Auctions

Michal Jakob
Agent Technology Center,
Dept. of Computer Science and Engineering,
FEE, Czech Technical University

AE4M36MAS Autumn 2014 - Lecture 12

Where are We?

Agent architectures (inc. BDI architecture)

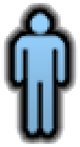

Logics for MAS

Non-cooperative game theory

Cooperative game theory

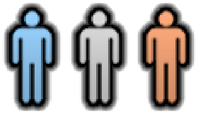

Resource allocation and Auctions

Social choice

Distributed constraint reasoning

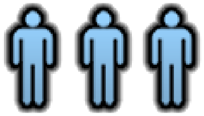

Lecture Online [TODO]

Introduction

Resource Allocation

- Type of resources
- Preference representation
- Social Welfare

Auction Mechanisms

- Basic Definitions
- Single-good auction mechanisms
- Analysis of auction mechanisms

What is an Auction?

*An **auction** is a protocol that allows agents (=bidders) to indicate their **interests** in one or more **resources** and that uses these indications of interest to determine both an **allocation** of resources and a set of **payments** by the agents.* [Shoham & Leyton-Brown 2009]

Why Auctions?

Market-based price setting: for objects of unknown value, the value is dynamically assessed by the market!

Flexible: any object type can be allocated

Can be **automated**

- use of simple rules reduces complexity of negotiations
- well-suited for computer implementation

Revenue-maximising and **efficient allocations** are achievable

Basic Single-Item Auction Mechanisms

English

Japanese

Dutch

First-Price

Second-Price

Analysing Auctions

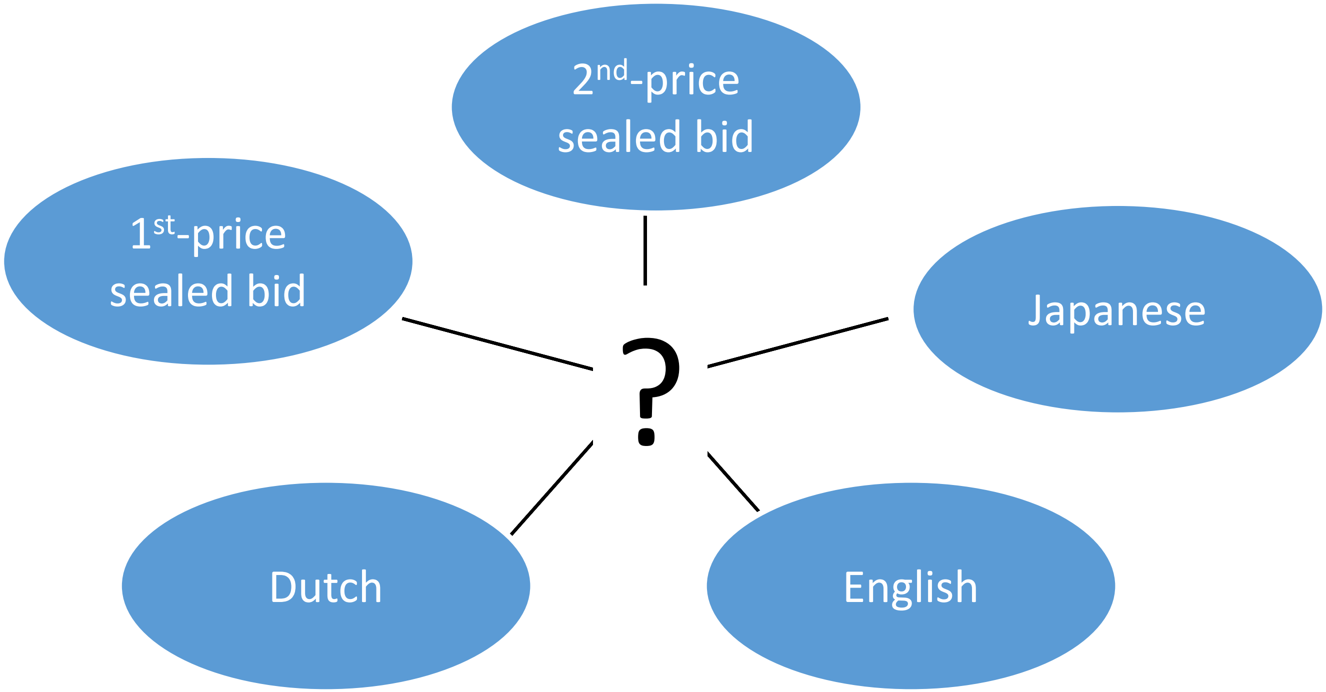

Are there fundamental similarities / differences between mechanisms described?

Mechanism Design: A *Very* Brief Intro

Bayesian Game

Definition (Bayesian game setting)

A Bayesian game setting is a tuple (N, O, Θ, p, u), where

- N is a finite set of n agents;
- O is a set of outcomes;
- $\Theta = \Theta_1 \times \cdots \times \Theta_n$ is a set of possible joint type vectors;
- p is a (common prior) probability distribution on Θ; and
- $u = (u_1, \ldots, u_n)$, where $u_i : O \times \Theta \mapsto \mathbb{R}$ is the utility function for each player i.

Mechanism

Definition (Mechanism)

A mechanism (for a Bayesian game setting (N, O, Θ, p, u)) is a pair (A, M), where

- $A = A_1 \times \cdots \times A_n$, where A_i is the set of actions available to agent $i \in N$; and
- $M : A \mapsto \Pi(O)$ maps each action profile to a distribution over outcomes.

Implementation

Definition (Implementation in dominant strategies)

Given a Bayesian game setting (N, O, Θ, p, u), a mechanism (A, M) is an implementation in dominant strategies of a social choice function C (over N and O) if for any vector of utility functions u, the game has an equilibrium in dominant strategies, and in any such equilibrium a^* we have $M(a^*) = C(u)$.

