VCG

Lecture 16

Lecture Overview

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.
- 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
- we can also define weak or ex ante variants

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 quasilinear mechanism is tractable when $\forall a \in A$, $\chi(a)$ and $p(a)$ 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

- 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.

Lecture Overview

VCG

Definition (Clarke tax)

The Clarke tax sets the h_i term in a Groves mechanism as

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

Definition (Vickrey-Clarke-Groves (VCG) mechanism)

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

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

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

VCG discussion

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

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

- You get paid everyone's utility under the allocation that is actually chosen
  - except your own, but you get that directly as utility
- Then you get charged everyone's utility in the world where you don't participate
- Thus you pay your social cost

VCG discussion

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

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

Questions:

- who pays 0?

VCG discussion

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

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

Questions:

- who pays 0?
  - agents who don't affect the outcome

VCG discussion

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

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

Questions:

- who pays 0?
  - agents who don't affect the outcome
- who pays more than 0?

VCG discussion

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

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

Questions:

- who pays 0?
  - agents who don't affect the outcome
- who pays more than 0?
  - (pivotal) agents who make things worse for others by existing

VCG discussion

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

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

Questions:

- who pays 0?
    - agents who don't affect the outcome
- who pays more than 0?
    - (pivotal) agents who make things worse for others by existing
- who gets paid?

VCG discussion

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

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

Questions:

- who pays 0?
  - agents who don't affect the outcome
- who pays more than 0?
  - (pivotal) agents who make things worse for others by existing
- who gets paid?
  - (pivotal) agents who make things better for others by existing

VCG properties

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

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

- Because only pivotal agents have to pay, VCG is also called the pivot mechanism
- It's dominant-strategy truthful, because it's a Groves mechanism

Lecture Overview

Selfish routing example

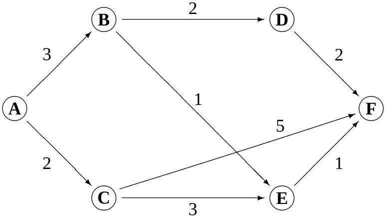

- What outcome will be selected by χ?

Selfish routing example

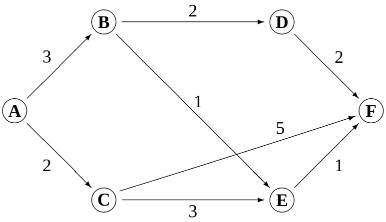

- What outcome will be selected by χ? path $ABEF$.

Selfish routing example

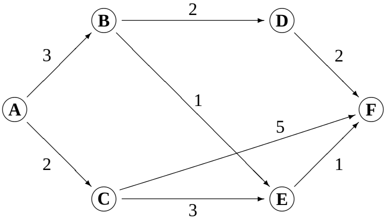

- What outcome will be selected by χ? path $ABEF$.
- How much will AC have to pay?

Selfish routing example

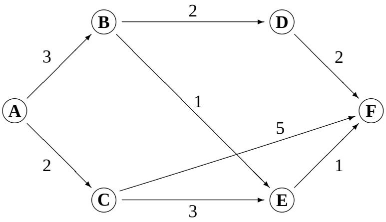

- What outcome will be selected by χ? path $ABEF$.
- How much will AC have to pay?
  - The shortest path taking his declaration into account has a length of 5, and imposes a cost of -5 on agents other than him (because it does not involve him). Likewise, the shortest path without AC's declaration also has a length of 5. Thus, his payment $p_{AC} = (-5) - (-5) = 0$.
  - This is what we expect, since AC is not pivotal.
  - Likewise, BD, CE, CF and DF will all pay zero.

Selfish routing example

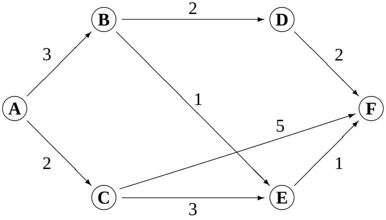

- How much will AB pay?

