Auctions 2

Michal Jakob
Artificial Intelligence Center,
Dept. of Computer Science and Engineering,
FEE, Czech Technical University

CGT Autumn 2023- Lecture 9

Efficiency of Single-Item Auctions?

Efficiency in single-item auctions: the item allocated to the agent who values it the most.

With independent private values (**IPV**):

Auction	Efficient
English (without reserve price)	yes
Japanese	yes
Dutch	no
Sealed bid second price	yes
Sealed bid first price	no

Note: Efficiency (often) lost in the **correlated** value setting.

Optimal Auctions

Optimal Auction Design

The seller's problem is to **design an auction mechanism** which has a Nash equilibrium giving him/her the **highest possible expected utility**.

- assuming individual rationality

Second-prize sealed bid auction **does not maximize** expected revenue → not the best choice if profit maximization is important (in the short term).

Designing an Optimum Auction

We assume the **IPV setting** and **risk-neutral bidders**.

Each bidder i's valuation is drawn from some **strictly increasing** cumulative density function $F_i(v)$, having probability density function $f_i(v)$ that is continuous and bounded below.

- Allow $F_i(v) \neq F_j(v)$: **asymmetric** valuations

The **risk neutral** seller knows each F_j and has **zero value** for the object.

The auction that maximizes the **seller's expected revenue** subject to **individual rationality** and **Bayesian incentive-compatibility** for the buyers is an **optimal auction**.

Example

2 bidders, v_i **uniformly** distributed on $[0,1]$.

Second-price sealed bid auction.

