Lec1: Introduction to Mechanism Design

Guest Lectures at ZJU Computer Science
(Summer 2024)

Instructor: Haifeng Xu

Outline

- Mechanism Design: Motivation and Examples
- Example Mechanisms for Single Item Allocation
- Example Mechanisms for Multiple Items Allocation

Mechanism Design: the Science of Rule Making

Mechanism Design (MD): designing a game by specifying its rules to induce a desired outcome among strategic participants

- So far, you have seen strategic game and their "equilibrium"
- In mechanism design, you design the game

Mechanism Design: the Science of Rule Making

Mechanism Design (MD): designing a game by specifying its rules to induce a desired outcome among strategic participants

- So far, you have seen strategic game and their “equilibrium”
- In mechanism design, you design the game
  - Specify game rules, player payoffs, allowable actions, etc.
  - Objective is to induce desirable outcome, e.g., incentivizing socially good or fair behaviors, maximizing revenue if selling goods
  - Typically, want the game to be easy to play
    - You don’t want it to be NP-hard for players to act!

Importance of Rule Making: Tale 1

A tale of horse racing

Importance of Rule Making: Tale 1

A tale of horse racing

- Two competitors; each has three horses of different levels: high, medium, low

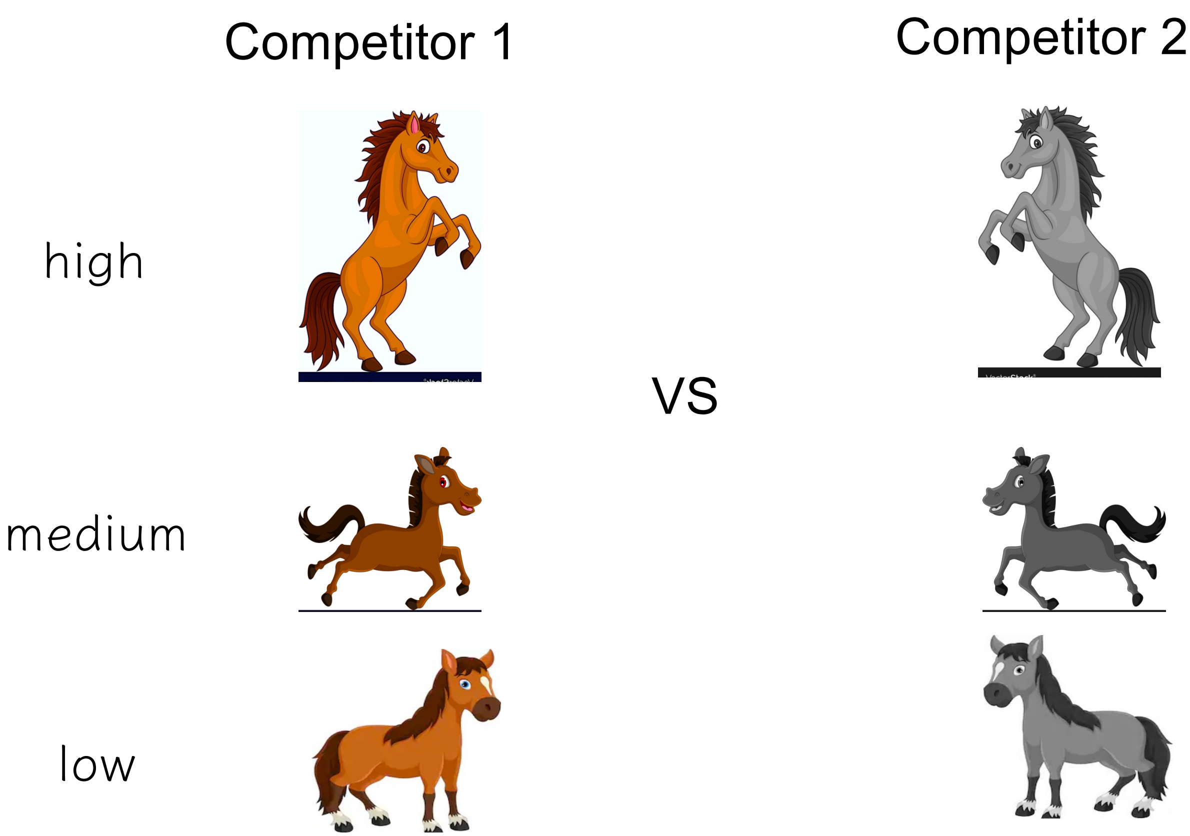

Importance of Rule Making: Tale 1

A tale of horse racing

- Two competitors; each has three horses of different levels: high, medium, low
- They need to compete at each horse level; whoever wins ≥ 2 times is the winner

