$ can be maximized by maximizing $c_{f(\theta)} + \sum_{j \in N} w_j v_j(f(\theta), \theta_j')$. This happens when $\theta_i = \theta_i'$. □

When there are at least three alternatives and preferences are unrestricted, affine maximizers are the only strategyproof mechanisms.

THEOREM 24.6 (Roberts, 1979). *Let* $|A| \geqslant 3$ *and* $\{(v_i(a, \theta_i))_{a \in A} : \theta \in \Theta\} = \mathbb{R}^{|A|}$ *for every* $i \in N$*. Let* $f : \theta \to A$ *be a surjective SCF,* $p : \Theta \to \mathbb{R}^n$ *a payment function. If* (f, p) *is strategyproof, then* f *is an affine maximizer.*