Quasilinear Mechanism Design

CMPUT 654: Modelling Human Strategic Behaviour

S&LB §10.3-10.4

Recap: Revelation Principle

Theorem: (Revelation Principle)
If there exists **any** mechanism that implements a social choice function C in dominant strategies, then there exists a **direct** mechanism that implements C in dominant strategies and is **truthful**.

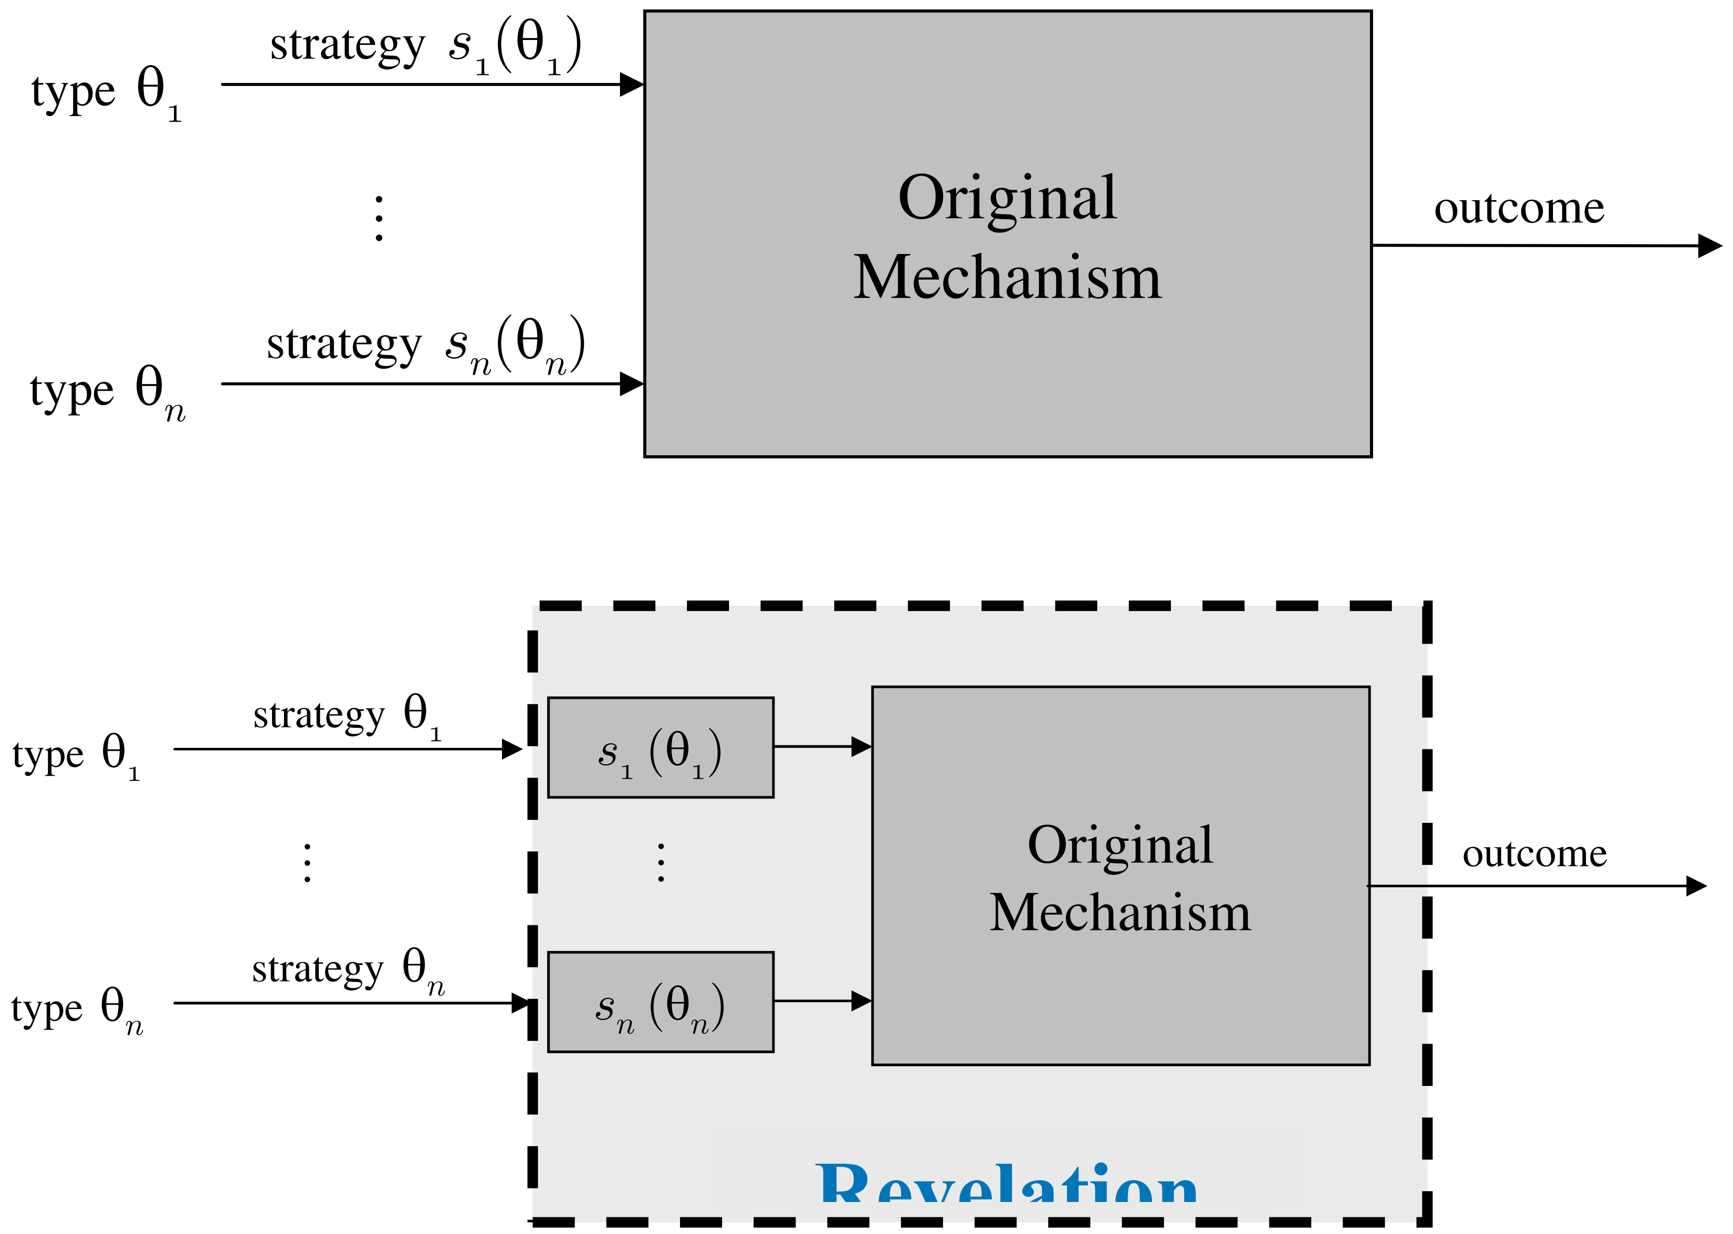

(Image: Shoham & Leyton-Brown 2008)

Recap: General Dominant-Strategy Implementation

Theorem: (Gibbard-Satterthwaite)

Consider any social choice function C over N and O. If $|O| > 2$ (there are at least **three** outcomes),

1. C is **onto**; that is, for every outcome $o \in O$ there is a preference profile $[\succ]$ such that $C([\succ]) = o$ (this is sometimes called **citizen sovereignty**), and
2. C is dominant-strategy **truthful**,

then C is **dictatorial**.

