Quasilinear Mechanisms; Groves Mechanism

Lecture 15

Lecture Overview

Revelation Principle

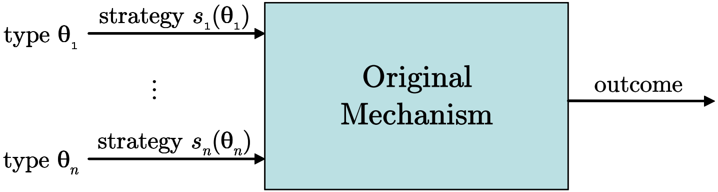

- It turns out that truthfulness can always be achieved!
- Consider an arbitrary, non-truthful mechanism (e.g., may be indirect)
- Recall that a mechanism defines a game, and consider an equilibrium $s = (s_1, \ldots, s_n)$

Revelation Principle

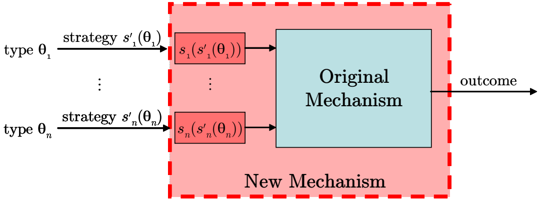

- We can construct a new direct mechanism, as shown above
- This mechanism is truthful by exactly the same argument that s was an equilibrium in the original mechanism
- "The agents don't have to lie, because the mechanism already lies for them."

Impossibility Result

Theorem (Gibbard-Satterthwaite)

Consider any social choice function C of N and O. If:

1. $|O| \geq 3$ *(there are at least three outcomes);*
2. C *is onto; that is, for every* $o \in O$ *there is a preference profile* $[\succ]$ *such that* $C([\succ]) = o$ *(this property is sometimes also called citizen sovereignty); and*
3. C *is dominant-strategy truthful,*

then C is dictatorial.

Quasilinear Utility

Definition (Quasilinear preferences)

Agents have quasilinear preferences in an n-player Bayesian game when the set of outcomes is

$$O = X \times \mathbb{R}^n$$

for a finite set X, and the utility of an agent i given joint type θ is given by

$$u_i(o, \theta) = u_i(x, \theta) - f_i(p_i),$$

where $o = (x, p)$ is an element of O, $u_i : X \times \Theta \mapsto \mathbb{R}$ is an arbitrary function and $f_i : \mathbb{R} \mapsto \mathbb{R}$ is a strictly monotonically increasing function.

Risk Neutrality

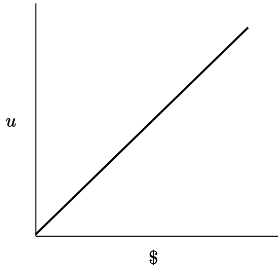

(a) Risk neutrality

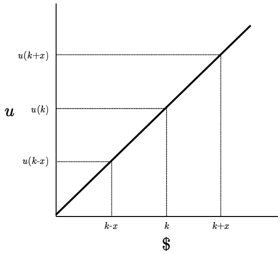

(b) Risk neutrality: fair lottery

Risk Aversion

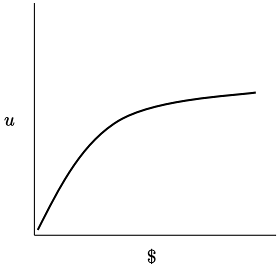

(c) Risk aversion

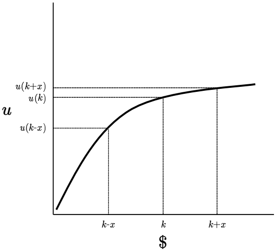

(d) Risk aversion: fair lottery

Risk Seeking

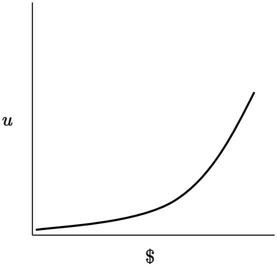

(e) Risk seeking

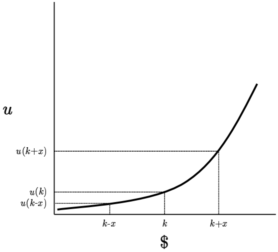

(f) Risk seeking: fair lottery

Lecture Overview

Quasilinear Mechanism

Definition (Quasilinear mechanism)

A mechanism in the quasilinear setting (for a Bayesian game setting $(N, O = X \times \mathbb{R}^n, \Theta, p, u)$) is a triple (A, χ, p), where

- $A = A_1 \times \cdots \times A_n$, where A_i is the set of actions available to agent $i \in N$,
- $\chi : A \mapsto \Pi(X)$ maps each action profile to a distribution over choices, and
- $p : A \mapsto \mathbb{R}^n$ maps each action profile to a payment for each agent.