Outcome without reserve price

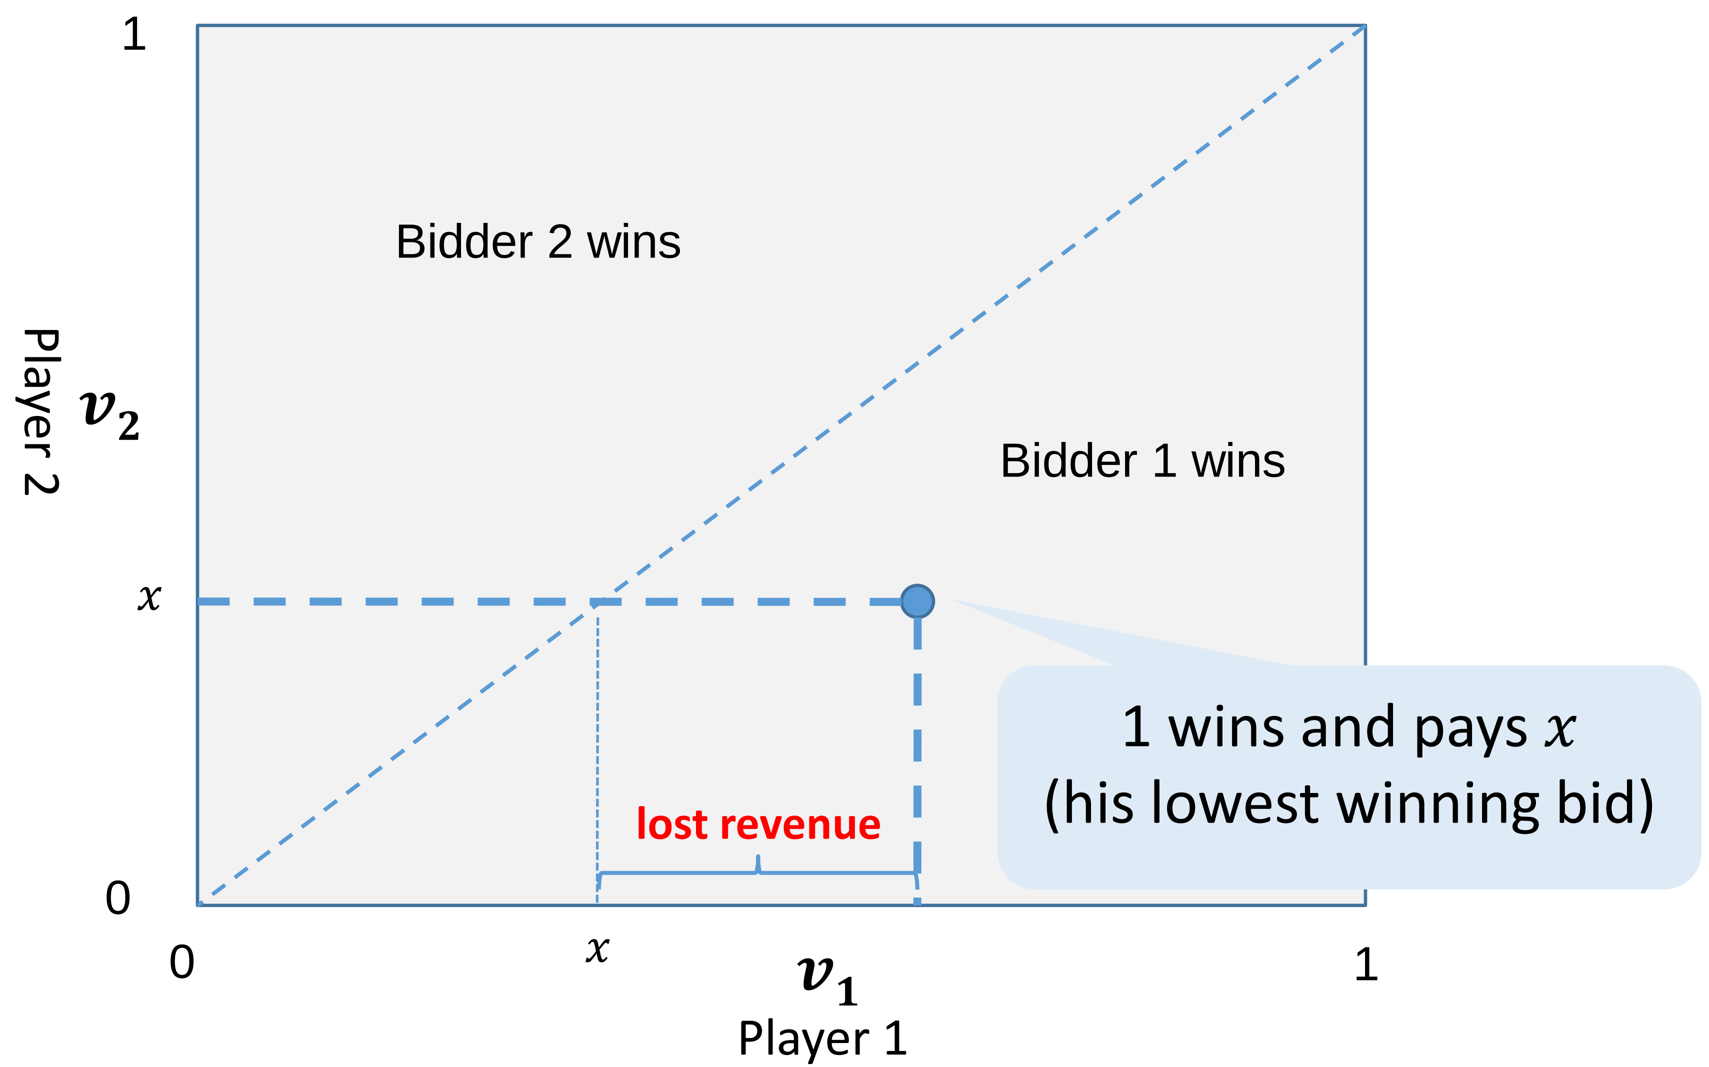

Outcome with reserve price

Some reserve price **improves revenue**.