Definition (Bayes–Nash implementation)

Given a Bayesian game setting (N, O, Θ, p, u), a mechanism (A, M) is an implementation in Bayes–Nash equilibrium of a social choice function C (over N and O) if there exists a Bayes–Nash equilibrium of the game of incomplete information (N, A, Θ, p, u) such that for every $\theta \in \Theta$ and every action profile $a \in A$ that can arise given type profile θ in this equilibrium, we have that $M(a) = C(u(\cdot, \theta))$.

Quasilinear Preferences

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

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

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

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 $(A, \chi, \boldsymbol{p})$, 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
- $\boldsymbol{p} : A \mapsto \mathbb{R}^n$ maps each action profile to a payment for each agent.

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 $(\chi, \boldsymbol{p})$. It defines a standard mechanism in the quasilinear setting, where for each i, $A_i = \Theta_i$.

Quasilinear Mechanisms with Conditional Utility Independence

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 $u_i(o, \theta) = u_i(o, \theta')$.

Given conditional utility independence, we can write i's utility $u_i(o, \theta)$ function as $u_i(o, \theta_i)$

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

Alternative 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
- also define tuples $\hat{v}$ and $\hat{v}_{-i}$

Direct Mechanism Redefined

Alternative 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
- also define tuples $\hat{v}$ and $\hat{v}_{-i}$

Mechanism Properties

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 $\hat{v}_i = v_i$.

Definition (Efficiency)

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

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

Others: Budget balance, Ex interim / Ex post individual rationality. tractability, ...

Design Objectives Mechanism

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.

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.

Analysing Auctions

Two Problems

Auction **mechanism analysis**

- determine the properties of a given auction mechanism
- methodology: treat auctions as (extended-form) *Bayesian games* and analyse players’ (i.e. bidders’) strategies

Auction **mechanism design**

- design the auction mechanism (i.e. the game for the bidders) with the desirable properties
- methodology: apply mechanism design techniques

(Desirable) Properties

Truthfulness: bidders are incentivized to bid their true valuations

Efficiency: the aggregated utility of bidders is maximized

Optimality: maximization of seller's revenue

Strategy: existence of a dominant strategy

Manipulation vulnerability: lying auctioner, shills, bidder collusion

Other consideration: communication complexity, private information revelation, ...

Second-Price Sealed Bid

Theorem

Truth-telling is a **dominant strategy** in a second-price sealed bid auction (assuming independent private values (IPV) model and risk neutral bidders).

Proof: Assume that the other bidders bid in some arbitrary way. We must show that i's best response is always to bid truthfully. We'll break the proof into two cases:

- Bidding honestly, i would win the auction
- Bidding honestly, i would lose the auction

Second-Price Sealed Bid Proof

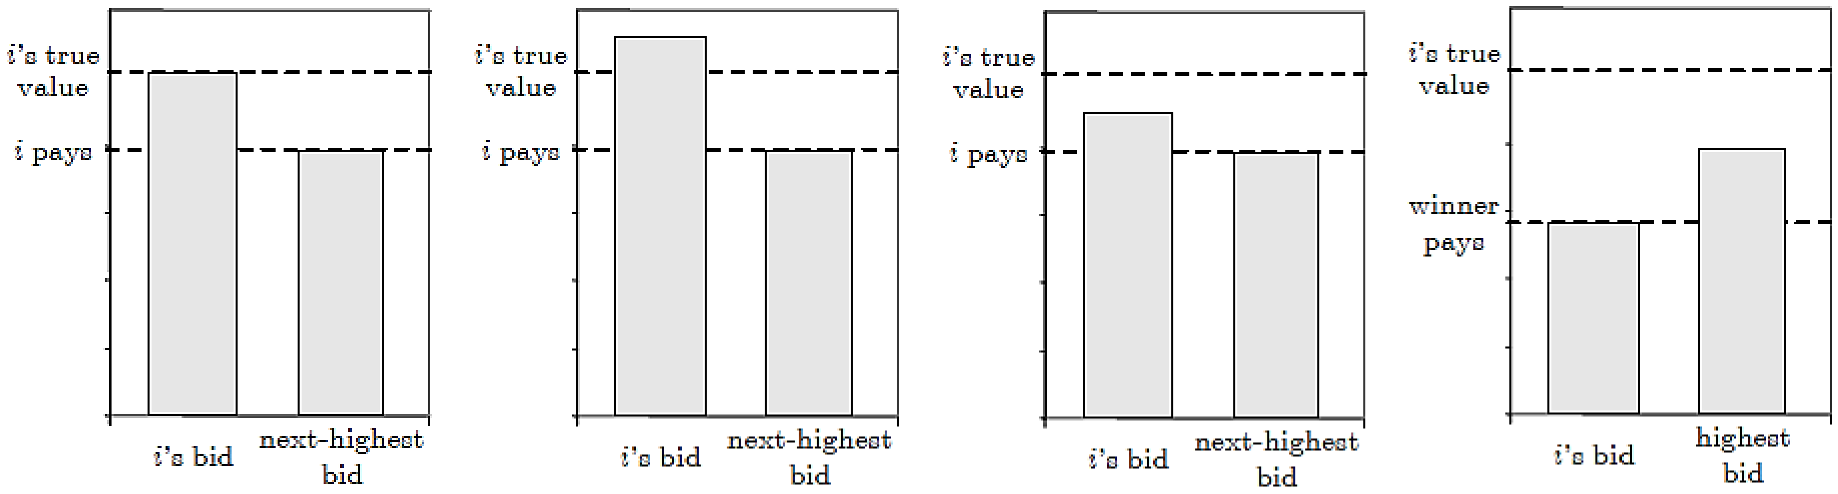

Bidding honestly, i is the winner

If i bids higher, he will still win and still pay the same amount

If i bids lower, he will either still win and still pay the same amount. . .