Direct Quasilinear Mechanism

Definition (Direct quasilinear mechanism)

A direct quasilinear mechanism (for a Bayesian game setting $(N, O = X \times \mathbb{R}^n, \Theta, p, u)$) is a pair (χ, p). It defines a standard mechanism in the quasilinear setting, where for each i, $A_i = \Theta_i$.

Definition (Conditional utility independence)

A Bayesian game exhibits conditional utility independence if for all agents $i \in N$, for all outcomes $o \in O$ and for all pairs of joint types θ and $\theta' \in \Theta$ for which $\theta_i = \theta'_i$, it holds that $u_i(o, \theta) = u_i(o, \theta')$.

Quasilinear Mechanisms with Conditional Utility Independence

- Given conditional utility independence, we can write i's utility function as $u_i(o, \theta_i)$
  - it does not depend on the other agents' types
- An agent's valuation for choice $x \in X$: $v_i(x) = u_i(x, \theta_i)$
  - the maximum amount i would be willing to pay to get x
  - in fact, i would be indifferent between keeping the money and getting x
- Alternate definition of direct mechanism:
  - ask agents i to declare $v_i(x)$ for each $x \in X$
- Define $\hat{v}_i$ as the valuation that agent i declares to such a direct mechanism
  - may be different from his true valuation v_i
- Also define the tuples $\hat{v}$, $\hat{v}_{-i}$

Lecture Overview

1. Recap
2. Quasilinear Mechanisms
3. Properties
4. The Groves Mechanism

Truthfulness

Definition (Truthfulness)

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

- Our definition before, adapted for the quasilinear setting

Efficiency

Definition (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 efficient mechanism selects the choice that maximizes the sum of agents' utilities, disregarding monetary payments.
- How is this related to Pareto efficiency from GT?

Efficiency

Definition (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 efficient mechanism selects the choice that maximizes the sum of agents' utilities, disregarding monetary payments.
- How is this related to Pareto efficiency from GT?
  - if we include the mechanism as an agent, all Pareto-efficient outcomes involve the same choice (and different payments)
  - any outcome involving another choice is Pareto-dominated: some agents could make a side-payment to others such that all would prefer the swap

Efficiency

Definition (Efficiency)

A quasilinear mechanism is strictly Pareto efficient, or just efficient, if in equilibrium it selects a choice x such that

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

- Called economic efficiency to distinguish from other (e.g., computational) notions
- Also called social-welfare maximization
- Note: defined in terms of true (not declared) valuations.

Budget Balance

Definition (Budget balance)

A quasilinear mechanism is budget balanced when

$$\forall v, \sum_i p_i(s(v)) = 0,$$

where s is the equilibrium strategy profile.

- regardless of the agents' types, the mechanism collects and disburses the same amount of money from and to the agents

Budget Balance

Definition (Budget balance)

A quasilinear mechanism is budget balanced when

$$\forall v, \sum_i p_i(s(v)) = 0,$$

where s is the equilibrium strategy profile.

- regardless of the agents' types, the mechanism collects and disburses the same amount of money from and to the agents
- relaxed version: weak budget balance:

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

  - the mechanism never takes a loss, but it may make a profit

Budget Balance

Definition (Budget balance)

A quasilinear mechanism is budget balanced when

$$\forall v, \sum_i p_i(s(v)) = 0,$$

where s is the equilibrium strategy profile.

- regardless of the agents' types, the mechanism collects and disburses the same amount of money from and to the agents
- Budget balance can be required to hold *ex ante*:

$$\mathbb{E}_v \sum_i p_i(s(v)) = 0$$

  - the mechanism must break even or make a profit only on expectation

Individual-Rationality

Definition (*Ex interim* individual rationality)

A mechanism is ex interim individual rational when
$\forall i \forall v_i,\ \mathbb{E}_{v_{-i}|v_i} v_i(\chi(s_i(v_i), s_{-i}(v_{-i}))) - p_i(s_i(v_i), s_{-i}(v_{-i})) \geq 0$,
where s is the equilibrium strategy profile.

