Multiagent Resource Allocation 1

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

AE4M36MAS Autumn 2014 - Lecture 11

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

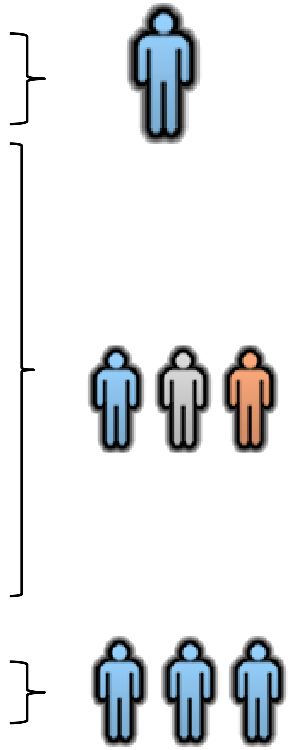

Motivating Example

Taxi is a scarce resource
Different value of using the taxi

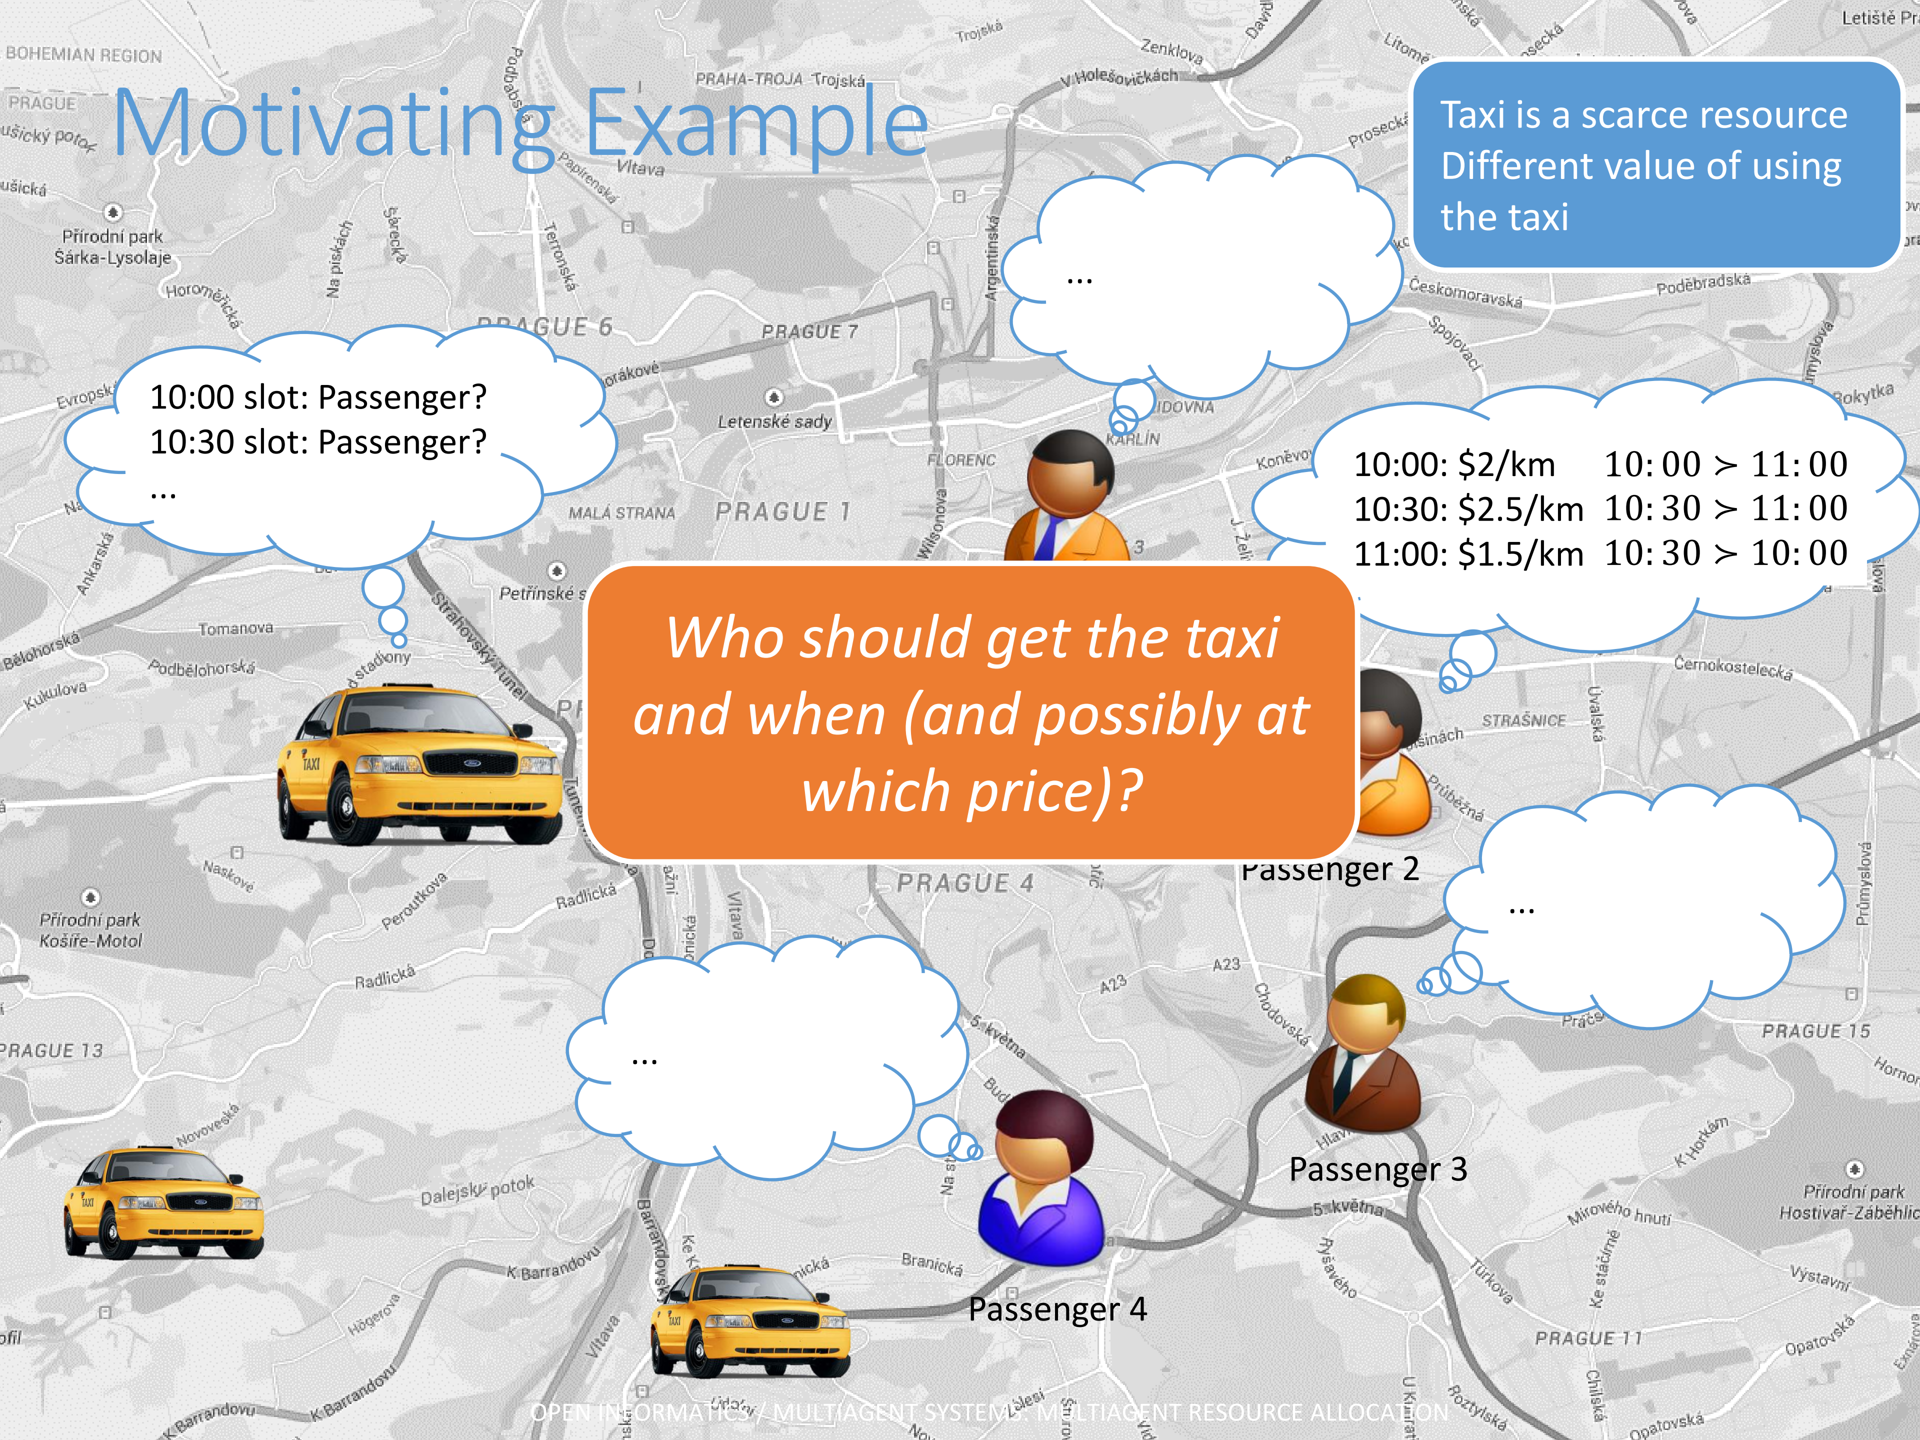

Who should get the taxi and when (and possibly at which price)?

Lecture Online

Introduction

Resource Allocation

- Type of resources
- Preference representation
- Social Welfare

Auction Mechanisms

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

Auctions: Traditional

