CS 598: AI Methods for Market Design

Lecture 5: Mechanism Design

Xintong Wang

Spring 2024

Announcements

- Two paper presentations today! One for next week
- HW1 will be out next week
  - You can work in pairs or individually
- Please NO ChatGPT (Gen AI) for CQs and homework
- You can drop two pre-class CQs, no more exceptions unless late enrollment
- Office hour today: after class to 2pm

Outline

- Recap: revenue equivalence
- Multi-round auctions
- Direct-revelation mechanism
- Revelation principle
- Design desiderata
- The VCG mechanism

Outline

- Recap: revenue equivalence
- Multi-round auctions
- Direct-revelation mechanism
- Revelation principle
- Design desiderata
- The VCG mechanism

Characterizing BNE in Auctions

Given an auction and an equilibrium strategy s^*, at an intermediate state where bidder i knows v_i

- Interim allocation for bidder i with v_i

The probability the item is allocate to bidder i in eq.

- Interim payment for bidder i with v_i

The expected payment made by bidder i in eq.

- Interim utility

Characterizing BNE in Auctions

Given an auction and an equilibrium strategy s^*, at an intermediate state where bidder i knows v_i

- Interim allocation for bidder i with v_i

$$x_i^*(v_i) = \mathbf{E}_{v_{-i}}\left[x_i(s_i^*(v_i), s_{-i}^*(v_{-i}))\right]$$

- Interim payment for bidder i with v_i

$$t_i^*(v_i) = \mathbf{E}_{v_{-i}}\left[t_i(s_i^*(v_i), s_{-i}^*(v_{-i}))\right]$$

- Interim utility then is $u_i^*(v_i) = v_i x_i^*(v_i) - t_i^*(v_i)$

Characterizing BNE in Auctions

Theorem. In any BNE of any sealed-bid auction, for bidder i with v_i, we have

- Interim monotonicity: The interim allocation $x_i^*(v_i)$ is monotone weakly increasing in value v_i
- Interim payment identity: For value v_i and interim allocation $x_i^*(v_i)$, the interim payment is

$$t_i^*(v_i) = v_i \times x_i^*(v_i) - \int_{z=v_{\min}}^{v_i} x_i^*(z)\mathrm{d}z - C_i$$

where C_i is a constant.

Rate of increase in interim payment = Rate of increase in interim value

Revenue Equivalence

A *normalized auction* is one where a bidder with value 0 (or v_{min}) has zero interim utility

Theorem. Any two normalized, sealed-bid auctions that each have a BNE with an identical interim allocation have the same expected revenue in these two BNE.

Proof:

same interim allocation → same interim payment

→ same revenue

$$\mathrm{Rev} = \sum_{i=1}^{n} \mathbf{E}_{v_i}[t_i^*(v_i)]$$

Finding a BNE in an Auction: Guess & Verify

- Guess that some auction design, A, with values i.i.d. sampled from distribution G has an efficient BNE and is normalized (i.e., $s_i(v_i) = 0$ for $v_i = 0$)

- Construct a strategy profile s, such that the interim payment in auction A at s is equal to the interim payment at truthful DSE of a SPSB auction

- Verify that s is a BNE in auction A. Confirm that auction A is efficient with s and normalized

Finding a BNE in an Auction: Guess & Verify

Question:

Follow guess & verify to derive the BNE for (1) the all-pay auction and (2) the FPSB auction

Outline

- Recap: revenue equivalence
- Multi-round auctions
- Direct-revelation mechanism
- Revelation principle
- Design desiderata
- The VCG mechanism

Multi-Round Auctions

- May allow bidders to respond to others' bids, esp. in interdependent or common value scenarios
- May have more flexible strategies
- Can be helpful in transparency and credibility