- no agent loses by participating in the mechanism.
- *ex interim* because it holds for *every* possible valuation for agent i, but averages over the possible valuations of the other agents.

Individual-Rationality

Definition (*Ex interim* individual rationality)

A mechanism is ex interim individual rational when
$\forall i \forall v_i,\ \mathbb{E}_{v_{-i}|v_i} v_i(\chi(s_i(v_i), s_{-i}(v_{-i}))) - p_i(s_i(v_i), s_{-i}(v_{-i})) \geq 0$,
where s is the equilibrium strategy profile.

- no agent loses by participating in the mechanism.
- *ex interim* because it holds for *every* possible valuation for agent i, but averages over the possible valuations of the other agents.

Definition (*Ex post* individual rationality)

A mechanism is ex post individual rational when
$\forall i \forall v,\ v_i(\chi(s(v))) - p_i(s(v)) \geq 0$, where s is the equilibrium strategy profile.

Tractability

Definition (Tractability)

A mechanism is tractable when $\forall \hat{v}$, $\chi(\hat{v})$ and $p(\hat{v})$ can be computed in polynomial time.

- The mechanism is computationally feasible.

Revenue Maximization

We can also add an objective function to our mechanism. One example: revenue maximization.

Definition (Revenue maximization)

A mechanism is revenue maximizing when, among the set of functions χ and p that satisfy the other constraints, the mechanism selects the χ and p that maximize $\mathbb{E}_\theta \sum_i p_i(s(\theta))$, where $s(\theta)$ denotes the agents' equilibrium strategy profile.

- The mechanism designer can choose among mechanisms that satisfy the desired constraints by adding an objective function such as revenue maximization.

Revenue Minimization

- The mechanism may not be intended to make money.
- Budget balance may be impossible to satisfy.
- Set weak budget balance as a constraint and add the following objective.

Definition (Revenue minimization)

A quasilinear mechanism is revenue minimizing when, among the set of functions χ and p that satisfy the other constraints, the mechanism selects the χ and p that minimize $\max_v \sum_i p_i(s(v))$ in equilibrium, where $s(v)$ denotes the agents' equilibrium strategy profile.

- Note: this considers the worst case over valuations; we could consider average case instead.

Fairness

- Fairness is hard to define. What is fairer:
  - an outcome that fines all agents \$100 and makes a choice that all agents hate equally?
  - an outcome that charges all agents \$0 and makes a choice that some agents hate and some agents like?

Fairness

- Fairness is hard to define. What is fairer:
  - an outcome that fines all agents \$100 and makes a choice that all agents hate equally?
  - an outcome that charges all agents \$0 and makes a choice that some agents hate and some agents like?
- Maxmin fairness: make the least-happy agent the happiest.

Definition (Maxmin fairness)

A quasilinear mechanism is maxmin fair when, among the set of functions χ and p that satisfy the other constraints, the mechanism selects the χ and p that maximize

$$\mathbb{E}_v \left[\min_{i \in N} v_i(\chi(s(v))) - p_i(s(v)) \right],$$

where $s(v)$ denotes the agents' equilibrium strategy profile.

Price of Anarchy Minimization

- When an efficient mechanism is impossible, we may want to get as close as possible
- Minimize the worst-case ratio between optimal social welfare and the social welfare achieved by the given mechanism.

Definition (Price-of-anarchy minimization)

A quasilinear mechanism minimizes the price of anarchy when, among the set of functions χ and p that satisfy the other constraints, the mechanism selects the χ and p that minimize

$$\max_{v \in V} \frac{\max_{x \in X} \sum_{i \in N} v_i(x)}{\sum_{i \in N} v_i\left(\chi(s(v))\right)},$$

where $s(v)$ denotes the agents' equilibrium strategy profile in the *worst* equilibrium of the mechanism—i.e., the one in which $\sum_{i \in N} v_i(\chi(s(v)))$ is the smallest.