high

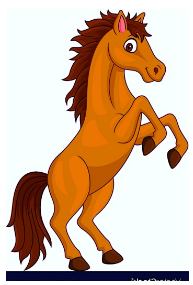

VS

medium

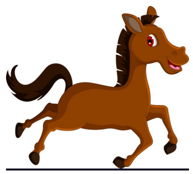

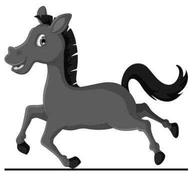

low

Importance of Rule Making: Tale 1

A tale of horse racing

- Two competitors; each has three horses of different levels: high, medium, low
- They need to compete at each horse level; whoever wins ≥ 2 times is the winner
  - Suppose horses are indistinguishable but a horse at a higher level will always beat any horse at a lower level
- Can we truly determine the winner?
  - Both want to use High horse against Medium and Medium against Low

Essentially no, winner will mainly depend on luck

Importance of Rule Making: Tale 1

A tale of horse racing

- Two competitors; each has three horses of different levels: high, medium, low
- What about the following rule?
  - They compete for 3 rounds
  - Winner of first round gains 3 points, winner of second round gains 2 points, and winner of the last round gains 1 point
  - Whoever gets the most points win

This is better; they will really compete at each level – designing right rules is important!

Importance of Rule Making: Tale 2

- Selling products
  - Post a price
  - Customers only get to choose buy or not buy
- Why not the following mechanism?

Importance of Rule Making: Tale 2

- Selling products
  - Post a price
  - Customers only get to choose buy or not buy
- Why not the following mechanism?
  - Customers will not be so honest