... or lose and get utility of zero.

Second-Price Sealed Bid Proof

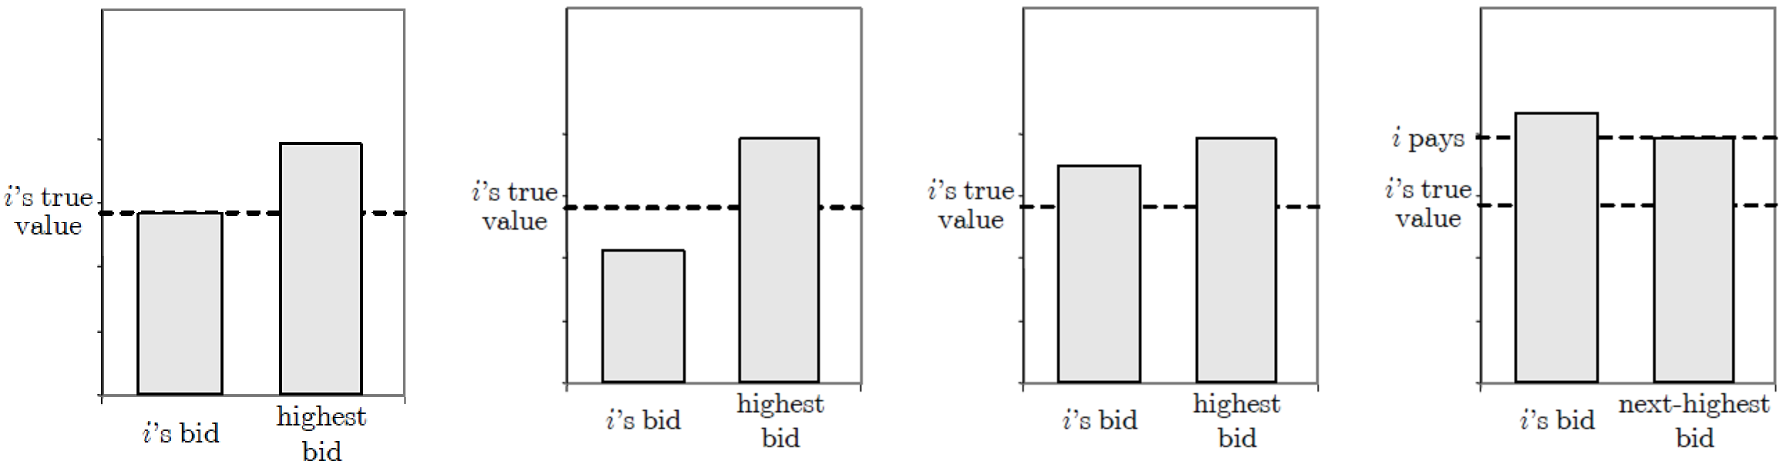

Bidding honestly, i is not the winner

If i bids lower, he will still lose and still pay nothing

If i bids higher, he will either still lose and still pay nothing...

... or win and pay more than his valuation.

Second-Price Sealed Bid

Advantages:

- Truthful bidding is dominant strategy
- No incentive for counter-speculation
- Computational efficiency

Disadvantages:

- Lying auctioneer
- Bidder collusion self-enforcing

Unfortunately, the auction is not very popular in real life due to its counter-intuitiveness

- but very successful in computational auction systems (e.g. Adwords)

Dutch and First-price Sealed Bid

Strategically equivalent: an agent bids without knowing about the other agents’ bids

- a bidder must decide on the amount he's willing to pay, conditional on having placed the highest bid

Differences

- First-price auctions can be held **asynchronously**
- Dutch auctions are **fast**, and require **minimal communication**

Bidding in Dutch / First Price Sealed Bid?

Bidders strategy?

- Bidders would normally bid less than own valuation but just enough to win ⇒ **not incentive compatible** and incentive to counter-speculate

Bidders don't have a **dominant strategy** any more:

- there's a **trade-off** between **probability of winning** vs. **amount paid** upon winning
- **individually optimal** strategy depends on **assumptions** about **others' valuations**

Theorem

In a first-price sealed bid auction with n **risk-neutral** agents whose valuations $v_1, v_2, \dots, v_n$ are **independently** drawn from a **uniform distribution** on the **same** bounded **interval** of the real numbers, the **unique** symmetric **equilibrium** is given by the **strategy profile** $\left(\frac{n-1}{n}v_1, \dots, \frac{n-1}{n}v_n\right)$.

English and Japanese Auctions Analysis

A much more complicated **strategy space**

- extensive-form game
- bidders are able to condition their bids on information revealed by others
- in the case of English auctions, the ability to place jump bids

Intuitively, though, the **revealed information** doesn't make any **difference** in the **independent-private value** (IPV) setting.

- proxy bidding

English and Japanese Auctions Analysis

Theorem

Under the IPV model, it is a **dominant strategy** for bidders to bid **up to** (and not beyond) their valuations in both Japanese and English auctions.

In correlated-value auctions, it can be worthwhile to counter-speculate

Revenue Equivalence

Which auction should an auctioneer choose?

To some extent, it doesn't matter...

Theorem (Revenue Equivalence)

Assume that each of n **risk-neutral** agents has an **independent private valuation** for a single good at auction, drawn from a **common cumulative distribution** $F(v)$ that is **strictly increasing** and **atomless** on $[\underline{v}, \overline{v}]$. Then any auction mechanism in which

1. the good will be allocated to the agent with the highest valuation; and
2. any agent with valuation $\underline{v}$ has an expected utility of zero

yields the **same expected revenue**, and hence results in any bidder with valuation v making the same expected payment.