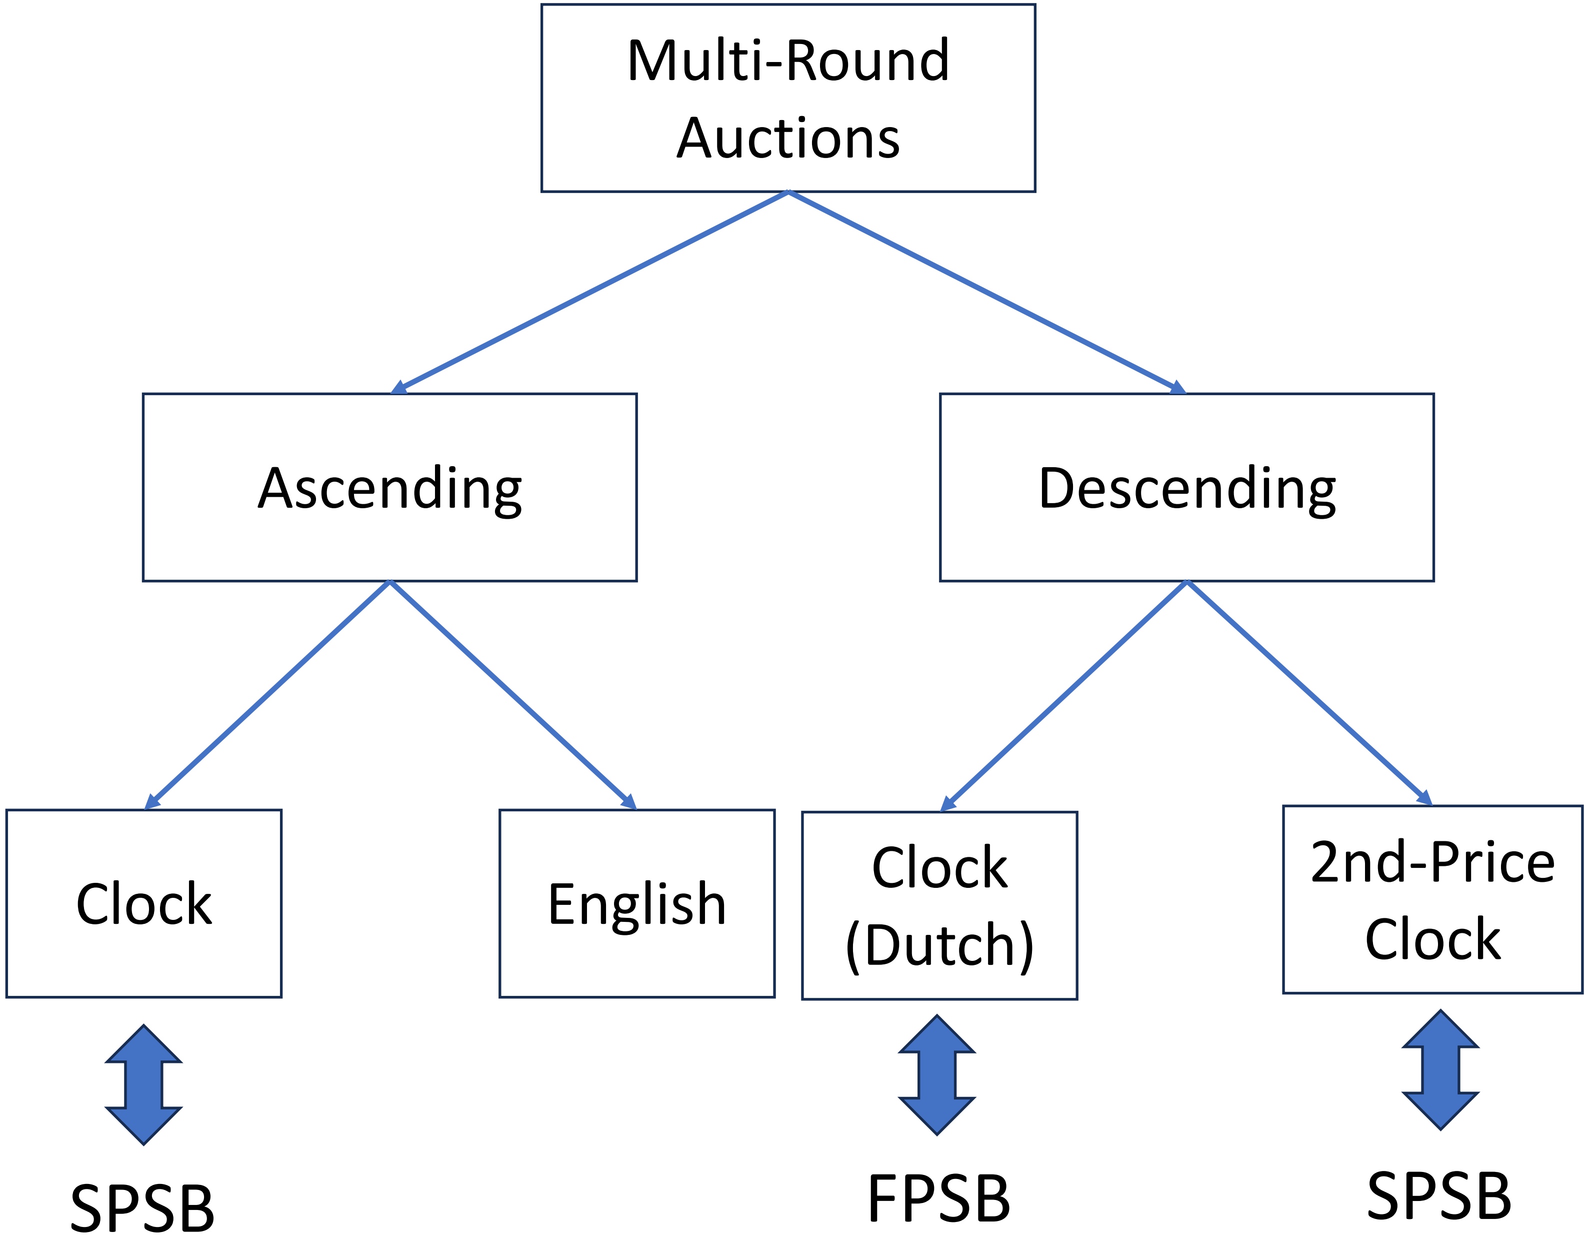
Multi-Round Auctions
Ascending
Descending
Clock
English
Clock (Dutch)
2nd-Price Clock
SPSB
FPSB
SPSB