Auctions used in Babylon as early as 500 B.C. but used to be rare (not so long ago)

Stage 0: No automation

Auctions: Partial Automation

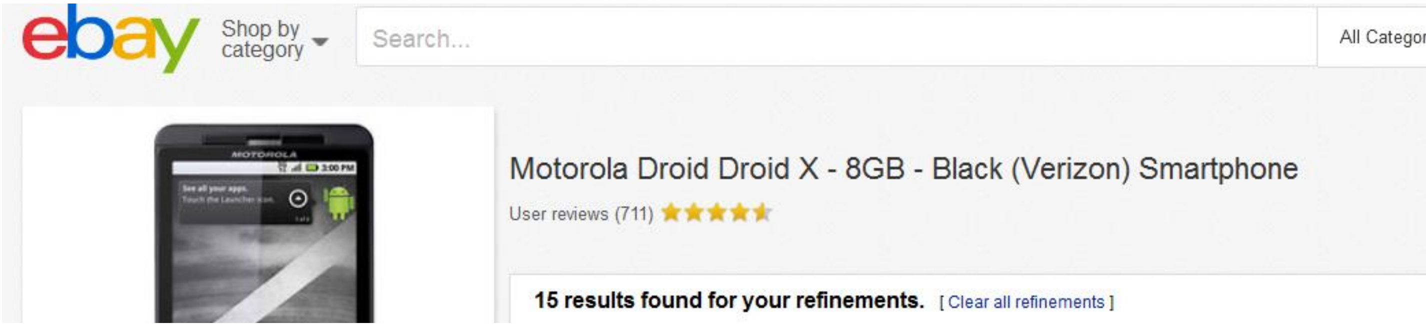

Grown massively with the Web/Internet

→ **Frictionless commerce**: feasible to auction things that weren't previously profitable