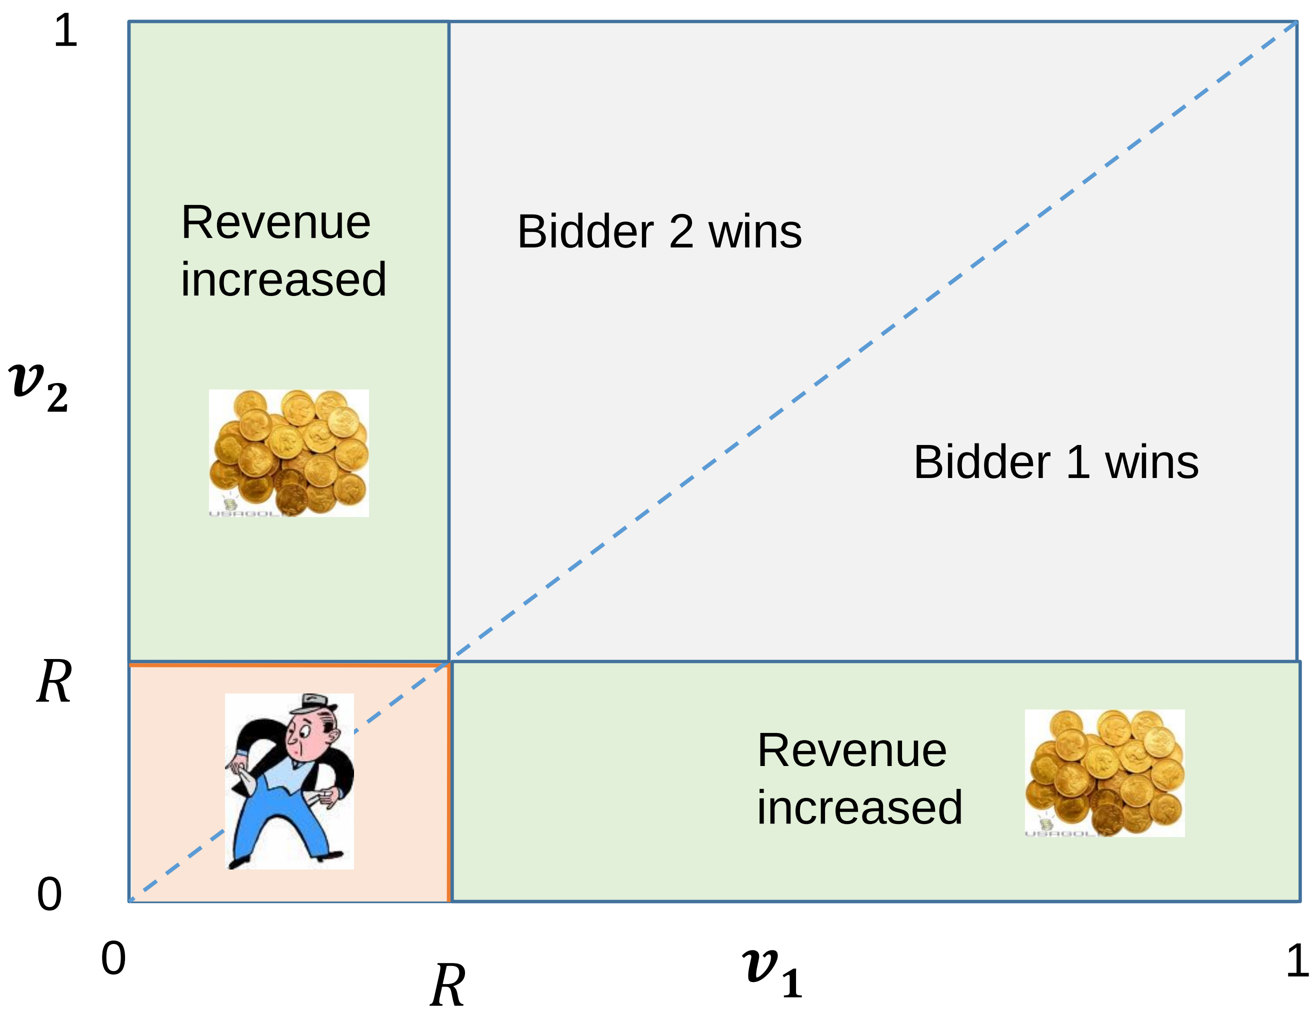

Outcome with reserve price

Bidding true value is still the dominant strategy, so:

1. [Both bides below R]: **No sale.**
   This happens with probability R^2 and then **revenue**=0
2. [One bid above the reserve and the other below]: Sale at **reserve price R**
   This happens with probability $2(1-R)R$ and the **revenue**= R
3. [Both bids above the reserve]: Sale at the **second highest bid**.
   This happens with probability $(1-R)^2$ and the
   revenue= $E[\min v_i \,|\, \min v_i \geq R] = \frac{1+2R}{3}$

Expected **revenue** $= 2(1-R)R^2 + (1-R)^2\frac{1+2R}{3}$

$$= \frac{1 + 3R^2 - 4R^3}{3}$$

Maximizing: $0 = 2R - 4R^2$, i.e., $R = \frac{1}{2}$

Outcome with reserve price

Reserve price of $1/2$: $\texttt{revenue}= 5/12$

Reserve price of 0: $\texttt{revenue}= 1/3 = 4/12$

Tradeoffs:

- **Lose the sale** when both bids below $1/2$: but low revenue then in any case and probability $1/4$ of happening.
- **Increase the sale price** when one bidder has low valuation and the other high: happens with probability $1/2$.