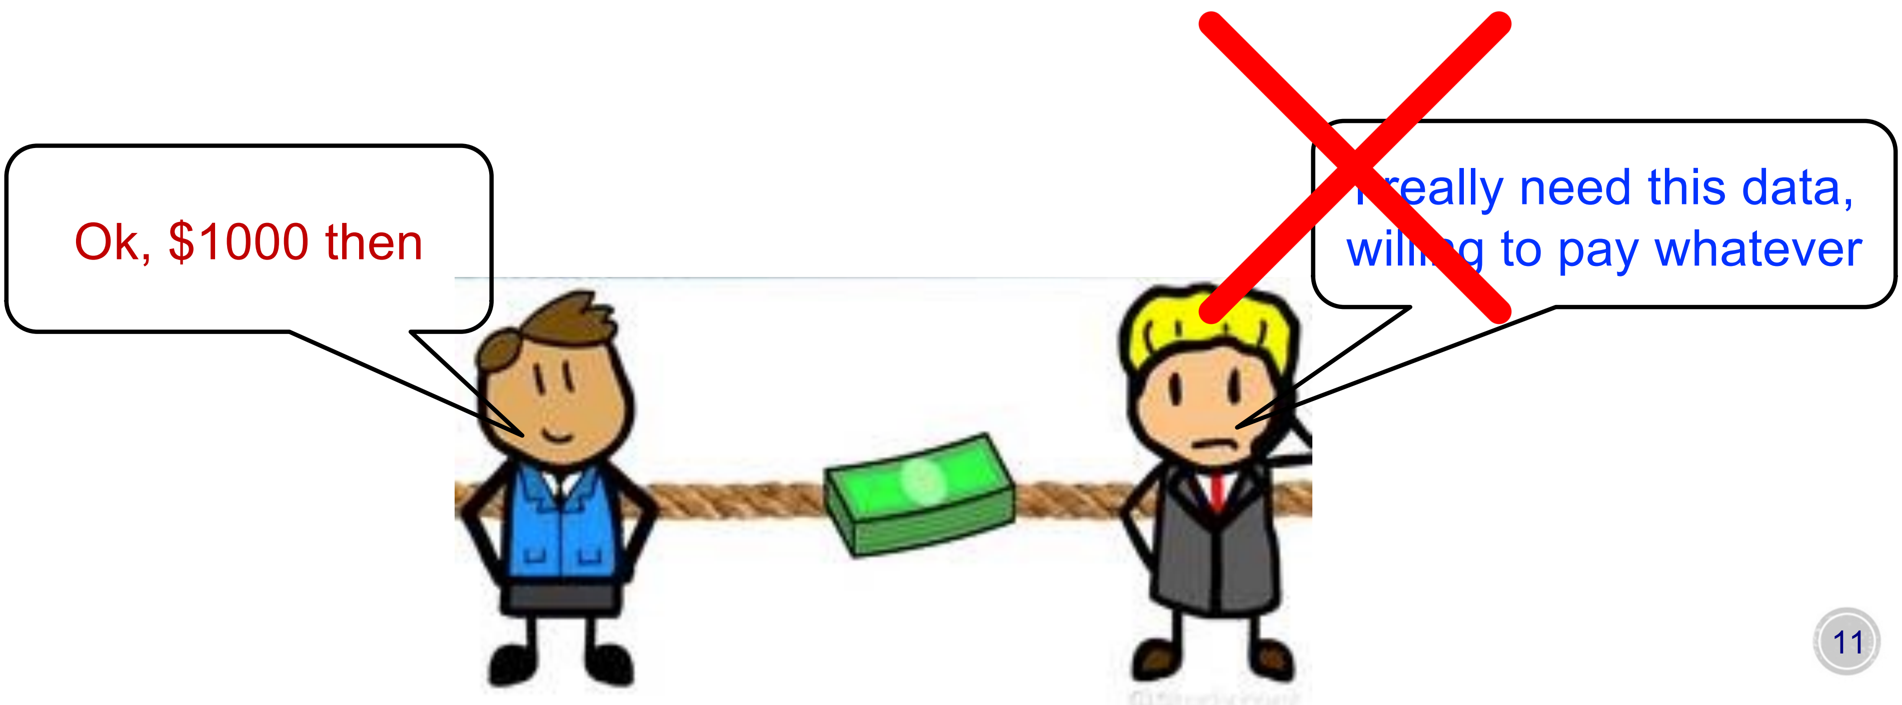

Importance of Rule Making: Tale 2

- Selling products
  - Post a price
  - Customers only get to choose buy or not buy
- Why not the following mechanism?
  - Customers will not be so honest

- Why not the following “bargaining” mechanism?

I have a price in mind, let me see whether your value is higher

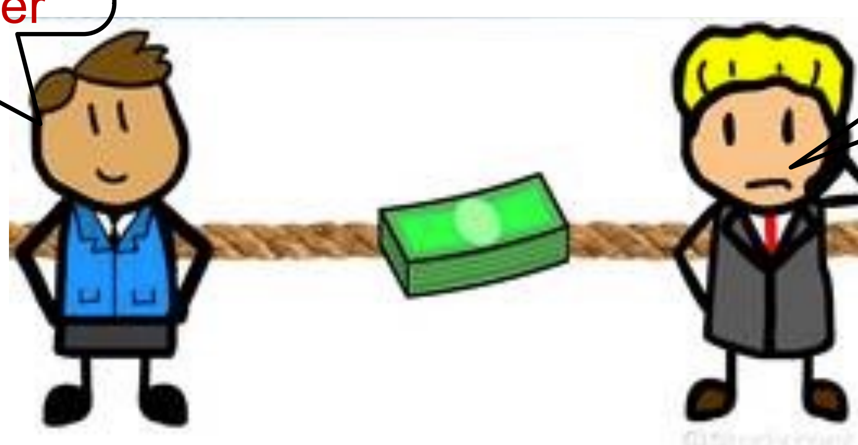

My value is $350

Importance of Rule Making: Tale 2

- Selling products
  - Post a price
  - Customers only get to choose buy or not buy
- Why not the following mechanism?
  - Customers will not be so honest
- Why not the following "bargaining" mechanism?

Ok, I will sell it to you for $350

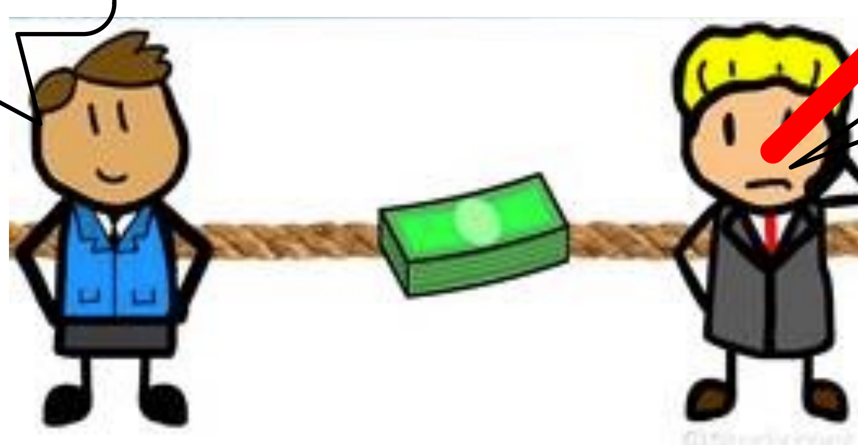

My value is $350

Importance of Rule Making: Tale 2

- Selling products
  - Post a price
  - Customers only get to choose buy or not buy
- Why not the following mechanism?
  - Customers will not be so honest
- Why not the following “bargaining” mechanism?