Phase 1: Computers used to manage auctions / **run** auction **protocols**

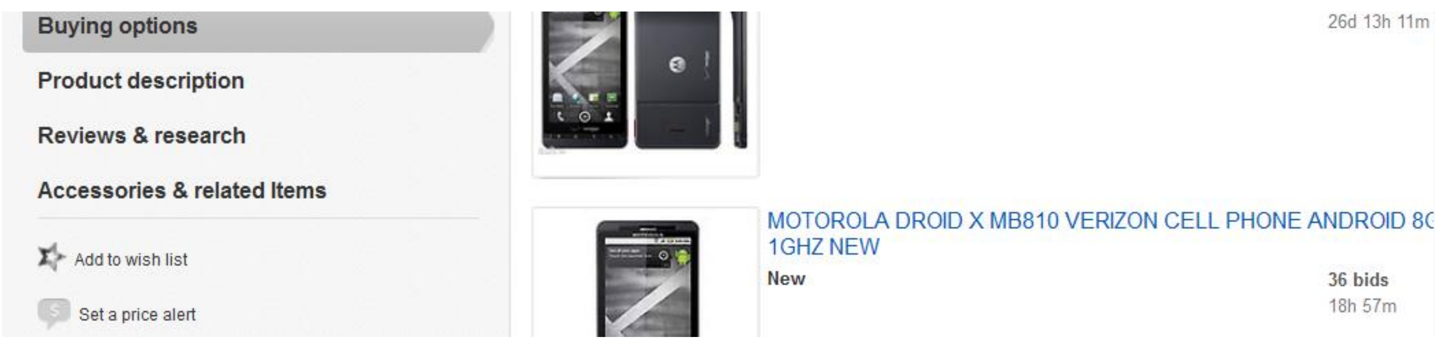

Auctions: (Almost) Full automation

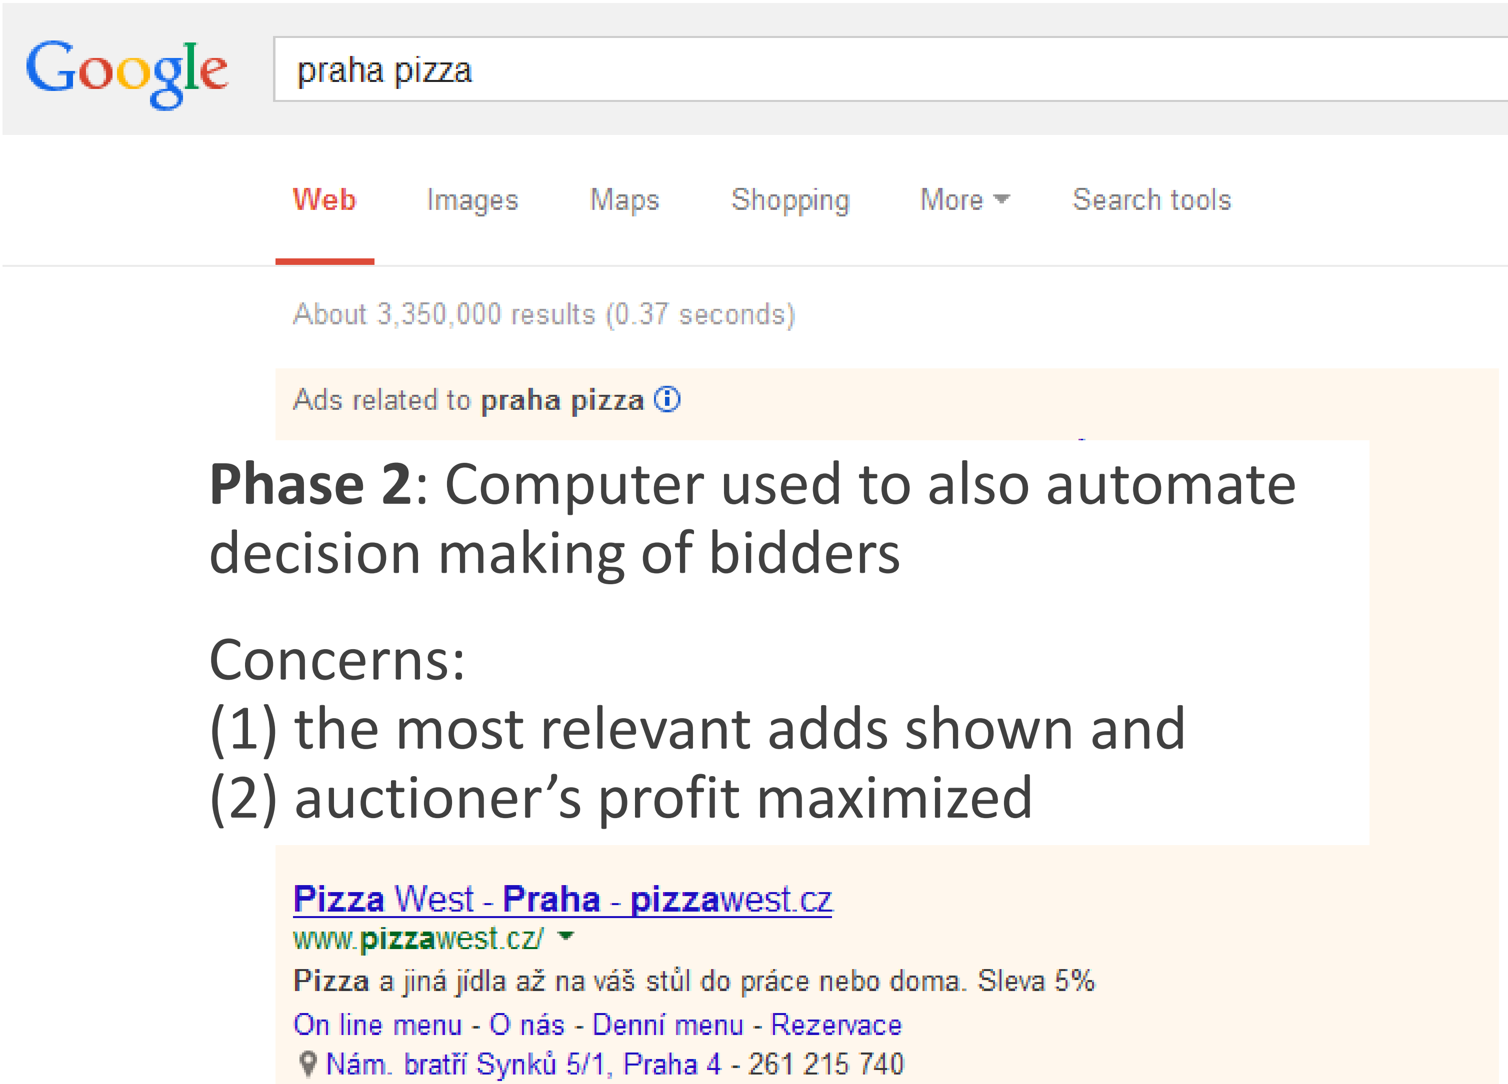

Phase 2: Computer used to also automate decision making of bidders

Concerns:
(1) the most relevant adds shown and
(2) auctioner's profit maximized

Lots of Applications

Industrial procurement

Transport and logistics

Energy markets

Cloud and grid computing

Internet auctions

Electromagnetic spectrum allocation

... and counting!