Setting a reserve price is like **adding another bidder**: it increases competition in the auction.

Optimal Single Item Auction

Definition (Virtual valuations)

Consider an **IPV setting** where bidders are **risk neutral** and each bidder i's valuation is drawn from some **strictly increasing** cumulative density function $F_i(v)$, having probability density function $f_i(v)$. We then define:
where

- Bidder i's **virtual valuation** is $\psi_i(v_i) = v_i - \frac{1-F_i(v_i)}{f_i(v_i)}$
- Bidder i's **bidder-specific reserve price** r_i^* is the value for which $\psi_i(r_i^*) = 0$

Example: uniform distribution over $[0,1]$: $\psi(v) = 2v - 1$

Example virtual valuation functions

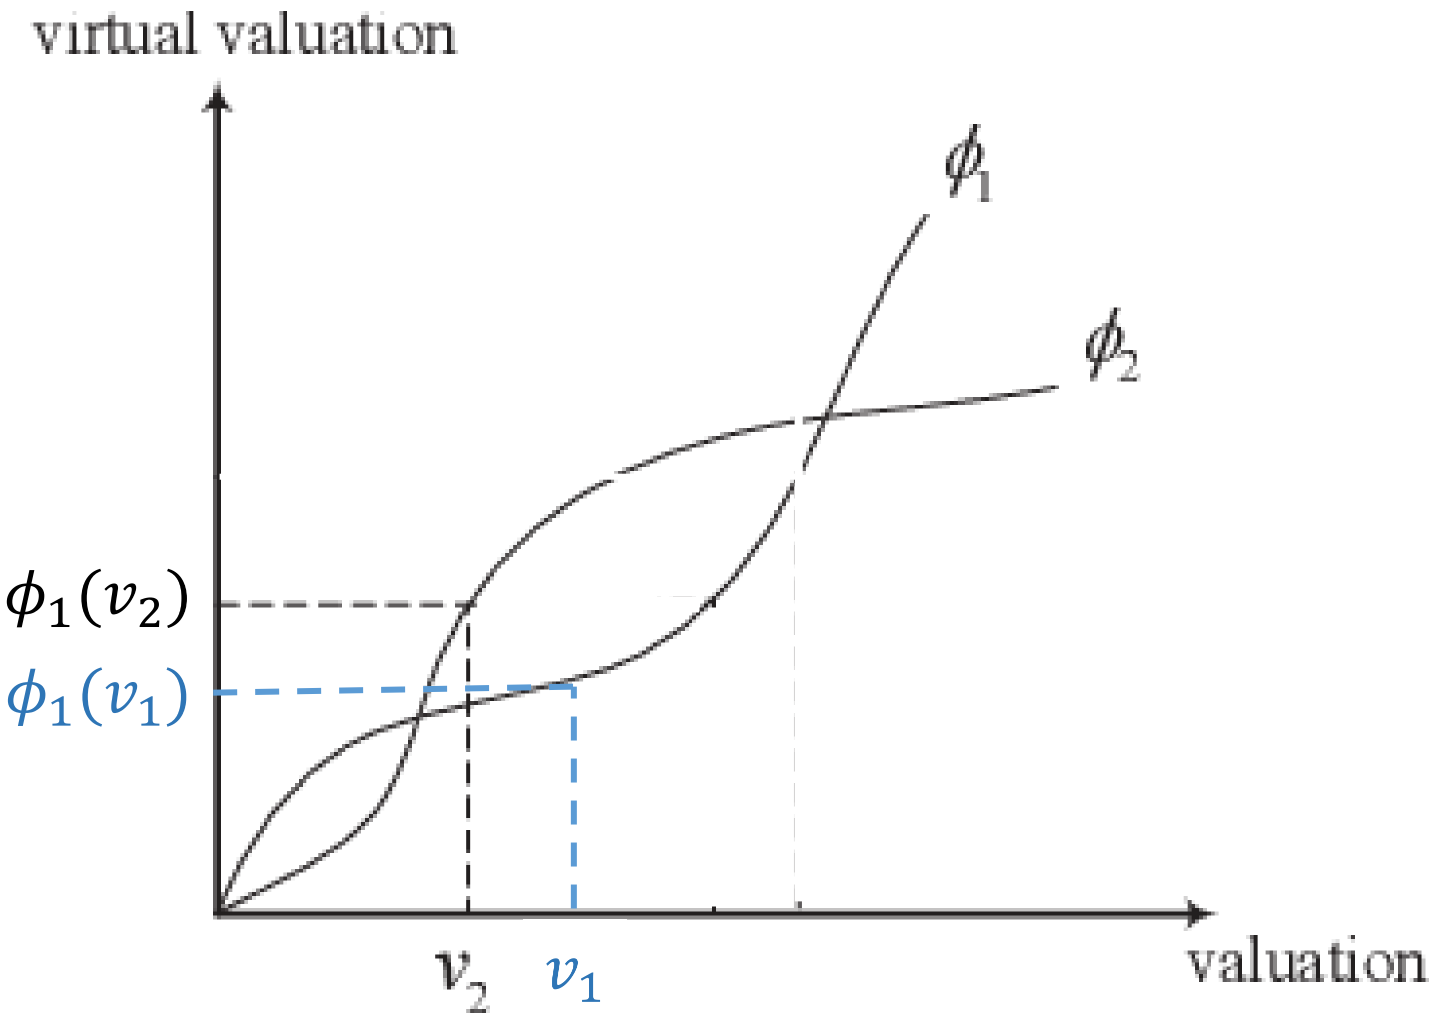

Optimal Single Item Auction

Theorem (Optimal Single-item Auction)

The **optimal (single-good) auction** is a sealed-bid auction in which every agent is asked to **declare his valuation**. The good is sold to the agent $\boldsymbol{i} = \mathbf{argmax_i}\boldsymbol{\psi_i}(\boldsymbol{\hat{v}_i})$**,** as long as $\hat{v}_i > r_i^*$.
If the good is sold, the winning agent i is charged the smallest valuation that it could have declared while still remaining the winner:

$$\inf\{v_i^* : \psi_i(v_i^*) \geq 0 \wedge \forall j \neq i, \psi_i(v_i^*) \geq \psi_j(\hat{v}_j)\}$$

Can be understood as a second-price auction with a reserve price, held **in virtual valuation space** rather than in the space of actual valuations.

Remains **dominant-strategy truthful.**

Second-Price Auction with Reservation Price

Symmetric case: second-price auction with reserve price r^* satisfying: $\psi(r^*) = r^* - \frac{1-F(r^*)}{f(r^*)} = 0$

- **Truthful** mechanism when $\psi(v)$ is non-decreasing.
- Uniform distribution over $[0, p]$: optimum reserve price $= p/2$.

Second-price sealed bid auction with Reserve Price is **not efficient**!

Second-Price Auction with Reservation Price

Why does this increase revenue?