Outline

- Recap: revenue equivalence
- Multi-round auctions
- Direct-revelation mechanism
- Revelation principle
- Design desiderata
- The VCG mechanism

Economic Environment

- A set of $N = \{1, \ldots, n\}$ agents
- A set of A alternatives
  - The time of a meeting, the assignment of ads to slots

Economic Environment

- A set of $N = \{1, \ldots, n\}$ agents
- A set of A alternatives
  - The time of a meeting, the assignment of ads to slots
- Settings without money
  - Each agent has a preference ordering: $a \succeq_i b$ for $a, b \in A$
  - Preference profile: $\succeq = (\succeq_1, \cdots, \succeq_n)$

Economic Environment

- A set of $N = \{1, \dots, n\}$ agents
- A set of A alternatives
  - The time of a meeting, the assignment of ads to slots
- Settings without money
  - Each agent has a preference ordering: $a \succeq_i b$ for $a, b \in A$
  - Preference profile: $\succeq = (\succeq_1, \cdots, \succeq_n)$
- Settings with money / payment
  - Each agent has a valuation function: $v_i: A \to \mathbb{R}$
  - Valuation profile: $v = (v_1, \dots, v_n)$
  - Each agent has a quasi-linear utility function:
    $u_i(a, p_i) = v_i(a) - p_i$, for alternative a and payment $p_i \in \mathbb{R}$

Direct-Revelation Mechanisms

- A direct-revelation mechanism (DRM) involves a single round of communication where each agent makes a simultaneous report of their preferences / valuation functions

Direct-Revelation Mechanisms

- The mechanism outcome (without money)