This is what’s really going to happen…

I have a price in mind, let me see whether your value is higher

My value is $100

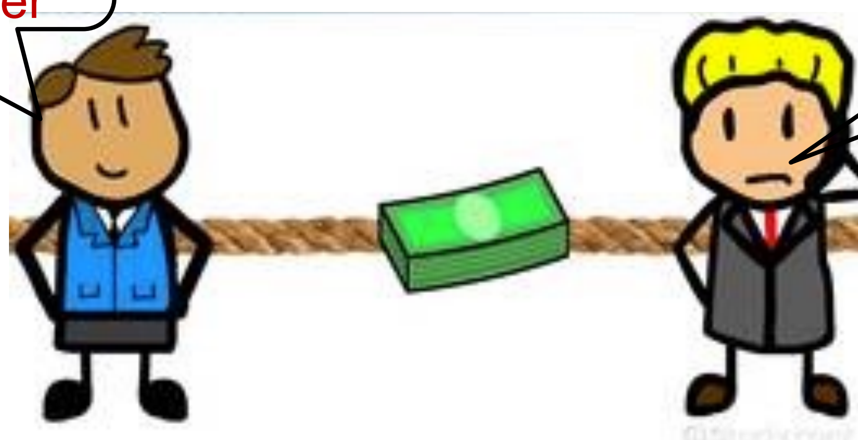

Importance of Rule Making: Tale 2

- Selling products
  - Post a price
  - Customers only get to choose buy or not buy
- Why not the following mechanism?
  - Customers will not be so honest
- Why not the following "bargaining" mechanism?

This is what's really going to happen…

Importance of Rule Making: Tale 2

- Selling products
  - Post a price
  - Customers only get to choose buy or not buy
- Why not the following mechanism?
  - Customers will not be so honest

- Why not the following “bargaining” mechanism?

This is what’s really going to happen…

Importance of Rule Making: Tale 2

- Selling products
  - Post a price
  - Customers only get to choose buy or not buy
- Why not the following mechanism?
  - Customers will not be so honest
- Why not the following “bargaining” mechanism?

This is what’s really going to happen…

Importance of Rule Making: Tale 2

- Selling products
  - Post a price
  - Customers only get to choose buy or not buy

- Why not the following mechanism?
  - Customers will not be so honest
- Why not the following "bargaining" mechanism?

This is what's really going to happen…

Importance of Rule Making: Tale 2

- Selling products
  - Post a price
  - Customers only get to choose buy or not buy
- Why not the following mechanism?
  - Customers will not be so honest
- Why not the following “bargaining” mechanism?

This is what’s really going to happen…

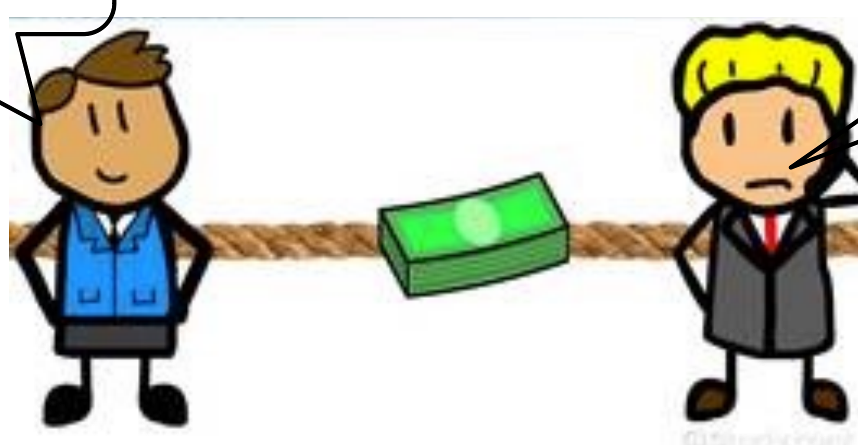

Importance of Rule Making: Tale 2

- Selling products
  - Post a price
  - Customers only get to choose buy or not buy
- Why not the following mechanism?
  - Customers will not be so honest
- Why not the following “bargaining” mechanism?
  - Too complex buyer behaviors, interactions are too time-consuming