- Reservation prices are like **competitors**: increase the payments of winning bidders.
- The virtual valuation can increase the impact of weak bidders' bids, making the **more competitive**.
- Bidders with higher expected valuations bid **more aggressively**.

Optimal Auctions: Remarks

For **optimal revenue** one needs to **sacrifice** some **efficiency**.

Optimal auctions are not **detail-free:**

- they require the seller to incorporate information about the bidders' valuation distributions into the mechanism
- → rarely used in practice

Theorem (Bulow and Klemperer): *revenue* of an efficiency-maximizing auction with *k+1* bidder is at least as high as that of the revenue-maximizing one with *k* bidders.

➔ better to spend energy on attracting more bidders

Multi-unit Auctions

Multi-unit Auctions

Multiple identical copies of the same good on sale.

Multi-unit Japanese auction:

- After each increment, the bidder specifies the amount he is willing to buy at that price
- The amount needs to decrease over time: cannot buy more at a higher pirce
- The auction is over when the supply equals or exceeds the demand.
  - Various options if supply exceeds demand

Similar extension possible for English and Dutch auctions.

Single-unit Demand

Assume there are k identical goods on sale and risk-neutral bidders who only want one unit each.

$k+1^{st}$-price auction is the equivalent of the second-price auction: sell the units to the k highest bidders for the same price, and to set this price at the amount offered by the highest losing bid.

Note: Seller will not always make higher profit by selling more items! Example:

Bidder	Bid amount
1	$25
2	$20
3	$15
4	$8

Combinatorial Auctions

Auctions for **bundles of goods.**

Let $\mathcal{G} = \{g_1, \dots, g_n\}$ be a set of items (goods) to be auctioned

A **valuation function** $v_i: 2^{\mathcal{G}} \mapsto \mathbb{R}$ indicates how much a bundle $G \subseteq \mathcal{G}$ is worth to agent i.