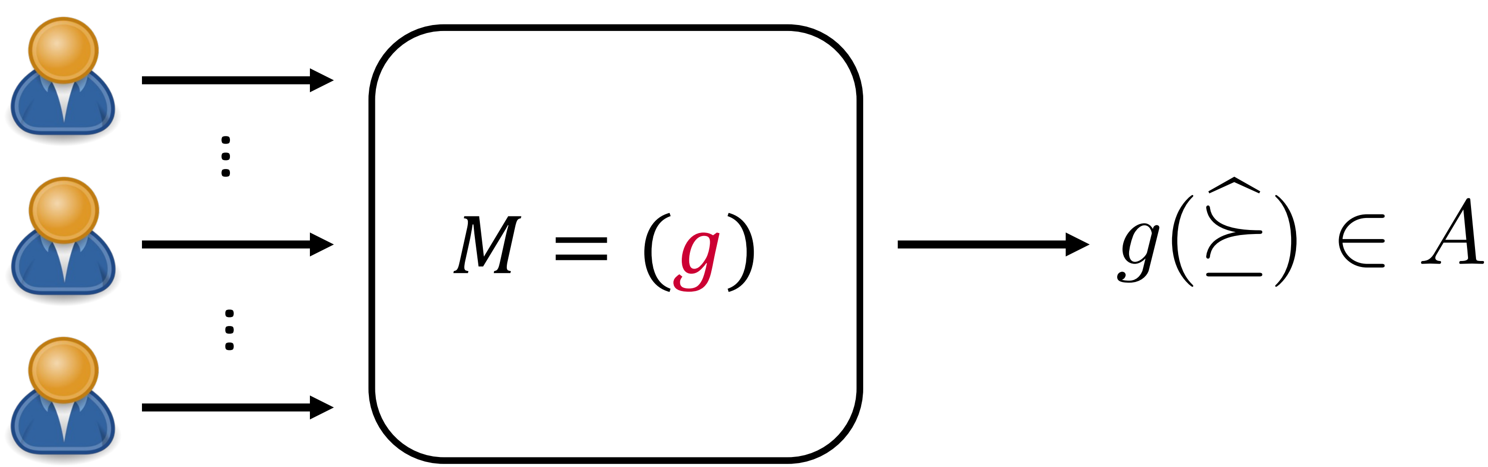

$\widehat{\succeq} = (\widehat{\succeq_1}, \cdots, \widehat{\succeq_n})$

Reported preference profile

An outcome rule $g: P^n \to A$

Direct-Revelation Mechanisms

- The mechanism outcome (without money)
- Example: meeting scheduling
  - Alternatives: 9am, 10am, 11am
  - Three agents with their preference orderings
    - 11am $\succ_1$ 10am $\succ_1$ 9am
    - 9am $\succ_2$ 11am $\succ_2$ 10am
    - 10am $\succ_3$ 9am $\succ_3$ 11am
  - Plurality rule, tie breaking in favor of earlier time
  - Under truthful report, $\mathrm{g}(\succ) =$ 9am
  - What would be a beneficial deviation for agent 1?

Direct-Revelation Mechanisms

- The mechanism outcome (with money)

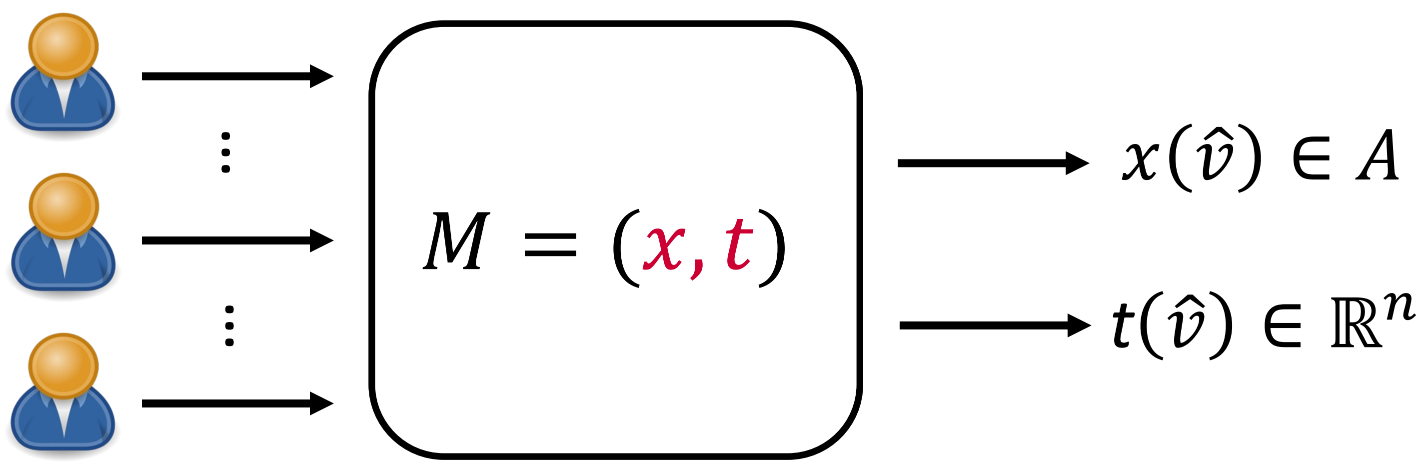

$\hat{v} = (\widehat{v_1}, \ldots, \widehat{v_n})$

Reported valuation profile

A choice rule $x: V \to A$

A payment rule $t: V \to \mathbb{R}^n$

Direct-Revelation Mechanisms

- The mechanism outcome (with money)
- Example: single-item auction
  - Alternatives: "do not allocation" & "allocate to each agent"
  - Three agents with their valuation functions
    - Agent 1, 2, 3 have values \$10, \$8, \$4 for the item

$$v_1(a) = \begin{cases} 10 & \text{if } a \in A \text{ assigns the item to agent 1} \\ 0 & \text{otherwise} \end{cases}$$

  - Under SPSB auction
    - The choice rule is $x(\hat{v}) = \text{argmax}_{a \in A}\left(\widehat{v_1}(a) + \widehat{v_2}(a) + \widehat{v_3}(a)\right)$
    - The payment rule is, for $x(\hat{v}) = i$,
      - $t_i(\hat{v}) = \max_{a \in A} \sum_{j \neq i} \widehat{v_j}(a)$ and $t_j(\hat{v}) = 0$ for $j \neq i$

Algorithm Design vs. Mechanism Design

Algorithm
- Fixed input
- Design won't change input
- E.g., route planning

Mechanism
- Strategic input
- Design may affect input and then outcome
- E.g., FPSB, SPSB

Mechanism Design as a Game

- A game of incomplete information
- Dominant-strategy equilibrium: a robust prediction of agent behavior and mechanism outcome

Mechanism Design as a Game

- DSE (no money)

A strategy profile $s^* = (s_1^*, \dots, s_n^*)$ is a dominant-strategy equilibrium in a DRM $M = (g)$ if and only if, for every agent i

$$g(s^*(\succeq_i), s_{-i}(\succeq_{-i})) \succeq_i g(\widehat{\succeq}_i, s_{-i}(\succeq_{-i})),$$

$$for\ all \succeq_i,\ all\ \widehat{\succeq}_i,\ all \succeq_{-i},\ all\ s_{-i}$$

No misreport will strictly improve agent utility.