Later, we will learn such complexity is not needed – posting a price is optimal among all possible ways of selling to a buyer

Examples of Mechanism Design Problems

Example 1: Single-Item Allocation

v_1

v_2

- A single and indivisible item, n agents
- Agent i has a (private) value v_i about the item
- Outcome: choice of the winner of the item, and possibly payment from each agent
  - Note: payments do not have to involve, e.g., allocating temporary residence to homeless individuals
- Typical objectives: maximize revenue, maximize social welfare (i.e., allocate to the one who values the item most)
- Applications: selling items (e.g., eBay), allocating scarce resources

Example 2: Multi-Item Allocation

  

- m items and n agents
- Agent i has (private) value $v_i(S)$ for any subset of items $S \subseteq [m]$
- Outcome: a partition of the items $[m]$ into $S_1, S_2, \dots, S_n$ and agent i gets items in set S_i
- Typical objectives: revenue, welfare, fairness
- Applications: rental room assignments, sell multiple products, dividing inheritance, etc.

Example 3: School Choice

- n students, m schools
- Each student has a (private) preference over schools
  - Preference $\neq$ value function as in previous item allocation
- Similarly, each school also has a (private) preference over students
- Outcome: match each student to a school
- Objective: maximize "happiness" or "fairness"
- Applications: school choice, marriage or online dating, job matching, assigning web users to distributed Internet services, etc.

Some Common Features

- Participants have private information (often called private types)
- Design objective depends on the private information
- Usually have to elicit such private information
- Participates are self-interested – they want to maximize their own utilities and may lie about their private information if helpful
  - Will be clear after we introduce mechanisms later

Outline

- Mechanism Design: Motivation and Examples
- Example Mechanisms for Single Item Allocation
- Example Mechanisms for Multiple Items Allocation

Single-Item Allocation

- A single and indivisible item, n agents
- Agent i has a (private) value v_i about the item
- Outcome: choice of the winner of the item, and possibly payment from each agent
- Typical objectives: maximize revenue, maximize social welfare
  - Social welfare equals total utility of all players, which in this case equals the value of the bidder who gets the item

Benign Designer: Welfare Maximization

- Want to give the item to the agent who values it the most, i.e., $i^* = \arg\max_{i \in [n]} v_i$
  - But v_i is i's private information
  - The mechanism needs to elicit this information
  - Do not care about revenue
- Each agent is self-interested and will maximize his own utility $v_i \cdot \mathbb{I}(\, i \text{ receives item}) - p_i$, where p_i is his payment (if any)

Benign Designer: Welfare Maximization

Q: what mechanism would work?

Trial 1: ask i to report his value b_i for all i; give the item to $i^* = \arg\max_i b_i$ (no payment)

- Use b_i because it may not equal v_i since agents may misreport
- Indeed, every one will report ∞

Can be proved that any mechanism without using payment cannot achieve the goal of welfare maximization

Ok, need payment, what is a natural mechanism with payment?

Benign Designer: Welfare Maximization

Q: what mechanism would work?

Trial 2: ask i to report his value b_i for all i; give the item to $i^* = \arg\max_i b_i$ and asks him to pay his own bid b_{i^*}

- This is called first-price auction
  - b_i called the “bid” and agents called the “bidders”
- Would agent report $b_i = v_i$?
  - They don’t want → unnecessarily paying too much
  - They dare not report too small neither → may miss out on the item
  - Lead to very intricate and unpredictable agent behaviors
  - Winner does not necessarily have the highest v_i

Benign Designer: Welfare Maximization

Q: what mechanism would work?

Trial 3: ask i to report his value b_i for all i; give the item to $i^* = \arg\max_i b_i$ and asks him to pay the second highest bid $\max2_{i \in [n]} b_i$

- This is called second-price auction

Fact. Truthful bidding is a dominant strategy equilibrium in second-price auctions.

- That is, bidding true value v_i is optimal for i, *regardless of how other people bid*
  - In such cases, we say the mechanism is Dominant-Strategy Incentive Compatible (DISC)
- Proof intuition: i'th payment does not depend on his own bid

Benign Designer: Welfare Maximization

Q: what mechanism would work?