Applying Revenue Equivalence TODO

Optimal Auctions

Optimal Auction Design

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

- assuming individual rationality

Second-prize sealed bid auction **does not maximize** expected revenue.

Can we get better revenue?

Let's have another look at 2nd price auctions:

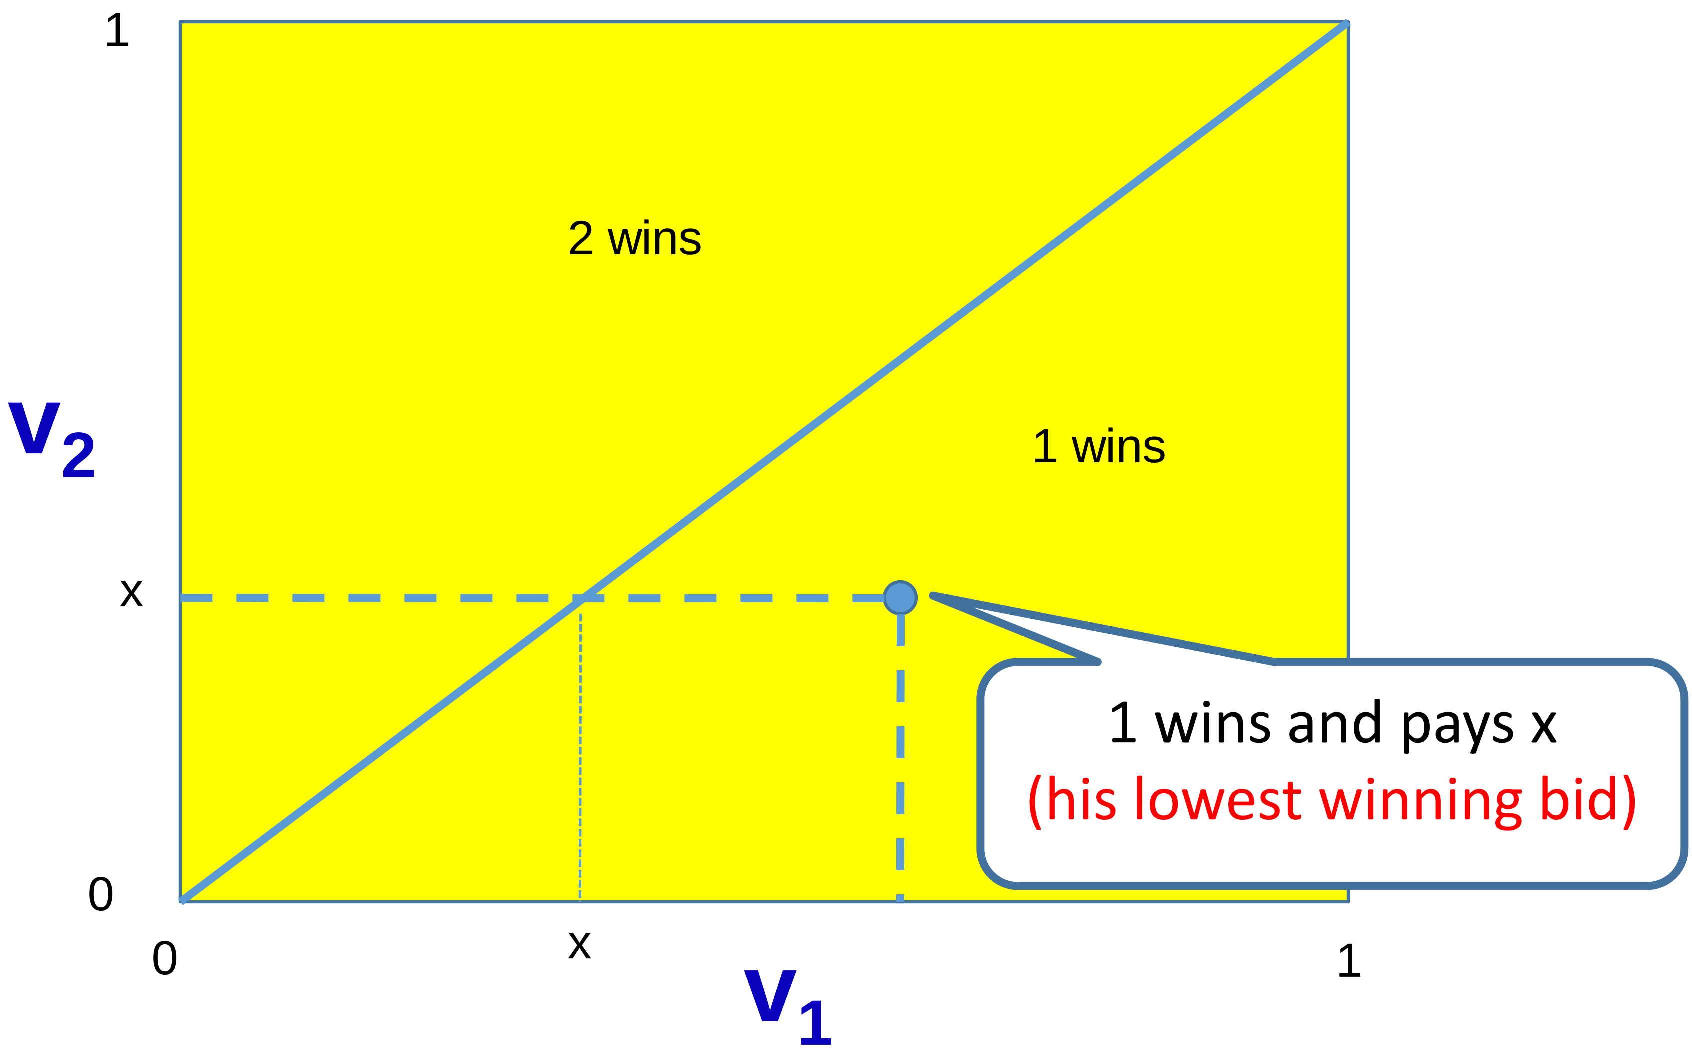

Can we get better revenue?