We typically assume the following properties:

- **normalization**: $v(\emptyset) = 0$
- **free disposal**: $G_1 \subseteq G_2$ implies $v(G_1) \leq v(G_2)$

Example

Buying a computer gaming rig: PC, Monitor, Keyboard and mouse.

Different types/brands available for each category of items.

Non-Additive Valuations

Combinatorial auctions are interesting when the valuation function is **not additive.**

Two main types on non-additivity.

Substitutability	Complementarity
The valuation function v exhibits **substitutability** if there exist two sets of goods $G_1, G_2 \subseteq G$ such that $G_1 \cap G_2 = \emptyset$ and $v(G_1 \cup G_2) < v(G_1) + v(G_2)$. Then this condition holds, we say that the valuation function v is **subadditive**.	The valuation function v exhibits **complementarity** if there exist two sets of goods $G_1, G_2 \subseteq G$ such that $G_1 \cap G_2 = \emptyset$ and $v(G_1 \cup G_2) > v(G_1) + v(G_2)$. Then this condition holds, we say that the valuation function v is **superadditive**.
Ex: Two different brands of TVs.	Ex: Left and right shoe.

How to Sell Goods with Non-Additive Valuations?

1. Ignore valuations dependencies and sell sequentially via a sequence of **independent single-item** auctions.

   → **Exposure problem**: A bidder may bid aggressively for a set of goods in the hope of winning a bundle but only succeed in winning a subset (a thus paying too much).

2. Run separate but **connected single-item** auctions **simultaneously**.
   - a bidder bids in one auction he has a reasonably good indication of what is transpiring in the other auctions of interest.

3. **Combinatorial auction**: bid directly on a **bundle of goods.**

Allocation in Combinatorial Auction

Allocation is a list of sets $G_1, \ldots, G_n \subseteq \mathcal{G}$, one for each agent i such that $G_i \cap G_j = \emptyset$ for all $i \neq j$ (i.e. not good allocated to more than one agent)

Which way to choose an allocation for a combinatorial auction?

→ The simples is to maximize **social welfare (efficient allocation)**:

$$U(G_1, \ldots, G_n, v_1, \ldots, v_n) = \sum_{i=1}^{n} v_i(G_i)$$

Simple Combinatorial Auction Mechanism

The mechanism determines the **social welfare maximizing allocation** and then **charges** the winners their **bid** (for the bundle they have won), i.e., $\rho_i = \hat{v}_i$.

Example:

Bidder 1	**Bidder 2**	**Bidder 3**
$v_1(x, y) = 100$	$v_2(x) = 75$	$v_3(y) = 40$
$v_1(x) = v_1(y) = 0$	$v_2(x, y) = v_2(y) = 0$	$v_3(x, y) = v_3(x) = 0$

Is this incentive-compatible? **No.**

VCG auction

A **Vickrey–Clarke–Groves (VCG) auction** is a type of sealed-bid auction of multiple items. Bidders submit bids that report their valuations for the items, without knowing the bids of the other bidders. The auction system assigns the items in a socially optimal manner: it charges each individual the harm they cause to other bidders.[1]

Vickrey–Clarke–Groves (VCG) auction, an analogy to **second-price** sealed bid single-unit auctions, exists for the combinatorial setting and it **is dominant-strategy truthful** and **efficient**.

VCG example

Suppose two apples are being auctioned among three bidders.

- **Bidder A** wants one apple and is willing to pay **$5** for that apple.
- **Bidder B** wants one apple and is willing to pay **$2** for it.
- **Bidder C** wants two apples and is willing to pay **$6** to have both of them but is uninterested in buying only one without the other.

First, the outcome of the auction is determined by maximizing social welfare:

- the **apples go to bidder A and bidder B**, since their combined bid of **$5 + $2 = $7** is greater than the bid for two apples by bidder C who is willing to pay only **$6**.
- Thus, after the auction, the value achieved by bidder **A is $5**, by bidder **B is $2**, and by bidder **C is $0** (since bidder C gets nothing).

VCG example

Payment of bidder **A**:

- an auction that excludes bidder A, the social-welfare maximizing outcome would assign both apples to bidder C for a total social value of $6.
- the total social value of the original auction *excluding A's value* is computed as **$7 - $5 = $2**.
- Finally, subtract the second value from the first value. Thus, the payment required of A is **$6 – $2 = $4**.