Multiagent Resource Allocation (MARA)

What is Multiagent Resource Allocation?

Multiagent Resource Allocation (MARA) is the process of distributing a number of items amongst a number of agents.

- **What** kind of items (resources) are being distributed?
- **How** are they being distributed?
- **Why** are they being distributed?

Classification of MARA

1. Resources
2. Agent preferences
3. Social welfare
4. Allocation mechanism

Type of Resources

Types of Resources

Central parameter in any resource allocation problem.

Different **types** of resources may require different resource allocation **techniques**.

Inherent **properties** of the **resource** vs. **characteristics** of the chosen **mechanism**.

Types of Resources (cntd.)

Continuous vs. **Discrete** resources

- physical property of the resource
- e.g. energy vs. fruit
- discrete resources are indivisible
- continuous resources may be treated either as being divisible or as being indivisible (→ discretisation)

Divisible vs. **Indivisible** resource

- resource may be treated as divisible or indivisible (e.g. 50l of juice)
- chosen feature of the allocation mechanism (possibly implied by resource physical properties)

Type of Resources (cntd.)

Sharable resource can be allocated to a number of different agents at the same time

- e.g. a photo taken by an earth observation satellite
- more often resource are **non-sharable**

Static resources do not change their properties during a negotiation process

- non-static: e.g. perishable goods (such as food)

Type of Resources (cntd.)

Single-unit vs. **Multi-unit** resources

- **single-unit**: exactly one copy of each type of good; all items are distinguishable (e.g. several houses).
- **multi-unit**: there may be several copies of the same type of good (e.g. 10 bottles of wine).

Distinction is only a matter of **representation**

- Every multi-unit problem can be translated into a single-unit problem by introducing new names (inefficient, but possible).
- Multi-unit problems allow a more **compact** representation but require a **richer** language.

Resources vs. Tasks

Tasks may be considered resources with **negative** utility (cost).

Task allocation may be regarded a multiagent resource allocation problem.

- However, tasks are often coupled with **constraints** regarding their **coherent combination** (timing and ordering).

Preference Representation

Preference Representation

The second important parameter in the specification of a MARA problem.

Agents may have **preferences** over

- the bundle of resources they receive
- the bundles of resources received by others (**externalities**)

What are suitable **languages** for representing agent **preferences**?

Preference Representation Languages

Expressive power

Succinctness

Complexity

Cognitive relevance

Elicitation

Cardinal vs. Ordinal Preferences

A **preference structure** represents an agent's preferences over a set of alternatives $\mathcal{X}$ (i.e. allocations).

- **Cardinal preference** structure is a function $u\colon \mathcal{X} \mapsto Val$, where Val is usually a set of numerical values such as $\mathbb{N}$ or $\mathbb{R}$.
- **Ordinal preference** structure is a *binary relation* $\preccurlyeq$ over the set of alternatives, that is *reflexive* and *transitive* (and connected).

If the alternatives over which agents have to express preferences are *bundles of indivisible resources* from the set $\mathcal{R}$, then we have $\mathcal{X} = 2^{\mathcal{R}}$.

Some Observations

Intrapersonal comparison: ordinal and cardinal preferences allow for comparing the satisfaction of an agent for different alternatives.

Interpersonal comparison: ordinal preferences don't allow for interpersonal comparison ("Ann likes x more than Bob likes y").

Preference intensity: ordinal preferences cannot express preference intensity; cardinal preferences can (subject to Val being numerical).

Representability: a connected ordinal preference relation $\preccurlyeq$ is representable by a utility function u: $x \preccurlyeq y$ iff $u(x) \leq u(x)$.

Cognitive relevance: hard to make general statements, but at least ordinal preferences don't require reasoning with numerical utilities.

