Mechanism Design: Groves Mechanisms and Clarke Tax

(Based on Shoham and Leyton-Brown (2008). *Multiagent Systems: Algorithmic, Game-Theoretic, and Logical Foundations*, Cambridge.)

Leen-Kiat Soh

Grove Mechanisms

- **Efficiency** (Definition 10.3.6) is often considered to be one of the most important properties for a mechanism to satisfy in the **quasilinear** setting
  - Research has considered the design of mechanisms that are **guaranteed** to select efficient choices when agents follow dominant or equilibrium strategies
- The most important family of efficient mechanisms are the **Groves mechanisms**

Quasilinear Preferences

Preferences that are quasilinear make analysis easier to handle: more flexible, more consistent, more simplistic

- First, we are in a setting in which the mechanism can choose to charge or reward the agents by an **arbitrary monetary amount**
- Second, an agent's degree of preference for the selection of any choice $x \in X$ is **independent** from his or her degree of preference for having to pay the mechanism some amount $p_i \in \mathbb{R}$.
  - Thus **an agent's utility for a choice *cannot* depend on the total amount of money that he or she has** (e.g., an agent cannot value having a yacht more if he/she is rich than if he/she is poor)
- Finally, agents care **only about the choice selected and about their own payments**
  - in particular, they do ***not*** care about the monetary payments **made or received by other agents**

Mechanism Efficiency

- **Definition 10.3.6 Efficiency.** *A quasilinear mechanism is* **strictly Pareto efficient,** *or just* **efficient**, *if in equilibrium it selects a choice x such that* $\forall v \forall x', \sum_i v_i(x) \geq \sum_i v_i(x')$.
  - An agent's *valuation* for choice $\in X$, written $v_i(x)$ should be thought of as the ***maximum amount of money*** that i would be willing to pay to get the mechanism designer to implement choice x

If the mechanism selects x and x is the choice that has the largest sum of all agents' valuation of a choice, then the mechanism is efficient

It does not mean that every agent's top choice is x: some agents might not like x at all.

Definition

Role of this payment?

- **Definition 10.4.1 (Groves mechanisms)** Groves mechanisms *are direct quasilinear mechanisms* $(\chi, \wp)$, *for which*

Social Choice

$$\chi(\hat{v}) = arg\max_{x} \sum_{i} \hat{v}_i(x),$$

Payment by Agent i

$$\wp_i(\hat{v}) = h_i(\hat{v}_{-i}) - \sum_{j \neq i} \hat{v}_j(\chi(\hat{v})).$$

How much all other agents as a whole value the social choice

  - Direct mechanisms in which agents can **declare any valuation function** $\hat{v}$ (may be different from their true valuation function, v)
  - The mechanism then **optimizes** its choice assuming that the agents disclosed their true utility function (arg max)
  - An agent is made to **pay an arbitrary amount $h_i(\hat{v}_{-i})$** which does ***not*** depend on its own declaration and **is paid the sum of every other agent's declared valuation for the mechanism's choice**

Properties

Role of this payment?

- The fact that the mechanism designer has the **freedom to choose the h_i functions** explains why we refer to the *family* of Groves mechanisms rather than to a single mechanism
- **Groves mechanisms provide a dominant strategy truthful implementation of a social-welfare-maximizing social choice function**
- **Theorem 10.4.2** *Truth telling is a dominant strategy under any Groves mechanism*
- Intuitively, the reason that Groves mechanisms are dominant-strategy truthful is that agents' externalities are ***internalized***
  - An agent's utility depends on the selected choice and **imposed payment**
  - Since **increasing the (reported) utility of all the other agents under the chosen allocation will decrease the imposed payment**, each agent is ***motivated*** to maximize the other agent's utilities just like his or her own

How do we set this function?

$$\wp_i(\hat{\boldsymbol{v}}) = \boldsymbol{h}_i(\hat{\boldsymbol{v}}_{-i}) - \sum_{j \neq i} \hat{\boldsymbol{v}}_j(\chi(\hat{\boldsymbol{v}}))$$

The VCG Mechanism (aka Pivot Mechanism)