Recap: Quasilinear Preferences

Definition:

Agents have **quasilinear preferences** in an n-player Bayesian game setting when

1. the set of outcomes is $O = X \times \mathbb{R}^n$ for a finite set X,
2. the utility of agent i given type profile θ for an element $(x, p) \in O$ is $u_i\big((x, p), \theta\big) = v_i(x, \theta) - f_i(p_i)$, where
3. $v_i : X \times \Theta \to \mathbb{R}$ is an **arbitrary** function, and
4. $f_i : \mathbb{R} \to \mathbb{R}$ is a **monotonically increasing** function.

Recap: Direct Quasilinear Mechanism

Definition:

A **direct quasilinear mechanism** is a pair (χ, p), where

- $\chi : \Theta \to \Delta(X)$ is the **choice rule** (often called the **allocation rule**), which maps from a profile of reported types to a distribution over nonmonetary outcomes, and

- $p : \Theta \to \mathbb{R}^n$ is the **payment rule**, which maps from a profile of reported types to a payment for each agent.

Logistics

- Next week is **reading week**; no lectures
- **Assignment 2** will be released at the end of this week
  - Due **Thu March 7**
- On **Tue Feb 27** (first day after reading week) we will choose the schedule for **project presentations**
  - Using the random dictatorship mechanism as for paper assignments

Paper Presentations

Paper presentations start after reading week:

- There will be **2 or 3** presentations per class
  - (Rabin 2000 is rescheduled to Oct 31)
- Each paper is allocated **25 minutes** for talk + questions
  - Budget for about a 15-20 minute talk and 5-10 minutes for questions
  - I will be **ruthless** about the 25 minute time limit
- Summarize the **important parts** of the paper
- **Paper summaries** are due **before class** starts (200-500 words)
  - Submit via eClass
  - The eClass assignment description will tell you what these should include

Lecture Outline

1. Recap & Logistics
2. Efficient Quasilinear Mechanisms
3. Properties of Quasilinear Mechanisms

Groves Mechanisms

Definition:

Groves mechanisms are direct quasilinear mechanisms (χ, p) for which

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

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

- Where h_i is an **arbitrary** function of the reports of the other agents
- Groves mechanisms implement **any social welfare maximizing** choice function in **dominant strategies**

Proof Sketch: Dominant Strategies

1. Suppose that every other agent j declares arbitrary $\hat{v}_j$
2. Agent i wants to report $\hat{v}_i$ that solves $\max_{\hat{v}_i} \left(v_i\left(\chi(\hat{v}_i, \hat{v}_{-i})\right) - p_i(\hat{v}_i, \hat{v}_{-i}) \right)$.
3. Substitute p_i: $\max_{\hat{v}_i} \left(v_i\left(\chi(\hat{v}_i, \hat{v}_{-i})\right) \cancel{- h_i(\hat{v}_{-i})} + \sum_{j \neq i} \hat{v}_j\left(\chi(\hat{v}_i, \hat{v}_{-i})\right) \right)$
4. $h_i(\hat{v}_{-i})$ doesn't depend on $\hat{v}_i$

Proof Sketch #2

5. So i should report $\arg\max_{\hat{v}_i} \left(v_i\left(\chi(\hat{v}_i, \hat{v}_{-i})\right) + \sum_{j \neq i} \hat{v}_j\left(\chi(\hat{v}_i, \hat{v}_{-i})\right) \right)$

6. But Groves will choose $\arg\max_{\chi(\hat{v}_i, \hat{v}_{-i})} \left(\hat{v}_i\left(\chi(\hat{v}_i, \hat{v}_{-i})\right) + \sum_{j \neq i} \hat{v}_j\left(\chi(\hat{v}_i, \hat{v}_{-i})\right) \right)$

7. So i should report $\hat{v}_i = v_i$. ∎

Dominant strategies, because this argument is for **arbitrary** $\hat{v}_{-i}$.