Lecture Overview

A positive result

- Recall that in the quasilinear utility setting, a mechanism can be defined as a choice rule and a payment rule.
- The Groves mechanism is a mechanism that satisfies:
  - dominant strategy (truthfulness)
  - efficiency
- In general it's not:
  - budget balanced
  - individual-rational

...though we'll see later that there's some hope for recovering these properties.

The Groves Mechanism

Definition (Groves mechanism)

The Groves mechanism is a direct quasilinear mechanism (χ, p), where

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

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

The Groves Mechanism

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

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

- The choice rule should not come as a surprise (why not?)

The Groves Mechanism

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

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

- The choice rule should not come as a surprise (why not?) because the mechanism is both truthful and efficient: these properties entail the given choice rule.

The Groves Mechanism

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

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

- The choice rule should not come as a surprise (why not?) because the mechanism is both truthful and efficient: these properties entail the given choice rule.
- So what's going on with the payment rule?
  - the agent i must pay some amount $h_i(\hat{v}_{-i})$ that doesn't depend on his own declared valuation
  - the agent i is paid $\sum_{j \neq i} \hat{v}_j(\chi(\hat{v}))$, the sum of the others' valuations for the chosen outcome

Groves Truthfulness

Theorem

Truth telling is a dominant strategy under the Groves mechanism.

Consider a situation where every agent j other than i follows some arbitrary strategy $\hat{v}_j$. Consider agent i's problem of choosing the best strategy $\hat{v}_i$. As a shorthand, we will write $\hat{v} = (\hat{v}_{-i}, \hat{v}_i)$. The best strategy for i is one that solves

$$\max_{\hat{v}_i} \left(v_i(\chi(\hat{v})) - p(\hat{v}) \right)$$

Substituting in the payment function from the Groves mechanism, we have

$$\max_{\hat{v}_i} \left(v_i(\chi(\hat{v})) - h_i(\hat{v}_{-i}) + \sum_{j \neq i} \hat{v}_j(\chi(\hat{v})) \right)$$

Since $h_i(\hat{v}_{-i})$ does not depend on $\hat{v}_i$, it is sufficient to solve

$$\max_{\hat{v}_i} \left(v_i(\chi(\hat{v})) + \sum_{j \neq i} \hat{v}_j(\chi(\hat{v})) \right).$$

Groves Truthfulness

$$\max_{\hat{v}_i} \left(v_i(\chi(\hat{v})) + \sum_{j \neq i} \hat{v}_j(\chi(\hat{v})) \right).$$

The only way the declaration $\hat{v}_i$ influences this maximization is through the choice of x. If possible, i would like to pick a declaration $\hat{v}_i$ that will lead the mechanism to pick an $x \in X$ which solves

$$\max_x \left(v_i(x) + \sum_{j \neq i} \hat{v}_j(x) \right). \quad (1)$$

Under the Groves mechanism,

$$\chi(\hat{v}) = \arg\max_x \left(\sum_i \hat{v}_i(x) \right) = \arg\max_x \left(\hat{v}_i(x) + \sum_{j \neq i} \hat{v}_j(x) \right).$$

The Groves mechanism will choose x in a way that solves the maximization problem in Equation (1) when i declares $\hat{v}_i = v_i$. Because this argument does not depend in any way on the declarations of the other agents, truth-telling is a dominant strategy for agent i.

Proof intuition

- externalities are internalized
  - agents may be able to change the outcome to another one that they prefer, by changing their declaration
  - however, their utility doesn't just depend on the outcome—it also depends on their payment
  - since they get paid the (reported) utility of all the other agents under the chosen allocation, they now have an interest in maximizing everyone's utility rather than just their own
- in general, DS truthful mechanisms have the property that an agent's payment doesn't depend on the amount of his declaration, but only on the other agents' declarations
  - the agent's declaration is used only to choose the outcome, and to set other agents' payments

Groves Uniqueness

Theorem (Green–Laffont)

An efficient social choice function $C : \mathbb{R}^{Xn} \rightarrow X \times \mathbb{R}^n$ can be implemented in dominant strategies for agents with unrestricted quasilinear utilities only if $p_i(v) = h(v_{-i}) - \sum_{j \neq i} v_j(\chi(v))$.

- it turns out that the same result also holds for the broader class of Bayes–Nash incentive-compatible efficient mechanisms.