- **Definition 10.4.4 (Clarke tax)** *The* Clarke tax *sets the h_i term in a Groves mechanism as $h_i(\hat{v}_{-i}) = \sum_{j \neq i} \hat{v}_j(\chi(\hat{v}_{-i}))$, where x is the Groves mechanism allocation function.*
- **Definition 10.4.5 (Vickrey–Clarke–Groves (VCG) mechanism)** *The* VCG mechanism *is a direct quasilinear mechanism $(\chi, \wp)$, where*

Equation same as before

$$\chi(\hat{v}) = arg \max_x \sum_i \hat{v}_i(x),$$

$$\wp_i(\hat{v}) = \sum_{j \neq i} \hat{v}_j(\chi(\hat{v}_{-i})) - \sum_{j \neq i} \hat{v}_j(\chi(\hat{v})).$$

The "Tax"

"rewards" same as before

What are these?

$$\wp_i(\hat{v}) = \sum_{j \neq i} \hat{v}_j(\chi(\hat{v}_{-i})) - \sum_{j \neq i} \hat{v}_j(\chi(\hat{v}))$$

Is it fair to require each agent to pay this amount?

The Clarke tax does *not* depend on an agent i's own declaration $\hat{v}_i$

Payment Rule's Intuition

$$\wp_i(\hat{v}) = \sum_{j \neq i} \hat{v}_j(\chi(\hat{v}_{-i})) - \sum_{j \neq i} \hat{v}_j(\chi(\hat{v}))$$

- Assume that all agents follow their dominant strategies and declare their valuations truthfully
- The second sum in the VCG payment rule **pays** each agent i the sum of every other agent $j \neq i$'s utility for the mechanism's choice
- The first sum **charges** each agent i the sum of every other agent's utility for the choice that ***would have been made*** **had *i* not participated in the mechanism**
- Thus, each agent is made to pay his or her ***social cost***—the aggregate impact that his or her participation has on other agents' utilities

Payment Rule’s Intuition 2

- If some agent i does not change the mechanism’s choice by his or her participation (i.e., if $\chi(v) = \chi(v_{-i})$), then the two sums will cancel out
  - **The social cost of *i*’s participation is zero**, and so he or she has to pay nothing
- In order for an agent i to be made to pay a nonzero amount, he or she must be ***pivotal*** in the sense that $\chi(v) \neq \chi(v_{-i})$
  - This is why VCG is sometimes called the ***pivot*** mechanism—***only pivotal agents are made to pay***
- It is possible that some agents will ***improve*** other agents’ utilities by participating
  - such agents will be made to pay a ***negative*** amount, or in other words ***will be paid by the mechanism***

$$\wp_i(\hat{v}) = \sum_{j \neq i} \hat{v}_j(\chi(\hat{v}_{-i})) - \boxed{\sum_{j \neq i} \hat{v}_j(\chi(\hat{v}))}$$

If this is greater than the Clarke tax, what happens?

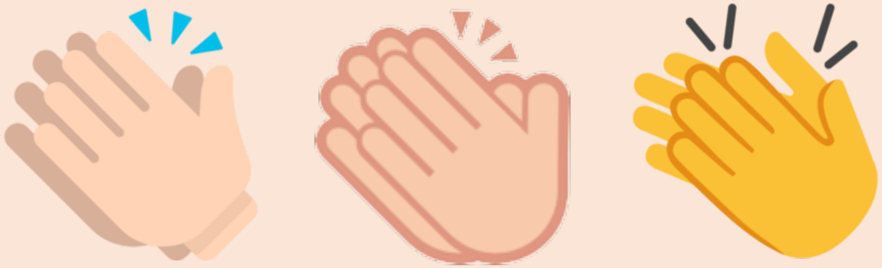

Drawbacks

- **Agents must fully disclose private information** (rationally motivated)
- **Susceptibility to collusion**
- **VCG is not frugal**
- **Dropping bidders can increase revenue**
  - *If we have agents that are not pivotal, then they don't have to pay …*
- **Cannot return all revenue to the agents**
- **Computational intractability**
  - *Evaluating the $argmax$ can require solving an NP-hard problem in many practical domains.*

Connection to MAS?

- Internalizing externalities can help design a mechanism to motivate agents to reveal their true preferences

- Mechanisms can be elegant and powerful for MAS designers, to achieve both local autonomy for agents and desired emergent behavior for the system (Recall our first handout on this tradeoff)