Selfish routing example

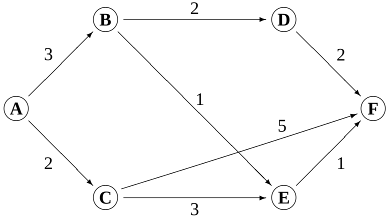

- How much will AB pay?
  - The shortest path taking AB's declaration into account has a length of 5, and imposes a cost of 2 on other agents.
  - The shortest path without AB is $ACEF$, which has a cost of 6.
  - Thus $p_{AB} = (-6) - (-2) = -4$.

Selfish routing example

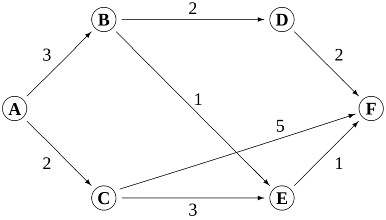

- How much will BE pay?

Selfish routing example

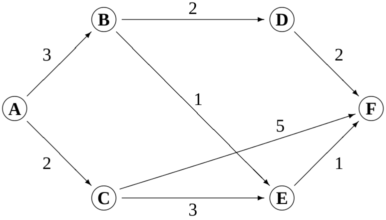

- How much will BE pay? $p_{BE} = (-6) - (-4) = -2.$

Selfish routing example

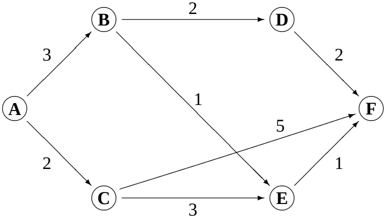

- How much will BE pay? $p_{BE} = (-6) - (-4) = -2.$
- How much will EF pay?

Selfish routing example

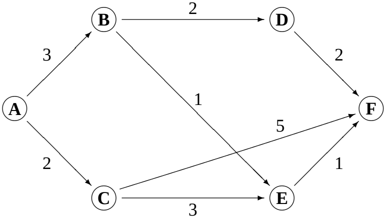

- How much will BE pay? $p_{BE} = (-6) - (-4) = -2.$
- How much will EF pay? $p_{EF} = (-7) - (-4) = -3.$

Selfish routing example

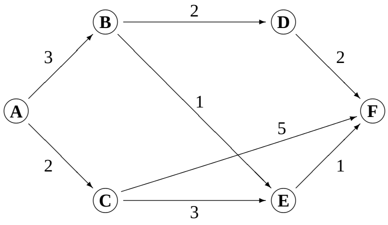

- How much will BE pay? $p_{BE} = (-6) - (-4) = -2$.
- How much will EF pay? $p_{EF} = (-7) - (-4) = -3$.
  - EF and BE have the same costs but are paid different amounts. Why?

Selfish routing example

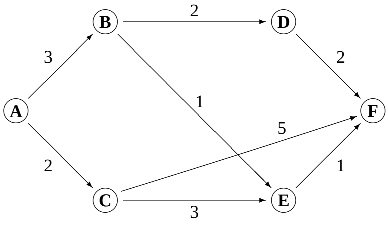

- How much will BE pay? $p_{BE} = (-6) - (-4) = -2$.
- How much will EF pay? $p_{EF} = (-7) - (-4) = -3$.
  - EF and BE have the same costs but are paid different amounts. Why?
  - EF has more *market power*: for the other agents, the situation without EF is worse than the situation without BE.

Lecture Overview

Two definitions

Definition (Choice-set monotonicity)

An environment exhibits choice-set monotonicity if $\forall i,\ X_{-i} \subseteq X$.

- removing any agent weakly decreases—that is, never increases—the mechanism's set of possible choices X

Definition (No negative externalities)

An environment exhibits no negative externalities if $\forall i \forall x \in X_{-i},\ v_i(x) \geq 0$.

- every agent has zero or positive utility for any choice that can be made without his participation

Example: road referendum

Example

Consider the problem of holding a referendum to decide whether or not to build a road.

- The set of choices is independent of the number of agents, satisfying choice-set monotonicity.
- No agent negatively values the project, though some might value the situation in which the project is not undertaken more highly than the situation in which it is.