Some reserve price improve revenue.

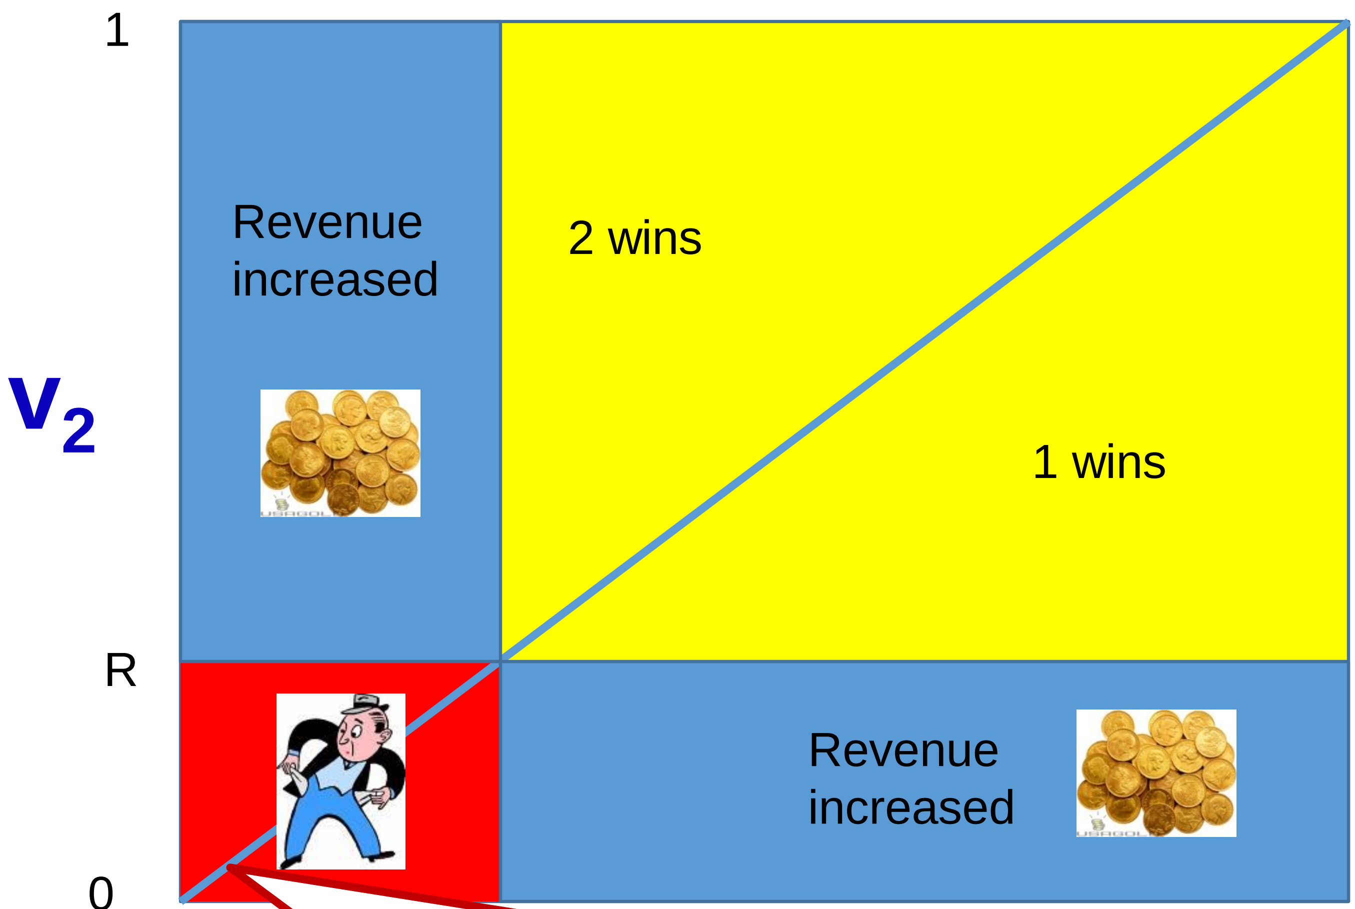

When comparing to the 2nd-price auction with no reserve price: Revenue loss here (efficiency loss too)

Can we get better revenue?