Mechanism Design as a Game

- DSE (with money)

A strategy profile $s^* = (s_1^*, \ldots, s_n^*)$ is a dominant-strategy equilibrium in a DRM $M = (x, t)$ if and only if, for every agent i

$$v_i(x(s_i^*(v_i), s_{-i}(v_{-i}))) - t_i(s_i^*(v_i), s_{-i}(v_{-i}))$$
$$\geq v_i(x(\hat{v}_i, s_{-i}(v_{-i}))) - t_i(\hat{v}_i, s_{-i}(v_{-i}))$$
$$\textit{for all } v_i, \textit{ all } \hat{v}_i, \textit{ all } v_{-i}, \textit{ all } s_{-i}$$

No misreport will strictly improve agent utility.

Mechanism Design as a Game

- Strategy-proof mechanism

A DRM is strategy-proof when truthful reporting is a DSE

- Terminology

Strategy-proof

Dominant-strategy incentive compatible (DSIC)

Truthful

Mechanism Design as a Game

- Example: a variant of meeting scheduling
  - Alternatives: 9am, 10am
  - Three agents with their preference orderings
  - Plurality rule, tie breaking in favor of earlier time
  - *What is the dominant strategy? Is the mechanism strategy-proof?*

Mechanism Design as a Game

- Example: a variant of meeting scheduling
  - Alternatives: 9am, 10am
  - Three agents with their preference orderings
  - Plurality rule, tie breaking in favor of earlier time
  - An agent’s report only matters when the others differ
  - Truthful reporting is a dominant strategy

Implementation by a Mechanism

- Social choice function $f(\succeq) \in A$ / $f(v) \in A$

Map true preferences of agents into an alternative

Implementation by a Mechanism

- A DRM implements a social choice function (SCF) f in dominant strategy, where s^* is the DSE and f is defined as the following:

  1) For a mechanism $M = (g)$ without money,

  $$f(\succeq_1, \cdots, \succeq_n) = g(s_1^*(\succeq_1), \cdots, s_n^*(\succeq_n))$$

  2) For a mechanism $M = (x, t)$ with money,

  $$f(v_1, \cdots, v_n) = x(s_1^*(v_1), \cdots, s_n^*(v_n))$$

Implementation by a Mechanism

- A DRM directly implements the outcome rule g or the choice rule x when the mechanism is strategy-proof (at the truthful DSE)

Direct-revelation mechanisms so far…

How about more complex, indirect mechanisms?

Indirect Mechanisms

- The English auction
- “Priority order”: the mechanism asks each agent in turn on which item they want from what is left
- Multiple rounds of negotiations and lots of computation
- …

What is the search space of a desirable mechanism?

Outline

- Recap: revenue equivalence
- Multi-round auctions
- Direct-revelation mechanism
- Revelation principle
- Design desiderata
- The VCG mechanism

Revelation Principle [Myerson 81]

- The Revelation Principle

For every mechanism where there is a dominant strategy for all players, there is an equivalent, strategy-proof direct-revelation mechanism

Revelation Principle [Myerson 81]

- The Revelation Principle

For every mechanism where there is a dominant strategy for all players, there is an equivalent, strategy-proof direct-revelation mechanism

Proof: (DSE, without money)

Given some mechanism M with a DSE s^*, how can we transform it into a strategy-proof DRM M'?

Revelation Principle [Myerson 81]

Proof: (DSE, without money)

Given some mechanism M with a DSE s^*, how can we transform it into a strategy-proof DRM M'?

(1) On any reported profile $\widehat{\succeq} = (\widehat{\succeq_1}, \cdots, \widehat{\succeq_n})$, let M' run M on input $s^*(\widehat{\succeq}) = (s^*(\widehat{\succeq_1}), \cdots, s^*(\widehat{\succeq_n}))$

Therefore, if reporting $s_i^*(\succeq_i)$ is a dominant strategy in M, then reporting $\succeq_i$ is a dominant strategy in M'

Revelation Principle [Myerson 81]

Proof: (DSE, without money)

Given some mechanism M with a DSE s^*, how can we transform it into a strategy-proof DRM M'?