Example: simple exchange

Example

Consider a market setting consisting of agents interested in buying a single unit of a good such as a share of stock, and another set of agents interested in selling a single unit of this good. The choices in this environment are sets of buyer-seller pairings (prices are imposed through the payment function).

- If a new agent is introduced into the market, no previously-existing pairings become infeasible, but new ones become possible; thus choice-set monotonicity is satisfied.
- Because agents have zero utility both for choices that involve trades between other agents and no trades at all, there are no negative externalities.

VCG Individual Rationality

Theorem

The VCG mechanism is ex-post individual rational when the choice set monotonicity and no negative externalities properties hold.

Proof.

All agents truthfully declare their valuations in equilibrium. Then

$$\begin{aligned} u_i &= v_i(\chi(v)) - \left(\sum_{j \neq i} v_j(\chi(v_{-i})) - \sum_{j \neq i} v_j(\chi(v)) \right) \\ &= \sum_i v_i(\chi(v)) - \sum_{j \neq i} v_j(\chi(v_{-i})) \end{aligned} \qquad (2)$$

$\chi(v)$ is the outcome that maximizes social welfare, and that this optimization could have picked $\chi(v_{-i})$ instead (by choice set monotonicity). Thus,

$$\sum_j v_j(\chi(v)) \geq \sum_j v_j(\chi(v_{-i})).$$

VCG Individual Rationality

Theorem

The VCG mechanism is ex-post individual rational when the choice set monotonicity and no negative externalities properties hold.

Proof.

$$\sum_j v_j(\chi(v)) \geq \sum_j v_j(\chi(v_{-i})).$$

Furthermore, from no negative externalities,

$$v_i(\chi(v_{-i})) \geq 0.$$

Therefore,

$$\sum_i v_i(\chi(v)) \geq \sum_{j \neq i} v_j(\chi(v_{-i})),$$

and thus Equation (2) is non-negative.

Lecture Overview

Another property

Definition (No single-agent effect)

An environment exhibits no single-agent effect if $\forall i$, $\forall v_{-i}$, $\forall x \in \arg\max_y \sum_j v_j(y)$ there exists a choice x' that is feasible without i and that has $\sum_{j \neq i} v_j(x') \geq \sum_{j \neq i} v_j(x)$.

Example

Consider a single-sided auction. Dropping an agent just reduces the amount of competition, making the others better off.

Good news

Theorem

The VCG mechanism is weakly budget-balanced when the no single-agent effect property holds.

Proof.

Assume truth-telling in equilibrium. We must show that the sum of transfers from agents to the center is greater than or equal to zero.

$$\sum_i p_i(v) = \sum_i \left(\sum_{j \neq i} v_j(\chi(v_{-i})) - \sum_{j \neq i} v_j(\chi(v)) \right)$$

From the no single-agent effect condition we have that

$$\forall i \sum_{j \neq i} v_j(\chi(v_{-i})) \geq \sum_{j \neq i} v_j(\chi(v)).$$

Thus the result follows directly.

More good news

Theorem (Krishna & Perry, 1998)

In any Bayesian game setting in which VCG is ex post individually rational, VCG collects at least as much revenue as any other efficient and ex interim individually-rational mechanism.

- This is somewhat surprising: does not require dominant strategies, and hence compares VCG to all Bayes–Nash mechanisms.
- A useful corollary: VCG is as budget balanced as any efficient mechanism can be
  - it satisfies weak budget balance in every case where *any* dominant strategy, efficient and *ex interim* IR mechanism would be able to do so.

Bad news

Theorem (Green–Laffont; Hurwicz)

No dominant-strategy incentive-compatible mechanism is always both efficient and weakly budget balanced, even if agents are restricted to the simple exchange setting.

Theorem (Myerson–Satterthwaite)

No Bayes-Nash incentive-compatible mechanism is always simultaneously efficient, weakly budget balanced and ex-interim individual rational, even if agents are restricted to quasilinear utility functions.