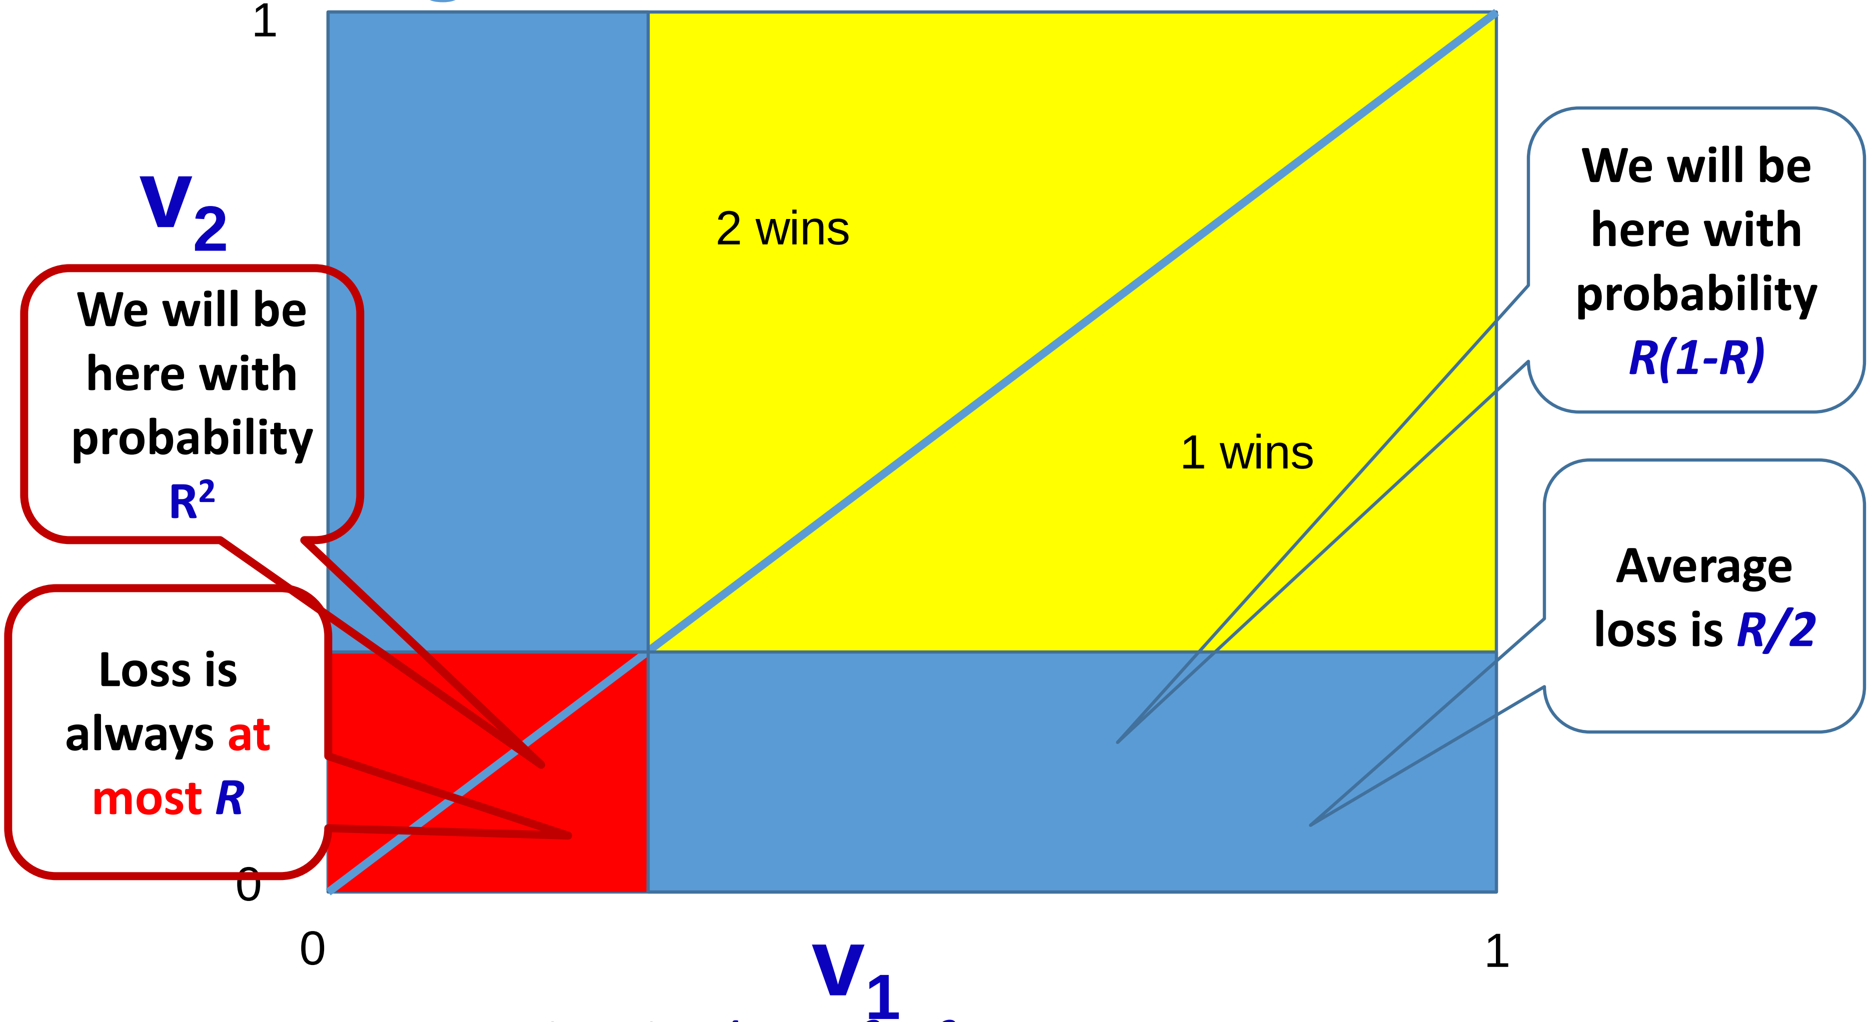

Gain is at least $2R(1-R)\ R/2 = R^2-R^3$

Loss is at most $R^2\ R = R^3$

→ When $R^2-2R^3>0$, reserve price of R is beneficial.

(for example, $R=1/4$)

Reservation price: Single Bidder

How do you sell one item to one bidder?

- Assume his value is drawn uniformly from [0,1].

Optimal way: reserve price.

- Take-it-or-leave-it-offer.

Probability that the buyer will accept th

The payment for the seller

Let's find the optimal reserve price.

E[revenue] = (1-F(R)) × R = (1-R) ×R

$$\frac{\partial(1-R)R}{\partial R} = 1 - 2R = 0$$ → *R=1/2*

Optimal Single Item Auction

Assumptions

- independent private valuations (IPV)
- risk-neutral bidders
- strictly increasing cumulative density function F_i (pdf f_i)

Definition (virtual valuation)

Bidder i's virtual valuation is $\psi_i(v_i) = v_i - \frac{1-F_i(v_i)}{f_i(v_i)}$.