(2) Let M' output the outcome rule that M outputs

That is $g'(\widehat{\succeq})) = g(s^*(\widehat{\succeq}))$

Revelation Principle [Myerson 81]

In other words, we have M' simulate M, implementing the same SCF

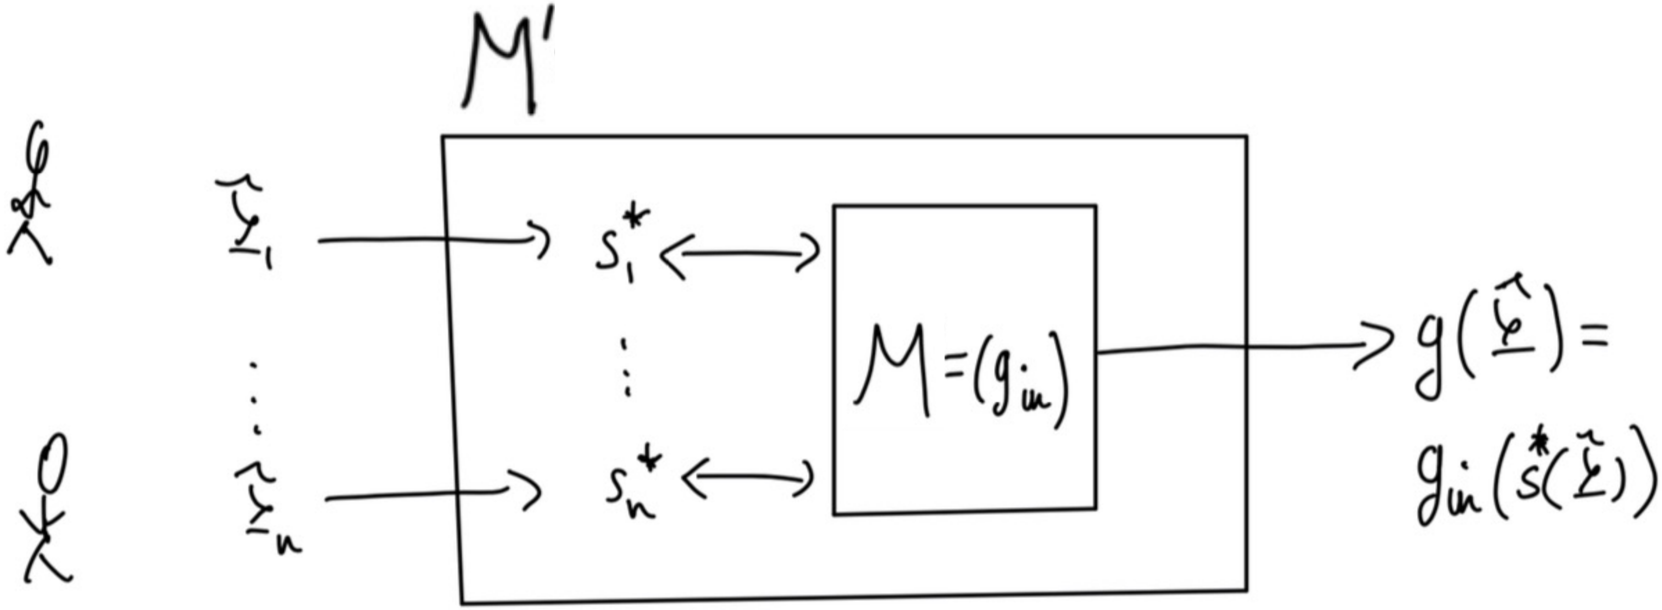

Revelation Principle [Myerson 81]

Example (DSE, with money):

Consider the ascending-clock auction for a single item.

Dominant strategy: stay until price reaches value.

Construct a DRM that simulates the ascending-clock auction, together with the dominant strategy

→ Strategy-proof SPSB auction

Revelation Principle [Myerson 81]

Some more comments:

- Can be extended to BNE
- A powerful theoretical construct: anything that can be achieved in the eq. of mechanism M can be achieved in the truthful eq. of a strategy-proof DRM M'!
- Vice versa: if there is no strategy-proof DRM that implements some SCF f, then it is impossible to implement f in the eq. of an indirect design

Outline

- Recap: revenue equivalence
- Multi-round auctions
- Direct-revelation mechanism
- Revelation principle
- Design desiderata
- The VCG mechanism

Mechanism Desiderata

- Pareto optimality
- Allocative efficiency
- Strategy proofness
- Individual rationality
- No deficit
- Budget balance

Mechanism Desiderata

- Pareto optimality
- Allocative efficiency
- Strategy proofness

(Pareto optimality, Allocative efficiency, Strategy proofness) } equilibrium

- Individual rationality — equilibrium strategy
- No deficit
- Budget balance

Outline

- Recap: revenue equivalence
- Multi-round auctions
- Direct-revelation mechanism
- Revelation principle
- Design desiderata
- The VCG mechanism