Trial 3: ask i to report his value b_i for all i; give the item to $i^* = \arg\max_i b_i$ and asks him to pay the second highest bid $\max 2_{i \in [n]} b_i$

- This is called second-price auction

Fact. Truthful bidding is a dominant strategy equilibrium in second-price auctions.

Formal proof:

- Fix a bidder i with true value v_i; let $b^* =$ highest bid among other bidders
- If $b^* < v_i$, any $b_i > b^*$ wins the item and pays b^*. So $b_i = v_i$ is also good
- If $b^* \geq v_i$, i prefers losing. Bidding $b_i = v_i$ indeed will make him lose
- Though i does not know b^*, the reasoning above shows, for whatever b^*, bidding $b_i = v_i$ is always optimal

Benign Designer: Welfare Maximization

Q: what mechanism would work?

Trial 3: ask i to report his value b_i for all i; give the item to $i^* = \arg\max_i b_i$ and asks him to pay the second highest bid $\max2_{i\in[n]} b_i$

- This is called second-price auction

Fact. Truthful bidding is a dominant strategy equilibrium in second-price auctions.

- So we generally expect truthful behaviors in second-price auctions
  - We will then give the item to the one with highest value
- This is the prototype of modern Ad Auctions used by Google, Microsoft, and many other ad exchange platforms
  - Handling gaming behaviors in ad auctions

Maximization: Welfare vs Revenue

- Both are important objectives
- Revenue-maximization turns out to be much more difficult
- Without additional assumptions, cannot obtain any guarantee
  - Typically, need to assume prior knowledge about each bidder's value
- Next, we show a simple example
  - Will see why second-price auction alone does not maximize revenue, and what is missing

Example: Sell to Two Uniform Bidders

- Two bidders; for $i = 1,2$, $v_i \sim U([0,1])$ independently
- What is the expected revenue of second price auction?
  - Since bidders bid truthfully, revenue equals the smaller bidder value

$$Rev = \mathbb{E}_{v_1,v_2} \min(v_1, v_2) = 1/3$$

- Consider the following slight auction variant: highest bidder still wins, but pays $\max(second\ highest\ bid,\ 1/2)$
  - If both v_1, v_2 are less than $1/2$, the item is not sold
  - $1/2$ is called the "reserve price"
  - Truthful bidding is still a dominant strategy (the same proof)

Example: Sell to Two Uniform Bidders

- Two bidders; for $i = 1,2$, $v_i \sim U([0,1])$ independently
- What is the expected revenue of second price auction?
  - Since bidders bid truthfully, revenue equals the smaller bidder value

$$Rev = \mathbb{E}_{v_1,v_2} \min(v_1, v_2) = 1/3$$

- Consider the following slight auction variant: highest bidder still wins, but pays $\max(second\ highest\ bid,\ 1/2)$
- What is the expected revenue of this modified auction