Vickrey-Clarke-Groves Mechanism

Definition:
The **Vickery-Clarke-Groves mechanism** is a direct quasilinear mechanism (χ, p), where

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

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

- i.e., it's a Groves mechanism with $h_i(\hat{v}_{-i}) = \Sigma_{j \neq i} \hat{v}_j(\chi(\hat{v}_{-i}))$.
- Each agent pays their **externality**: difference between *other* agents' utility **if i weren't there** and the *other* agents' utility given that ***i* is there.**
- **Question:** Why don't we use this for **everything**?

Second Price Auctions *Are* VCG

The **second price auction** is **VCG** in the quasilinear **single-item auction setting**:

- Agents are **not permitted unrestricted** preferences over the outcome space of allocations and payments
- Object is awarded to agent with **highest valuation**; this maximizes the sum of (non-monetary) agent valuations for the outcome
- Externality of **winning agent** is the value that **next-highest-valuation agent** could have gotten by winning the auction
- Externality of **losing agent** is nothing; if they weren't there, the outcome would be no different

Externalities: Example

1. Who wins the second-price auction?
   i.e., $\chi(\hat{v})$
2. Who would win if Alice weren't in the auction?
   i.e., $\chi(\hat{v}_{-Alice})$
3. How much does Alice pay?
4. What is the VCG payment?

$$v_{Alice}(\text{Alice gets object}) = 10$$
$$v_{Bob}(\text{Bob gets object}) = 6$$
$$v_{Carol}(\text{Carol gets object}) = 3$$
$$v_{Dave}(\text{Dave gets object}) = 1$$

$$\sum_{j \neq Alice} \hat{v}_j(\text{Bob}) - \sum_{j \neq Alice} \hat{v}_j(\text{Alice}) = (6 + 0 + 0) - (0 + 0 + 0) = 6$$

Mechanism Properties

Definition:
A quasilinear mechanism is **truthful** if it is direct and $\forall i \in N, \forall v_i$, agent i's equilibrium strategy is to adopt the strategy $\hat{v}_i = v_i$.

Definition:
A quasilinear mechanism is **Pareto efficient**, or just **efficient**, if for all v in equilibrium it selects a choice x such that

$$\forall x' \sum_i v_i(x) \geq \sum_i v_i(x').$$

Budget Balance

Definition:
A quasilinear mechanism is **weakly budget balanced** when

$$\forall v, \sum_i p_i(s^*(v)) \geq 0,$$

where s^* is the equilibrium strategy profile.

Individual Rationality

Definition:
A quasilinear mechanism is ***ex-interim* individually rational** when

$$\forall i \, \forall v_i, \mathbb{E}_{v_{-i}|v_i}\left[v_i(\chi(s_i(v_i), s_{-i}(v_{-i}))) - p_i(s_i(v_i), s_{-i}(v_{-i}))\right] \geq 0.$$

All Efficient Dominant Strategy Mechanisms are Groves Mechanisms

Theorem: (Green-Laffont)

An **efficient** social choice function $C : \mathbb{R}^{X \times N} \to X \times \mathbb{R}^N$ can be implemented in dominant strategies for agents with **unrestricted quasilinear utilities only if**

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

One Last Impossibility Result

Theorem: (Myerson-Satterthwaite)
No Bayes-Nash incentive-compatible mechanism is always simultaneously efficient, weakly budget-balanced, and ex-interim individually rational, even if agents are restricted to quasilinear utility functions.

- It does turn out to be possible to get any **two of the three**
- **Question:** Wait a minute, doesn't the second-price auction satisfy all three conditions?

Summary

- When agents are restricted to **quasilinear preferences**, social choice functions *can* be implemented in **dominant strategies**
- **Groves mechanisms** are the unique class of mechanisms that implement efficient social choice functions in dominant strategies
  - **VCG** is the pre-eminent Groves mechanism
  - Second-price auctions turn out to be VCG in the single-item auction setting
- You can only have **two** of **efficiency**, **weak budget balance**, and ***ex-interim* individual rationality**, even in the quasilinear setting