VCG Mechanism

- The Vickrey-Clarke-Groves (VCG) mechanism
- A DRM that achieves many good properties
    - Strategy-proof (incentive compatible)
    - Allocative efficient (welfare maximizing)
    - Individually rational

VCG Mechanism

Given reported valuation profile $\hat{v} = (\widehat{v_1}, \ldots, \widehat{v_n})$, the VCG mechanism on a set of alternatives A is defined by

- A choice rule

$$x(\hat{v}) = \mathrm{argmax}_{a \in A} \sum_{i \in N} \widehat{v_i}(a)$$

with selected alternative $\mathrm{a}^* = x(\hat{v})$

VCG Mechanism

Given reported valuation profile $\hat{v} = (\widehat{v_1}, \dots, \widehat{v_n})$, the VCG mechanism on a set of alternatives A is defined by

- A choice rule

$$x(\hat{v}) = \text{argmax}_{a \in A} \sum_{i \in N} \widehat{v_i}(a)$$

  with selected alternative $\mathrm{a}^* = x(\hat{v})$

- A payment rule: charge agent i

$$t_i(\hat{v}) = \max_{\mathrm{a}^{-i} \in A^{-i}} \sum_{j \neq \mathrm{i}} \widehat{v_j}\left(a^{-i}\right) - \sum_{j \neq \mathrm{i}} \widehat{v_j}(a^*)$$

Opportunity cost incurred by agent i = The max total value to others without agent i − The total value to others under a^* without agent i

A^{-i} denotes the set of alternatives when agent i is not present

VCG Mechanism

Example: VCG mechanism on a single item Second-price auction

- Alternatives: "do not allocation" & "allocate to each agent"
- Three agents with their bids \$10, \$8, \$4 for the item
- The choice rule is $x(\hat{v}) = \text{argmax}_{a \in A}\left(\widehat{v_1}(a) + \widehat{v_2}(a) + \widehat{v_3}(a)\right)$
- The payment rule

Agent 1: $t_1(\hat{v})$ = max total value w/o 1 – current total value w/o 1

= 8 - 0 = 8 Pivotal: $a^{-i} \neq a^*$

Agent 2: $t_2(\hat{v})$ = max total value w/o 2 – current total value w/o 2

= 10 - 10 = 0 Non-pivotal: $a^{-i} = a^*$

VCG Mechanism

Example: VCG mechanism on a single item

- Alternatives: “do not allocation” & “allocate to each agent”
- Three agents with their bids \$10, \$8, \$4 for the item
- The choice rule is $x(\hat{v}) = \text{argmax}_{a \in A}\left(\widehat{v_1}(a) + \widehat{v_2}(a) + \widehat{v_3}(a)\right)$
- The payment rule

Agent 1: $t_1(\hat{v})$ = max total value w/o 1 – current total value w/o 1

= 8 - 0 = 8

Agent 2: $t_2(\hat{v})$ = max total value w/o 2 – current total value w/o 2

= 10 - 10 = 0

VCG Mechanism

Example: VCG mechanism on a single item

Second-price auction

- Alternatives: "do not allocation" & "allocate to each agent"
- Three agents with their bids \$10, \$8, \$4 for the item
- The choice rule is $x(\hat{v}) = \mathrm{argmax}_{a \in A}\left(\widehat{v_1}(a) + \widehat{v_2}(a) + \widehat{v_3}(a)\right)$
- The payment rule

Agent 1: $t_1(\hat{v})$ = max total value w/o 1 – current total value w/o 1

= 8 - 0 = 8

Pivotal: $a^{-i} \neq a^*$

Agent 2: $t_2(\hat{v})$ = max total value w/o 2 – current total value w/o 2

= 10 - 10 = 0

Non-pivotal: $a^{-i} = a^*$

VCG Mechanism

Example: VCG mechanism, scheduling

	9am	10am	11am
Agent 1	-5	1	2
Agent 2	20	5	10
Agent 3	5	11	2

What would be the selected alternative?

What would be the payment for each agent?

VCG Mechanism

Theorem: The VCG mechanism is strategy-proof, allocative efficient, and individually rational.

VCG Mechanism

Theorem: The VCG mechanism is strategy-proof, allocative efficient, and individually rational.

Proof (strategy-proof): Being truthful is dominant strategy.