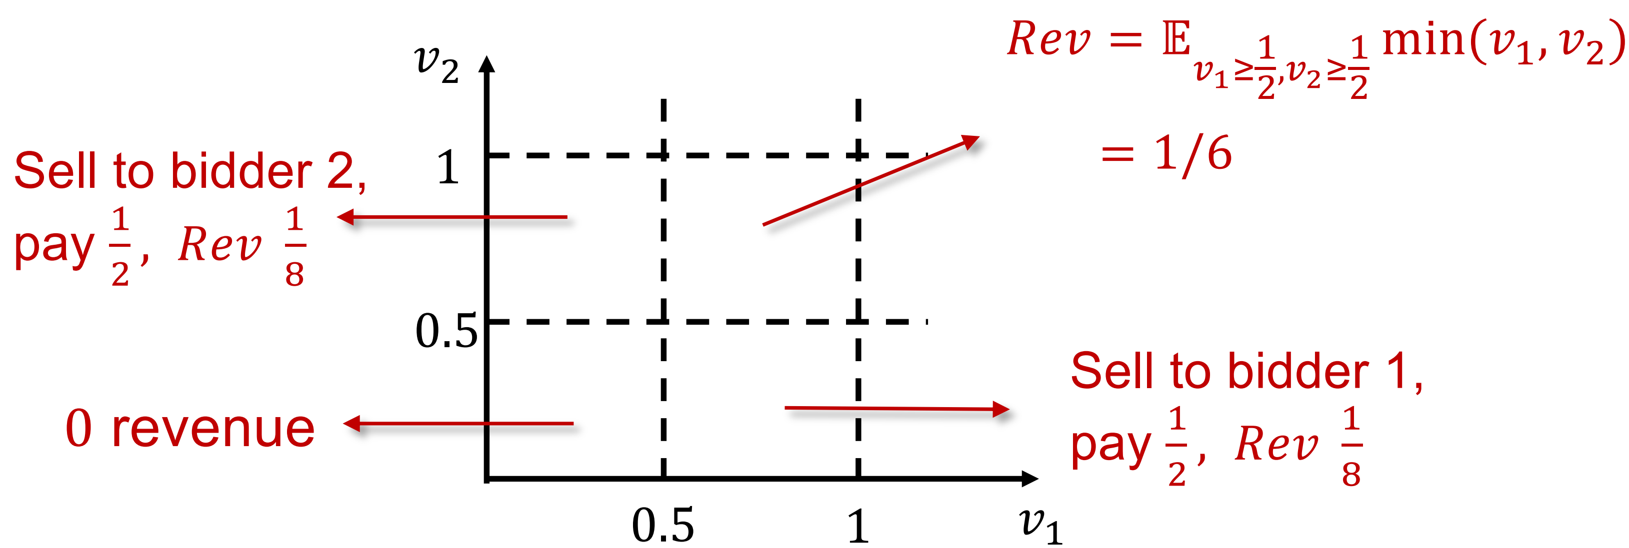

Example: Sell to Two Uniform Bidders

- Two bidders; for $i = 1,2$, $v_i \sim U([0,1])$ independently
- What is the expected revenue of second price auction?
  - Since bidders bid truthfully, revenue equals the smaller bidder value

$$Rev = \mathbb{E}_{v_1,v_2} \min(v_1, v_2) = 1/3$$

- Consider the following slight auction variant: highest bidder still wins, but pays $\max(\mathit{second\ highest\ bid},\ 1/2)$
- What is the expected revenue of this modified auction
  - Total revenue is $\frac{1}{8} + \frac{1}{8} + \frac{1}{6} = \frac{5}{12}$, which turns out to be optimal revenue
  - Second price auction is not optimal because it charges too little when $v_1 > 1/2 > v_2$
  - ½ here is not arbitrary → it equals $\arg\max_{x\in[0,1]} x(1 - F(x))$

Outline

- Mechanism Design: Motivation and Examples
- Example Mechanisms for Single Item Allocation
- Example Mechanisms for Multiple Items Allocation

Multi-Item Allocation

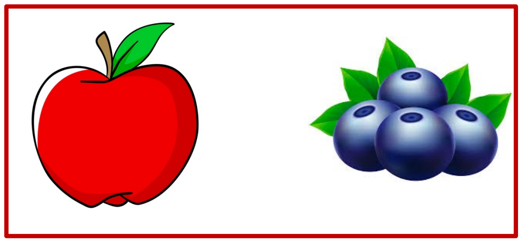

S_1

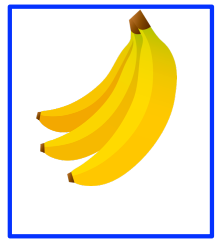

S_2

$v_1(S)$

$v_2(S)$

- m items and n agents
- Agent i has (private) value $v_i(S)$ for any subset of items $S \subseteq [m]$
- Outcome: a partition of the items $[m]$ into $S_1, S_2, \dots, S_n$ and agent i gets items in set S_i
- Typical objectives: revenue, welfare, fairness
  - Revenue-maximizing is extremely challenging (still a major open question despite extensive study)
  - But welfare maximization can be solved via an elegant generalization of second-price auction

Multi-Item Allocation: Welfare Maximization

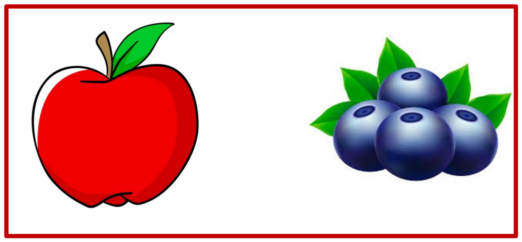

S_1^*

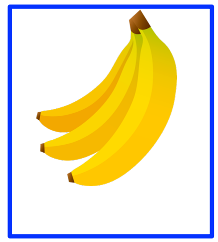

S_2^*

$v_1(S)$

$v_2(S)$

- The Vickrey-Clarke-Groves (VCG) mechanism
  1. Ask each bidder to report their value function $b_i(S)$
  2. Compute optimal allocation $(S_1^*, \cdots, S_n^*) = \arg \max_{(S_1,\cdots,S_n)} \sum_{i=1}^{n} b_i(S_i)$
  3. Allocate S_i^* to bidder i, charge i the following amount