Payment of bidder **B**:

- the best outcome for an auction that excludes bidder B assigns both apples to bidder **C for $6**.
- The total social value of the original auction *minus B's portion* is $5. Thus, the payment required of B is **$6 – $5 = $1**.

Finally, the payment for bidder **C** is **($5 + $2) – ($5 + $2) = $0.**

After the auction, A is $1 better off than before (paying $4 to gain $5 of utility), B is $1 better off than before (paying $1 to gain $2 of utility), and C is neutral (having not won anything).

Winner Determination Problem

Definition

The **winner determination problem** for a combinatorial auctions, given the agents' declared valuations $\hat{v}_i$ is to find the social-**welfare-maximizing allocation** of goods to agents. This problem can be expressed as the following integer program

$$\text{maximize} \sum_{i \in N} \sum_{Z \subseteq \mathcal{Z}} \hat{v}_i(Z) x_{Z,i}$$

$$\text{subject to} \sum_{Z, j \in Z} \sum_{i \in N} x_{Z,i} \leq 1 \quad \forall j \in \mathcal{Z}$$

$$\sum_{Z \subseteq \mathcal{Z}} x_{Z,i} \leq 1 \qquad \forall i \in N$$

$$x_{Z,i} = \{0,1\} \qquad \forall Z \subseteq \mathcal{Z}, i \in N$$

Complexity of the Winner Determination Problem

Equivalent to a **set packing problem** (SSP) which is known to be **NP-complete**.

Worse: SSP cannot be **approximated uniformly** to a fixed constant.

Two possible solutions:

- **Limit** to instance where polynomial-time solutions exist.
- **Heuristic methods** that drop the *guarantee* of polynomial runtime, optimality or both.

Restricted instances

Use **relaxation** to solve WDP in polynomial time: Drop the integrality constraint and solve as a **standard** linear program.

The solution is guaranteed to be integral when the constraints matrix is **unimodular**.

Two important real-world cases fulfills this condition.

Contiguous ones property (continuous bundles of goods)

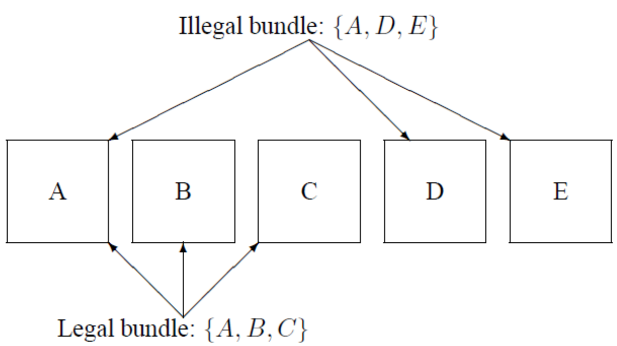

Tree-structured bids

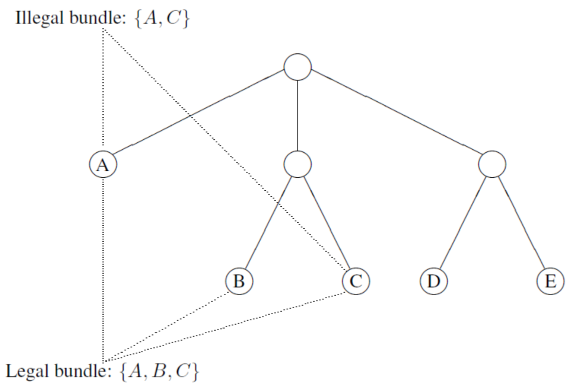

Heuristics Methods

Incomplete methods **do not guarantee** to find optimal solution.

Methods do exist that can **guarantee** a solution that is within $1/\sqrt{k}$ of the optimal solution, where k is the number of goods.

Works well in practice, making it possible to solve WDPs with **many hundreds of goods** and **thousands of bids**.

Auctions Summary

Auctions are mechanisms for **allocating scarce resource** among **self-interested agent**

Mechanism-design and game-theoretic perspective

Many auction mechanisms: English, Dutch, Japanese, First-price sealed bid, Second-price sealed bid

Desirable properties: truthfulness, efficiency, optimality, ...

Rapidly expanding list of **applications** worth billions of dollars

Reading:

- [Shoham] – Chapter 11
- [Maschler] – Chapter 12