Definition (bidder-specific reserve price)

Bidder i's bidder-specific reserve price r_i^* is the value for which $\psi_i(r_i^*) = 0$.

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

Optimal Single Item Auction

Theorem

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 $i = \arg\max_i \psi_i(\hat{v}_i)$, *as long as* $v_i > r_i^*$. *If the good is sold, the winning agent* i *is charged the smallest valuation that he could have declared while still remaining the winner:* $\inf\{v_i^* : \psi_i(v_i^*) \geq 0$ *and* $\forall j \neq i, \psi_i(v_i^*) \geq \psi_j(\hat{v}_j)\}$.

The virtual valuations also increase weak bidders' bids, making them more competitive.

- Low bidders can win, paying less.
- However, bidders with higher expected valuations must bidmore aggressively

Second-Prized 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 $p/2$.

The SPSB with Reserve Price is **not efficient**!

Optimal Auctions: Remarks

Always: **revenue ≤ efficiency**

- Due to **individual rationality**
- More efficiency makes the pie larger!

However, for optimal revenue one needs to sacrifice some efficiency.

Optimal auctions are not **detailed-free** ➔rarely used in practice

- better to spend energy on attracting more bidders

Multi-Item Auctions

Multi-Item Auctions

Combinatorial Auctions

Auctions for **bundles of goods**

Let $\mathcal{Z} = \{z_1, \ldots, z_n\}$ be a set of items to be auctioned

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

Properties

- **normalization**: $v(\emptyset) = 0$
- **free disposal**: $Z_1 \subseteq Z_2$ implies $v(Z_1) \leq v(Z_2)$

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

- **complementarity**: $v(Z_1 \cup Z_2) > v(Z_1) + v(Z_2)$ (e.g. left and right shoe)
- **substitutability**: $v(Z_1 \cup Z_2) < v(Z_1) + v(Z_2)$ (e.g. cinema tickets for the same time)

Allocation

Allocation is a list of sets $Z_1, \ldots, Z_n \subseteq \mathcal{Z}$, one for each agent i such that $Z_i \cap Z_j = \emptyset$ for all $i \neq j$ (i.e. not good allocated to more than one agent)

Allocation is determined by the auction mechanism

- trivial for single-good auctions

How to define allocation for combinatorial auction?

Maximize **social welfare**: $U(Z_1, \ldots, Z_n, v_1, \ldots, v_n) = \sum_{i=1}^{n} v_i(Z_i)$

Winner Determination Problem

Definition

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

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

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

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

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

Issues with Winner Determination

Communication complexity

Computation complexity

- Solution 1: Require bids to come from a restricted set, guaranteeing that the WDP can be solved in polynomial time
  - problem: these restricted sets are very restricted...
- Solution 2: Use heuristic methods to solve the problem
  - this works pretty well in practice, making it possible to solve WDPs with many hundreds of goods and thousands of bids.