$$p_i = \underbrace{\left[\max_{S_{-i}} \sum_{j \neq i} b_j(S_j) \right]}_{\text{Maximum welfare without } i} - \underbrace{\sum_{j \neq i} b_j(S_j^*)}_{\text{Current welfare without } i \text{ and } S_i^*}$$

p_i = how much i “hurts” all the others’ welfare due to his participation

Multi-Item Allocation: Welfare Maximization

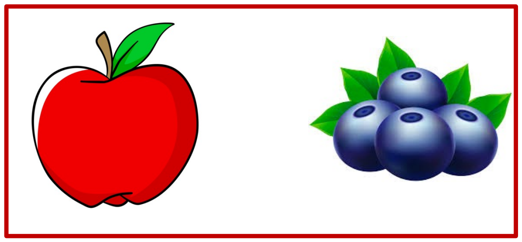

S_1^*

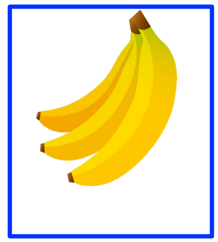

S_2^*

$v_1(S)$

$v_2(S)$

- The Vickrey-Clarke-Groves (VCG) mechanism
  1. Ask each bidder to report their value function $b_i(S)$
  2. Compute optimal allocation $(S_1^*, \cdots, S_n^*) = \arg\max_{(S_1,\cdots,S_n)} \sum_{i=1}^n b_i(S_i)$
  3. Allocate S_i^* to bidder i, charge i the following amount

$$p_i = \left[\max_{S_{-i}} \sum_{j \neq i} b_j(S_j)\right] - \sum_{j \neq i} b_j(S_j^*)$$

Q: what is p_i if there is only a single item for sale?

1. The item will be allocated to largest $b_i(\text{item})$
2. Winner pays the second highest bid; others pay 0
3. Degenerate to a second price auction

Multi-Item Allocation: Welfare Maximization

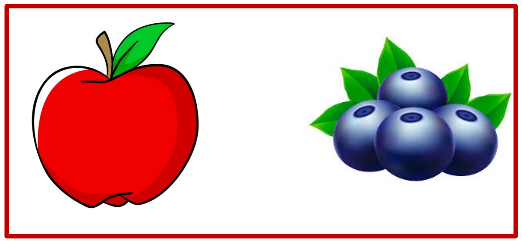

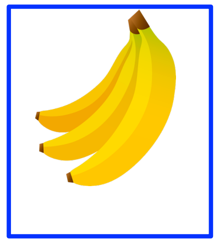

S_1^* S_2^* $v_1(S)$ $v_2(S)$

- The Vickrey-Clarke-Groves (VCG) mechanism
  1. Ask each bidder to report their value function $b_i(S)$
  2. Compute optimal allocation $(S_1^*, \cdots, S_n^*) = \arg\max_{(S_1,\cdots,S_n)} \sum_{i=1}^n b_i(S_i)$
  3. Allocate S_i^* to bidder i, charge i the following amount

$$p_i = \left[\max_{S_{-i}} \sum_{j \neq i} b_j(S_j)\right] - \sum_{j \neq i} b_j(S_j^*)$$

Fact. Truthful bidding is a dominant strategy equilibrium in VCG.

- VCG is DISC, and it does maximize welfare at equilibrium
- Proof: almost the same as truthfulness of second price (HW exercise)

Thank You

Haifeng Xu

University of Chicago

haifengxu@uchicago.edu

Appendix: the Mechanism Design Problem

Mechanism Design for Single-Item Allocation

Described by $\langle n, V, X, u \rangle$ where:

- $[n] = \{1, \cdots, n\}$ is the set of n participants
- $V = V_1 \times \cdots \times V_n$ is the set of all possible value profiles
- $X = \{e_0, e_1, \cdots, e_n\}$ is the set of all possible allocation outcomes
- $u = (u_1, \cdots, u_n)$ where $u_i = v_i x_i - p_i$ is the utility function of i for any outcome $x \in X$ and payment p_i required from i

Remarks:

- Assume risk neural players – i.e., all players maximize expected utilities
- Will guarantee $\mathbb{E}[u_i] \geq 0$ (a.k.a., individually rational or IR)
  - Otherwise, players would not bother coming to your auction even