1) Fix other reports v_{-i}; a is the selected alternative under (v_i, v_{-i})

$$v_i(a) - \left(\max_{\mathrm{a^{-i}} \in A^{-i}} \sum_{\mathrm{j} \neq i} v_j\left(a^{-i}\right) - \sum_{j \neq i} v_j(a) \right) = \sum_i v_i(a) - \max_{\mathrm{a^{-i}} \in A^{-i}} \sum_{\mathrm{j} \neq i} v_j\left(a^{-i}\right)$$

2) Fix other reports v_{-i}; a' is the selected alternative under $(v_i{}', v_{-i})$

$$v_i(a') - \left(\max_{\mathrm{a^{-i}} \in A^{-i}} \sum_{\mathrm{j} \neq i} v_j\left(a^{-i}\right) - \sum_{j \neq i} v_j(a') \right) = \sum_i v_i(a') - \max_{\mathrm{a^{-i}} \in A^{-i}} \sum_{\mathrm{j} \neq i} v_j\left(a^{-i}\right)$$

$$\sum_i v_i(a) - \sum_i v_i(a') = \max_{\mathrm{a} \in A} \left(v_i(a) + \sum_{j \neq \mathrm{i}} v_j(a) \right) - \left(v_i(a') + \sum_{j \neq \mathrm{i}} v_j(a') \right) \geq 0$$

VCG Mechanism

Theorem: The VCG mechanism is strategy-proof, allocative efficient, and individually rational.

Proof (allocative efficiency):

This is by construction: $x(\hat{v}) = \mathrm{argmax}_{a \in A} \sum_i v_i(a)$

VCG Mechanism

Theorem: The VCG mechanism is strategy-proof, allocative efficient, and individually rational.

Proof (individual rationality): Agent i's utility of truthfulness

$$v_i(a) - \left(\max_{a^{-i} \in A^{-i}} \sum_{j \neq i} v_j(a^{-i}) - \sum_{j \neq i} v_j(a) \right)$$

$$= \sum_i v_i(a) - \max_{a^{-i} \in A^{-i}} \sum_{j \neq i} v_j(a^{-i})$$

$$\geq \sum_i v_i(a) - \max_{a^{-i} \in A^{-i}} \sum_i v_i(a^{-i})$$

$$\geq 0$$

VCG Mechanism

Recap with some more comments:

- Choose a welfare maximizing outcome
- Charge each agent i the welfare had agent i not participate minus the welfare of everyone else given agent i participates
- Charge each agent the “harm” it does on the welfare of everyone else (i.e., *externality*)

Computational Aspects of VCG

Computational Aspects of VCG

Example: Two items, three bidders

- A wants one apple and is willing to pay $5
- B wants one apple and is willing to pay $2
- C wants two apples and is willing to pay $6 for both but is uninterested in buying one without the other

Computational Aspects of VCG

Example: Two items, three bidders

- A wants one apple and is willing to pay $5
- B wants one apple and is willing to pay $2
- C wants two apples and is willing to pay $6 for both but is uninterested in buying one without the other

Outcome:

A and B get the two apples

A pays $4 = $6 (max value w/o A) - $2 (current value w/o A)

B pays $1 = $6 (max value w/o B) - $5 (current value w/o B)

Computational Aspects of VCG

Example: Two items, three bidders

- A wants one apple and is willing to pay $5
- B wants one apple and is willing to pay $2
- C wants two apples and is willing to pay $6 for both but is uninterested in buying one without the other

What is the computational complexity of finding the welfare-maximizing outcome?

Computational Aspects of VCG

Example: Two items, three bidders

- A wants one apple and is willing to pay \$5
- B wants one apple and is willing to pay \$2
- C wants two apples and is willing to pay \$6 for both but is uninterested in buying one without the other

What is the computational complexity of finding the welfare-maximizing outcome? A knapsack problem!

$$\text{maximize} \sum_{i=1}^{n} v_i x_i$$

$$\text{subject to} \sum_{i=1}^{n} w_i x_i \leq W \text{ and } x_i \in \{0, 1\}$$

Computational Aspects of VCG

Example: Two items, three bidders

- A wants one apple and is willing to pay \$5
- B wants one apple and is willing to pay \$2
- C wants two apples and is willing to pay \$6 for both but is uninterested in buying one without the other

What is the computational complexity of finding the welfare-maximizing outcome? A knapsack problem!

$$\text{maximize} \sum_{i=1}^{n} v_i x_i$$

$$\text{subject to} \sum_{i=1}^{n} w_i x_i \leq W \text{ and } x_i \in \{0, 1\}$$

$n = 3$ bidders

$v_1 = 5, v_2 = 2, v_3 = 6$

$w_1 = 1\ w_2 = 1, w_3 = 2$

$W = 2$

Computational Aspects of VCG

- Compute the welfare-maximizing outcome is NP-hard

- Communication cost of each agent's valuation function
  - Allocate m items to n participants
  - Each agent's valuation function consists 2^m numbers
  - Communication requirement is then $n2^m$

- What if the goal is revenue-maximizing?