Example Application

Taxi is a scarce resource
Different value of using the taxi

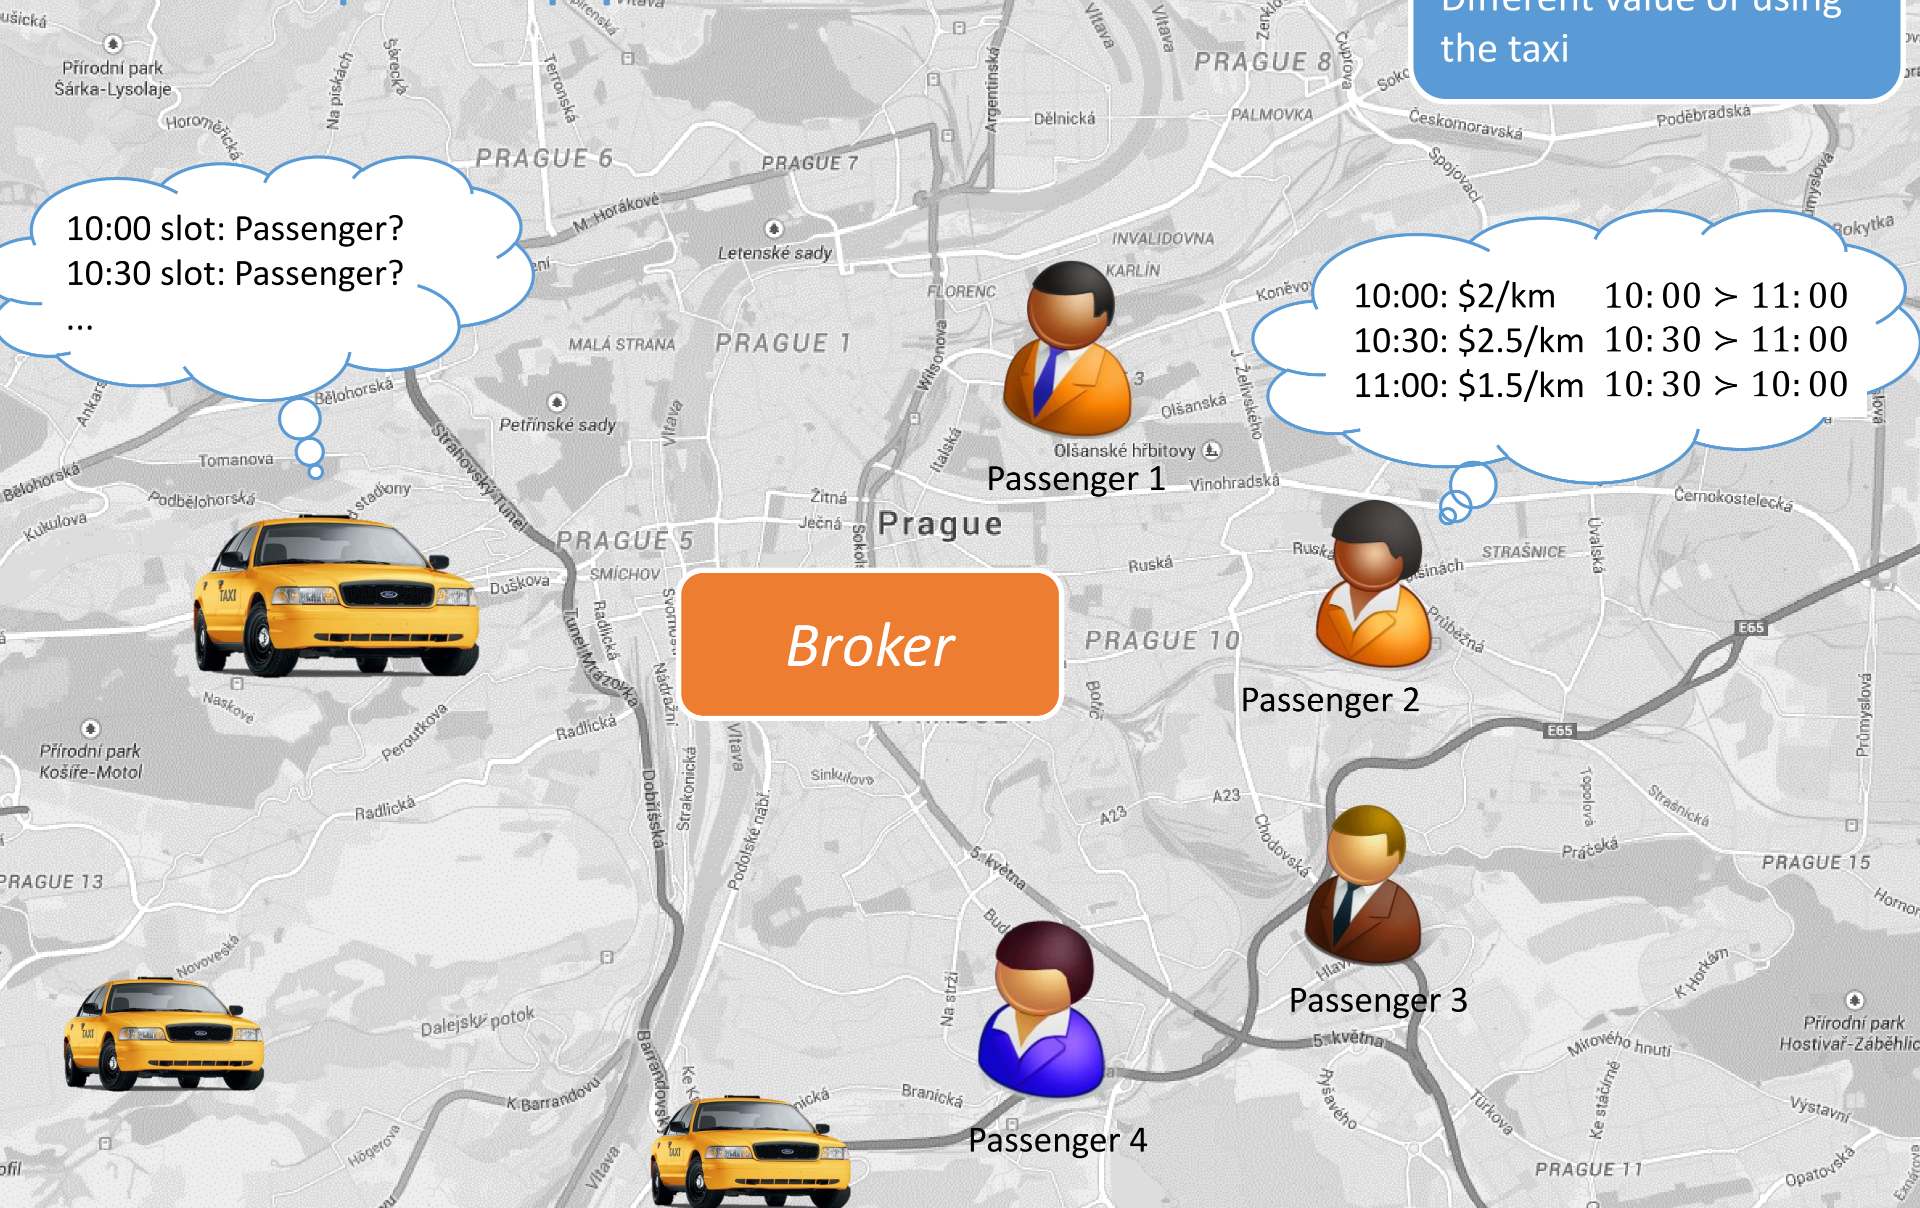

Auctions Summary

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

Mechanism-design and game-theoretic perspective

Vast range of auctions 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

MAS Course Summary

Logics for MAS: Formally describe and analyze (multiple) agents

Agent architectures: acting rationally in an environment

Non-cooperative game theory: acting rationally in strategic interactions

Coalitional game theory: making rational decisions about collaboration

Distributed constraint reasoning: coordinating cooperative action

Social choice: aggregating individual preferences into a collective choice

Multiagent Resource Allocation and Auctions: distributing scarce resources

Many topics not covered: bargaining / negotiation, multiagent learning, multiagent planning, mechanism design, agent-oriented software engineering

Many interconnections

Final Notes

Rapidly evolving field with the exploding number of applications
→ http://agents.cz for (Ph.D.) opportunities

Exam

- 8th Jan + 2 more dates
- mostly written

Survey/Anketa: be as specific possible: we *do* care