Explicit representation: the explicit representations of cardinal and ordinal preferences have space complexity $\mathcal{O}(|\mathcal{X}|)$ and $\mathcal{O}(|\mathcal{X}|^2)$, respectively.

Example

Hanging a frame (f) with a hammer (h) and a nail (n) . . .

Cardinal

B	$u(B)$
$\{\}$	0
$\{f\}$	10
$\{h\}$	5
$\{n\}$	0
$\{f,n\}$	10
$\{f,h\}$	15
$\{h,n\}$	8
$\{f,h,n\}$	20

Ordinal

$\succeq$	$\{\}$	$\{f\}$	$\{h\}$	$\{n\}$	$\{f,n\}$	$\{f,h\}$	$\{h,n\}$	$\{f,h,n\}$
$\{\}$	1	0	0	1	0	0	0	0
$\{f\}$	1	1	1	1	1	0	1	0
$\{h\}$	1	0	1	1	0	0	0	0
$\{n\}$	1	0	0	1	0	0	0	0
$\{f,n\}$	1	1	1	1	1	0	1	0
$\{f,h\}$	1	1	1	1	1	1	1	0
$\{h,n\}$	1	0	1	1	0	0	1	0
$\{f,h,n\}$	1	1	1	1	1	1	1	1

Social Welfare

A third parameter in the specification of a MARA problem concerns our goals: **what kind of allocation do we want to achieve?**

We use the term **social welfare** in a very broad sense to describe **metrics** for assessing the **quality** of an **allocation** of resources.

Efficiency and Fairness

Two key indicators of social welfare.

Aspects of **efficiency*** include:

- The chosen agreement should be such that there is no alternative agreement that would be better for some and not worse for any of the other agents (**Pareto optimality**).
- If preferences are quantitative, the sum of all payoffs should be as high as possible (**utilitarianism**).

Aspects of **fairness** include:

- No agent should prefer to take the bundle allocated to one of its peers rather than keeping their own (**envy-freeness**).
- The agent that is going to be worst off should be as well off as possible (**egalitarianism**).

*not in the computational sense

Notation

Set of **agents** $\mathcal{A} = \{1, \dots, n\}$

Agents have **preferences over allocations**:

- **ordinal**: $A \preccurlyeq_i A'$ means agent i likes A no less than A'
- **cardinal**: $u_i(A) = x \in \mathbb{R}$ means agent i assigns utility x to A

Utilitarian Social Welfare

Utilitarian Collective Utility

The **utilitarian** collective utility function sw_u is defined as the sum of individual utilities:

$$sw_u(A) = \sum_{i \in \mathcal{A}} u_i(A)$$

Maximizing utilitarian CUF improves **efficiency**.

The utilitarian CUF is **zero-independent**: adding a constant value to your utility function won't a affect social welfare judgements.

Egalitarian Social Welfare

Egalitarian Collective Utility

The **egalitarian** collective utility function sw_e is defined as the sum of individual utilities:

$$sw_e(A) = \min\{u_i(A) | i \in \mathcal{A}\}$$

Maximising this function amounts to improving the situation of the weakest members of society (→ **fairness**).

Allocation A' is strictly preferred over allocation A (by society) iff $sw_e(A) < sw_e(A')$ holds (so-called **maximin**-ordering).

Nash Product Social Welfare

Nash Collective Utility

The **Nash** collective utility function sw_e is defined as the sum of individual utilities:

$$sw_e(A) = \prod_{i \in \mathcal{A}} u_i(A)$$

This is a useful measure of social welfare as long as all utility functions can be assumed to be **positive**.

Nash CUF favours increases in overall utility, but also inequality-reducing redistributions ($2 \cdot 6 < 4 \cdot 4$).

The Nash CUF is **scale independent**: whether a particular agent measures their own utility in euros or dollars does not affect social welfare judgements.

Social Welfare Curve Illustration

Allocation Procedures

Allocation Procedures

Solution of a MARA problem need to determine at least:

- **Protocols**: What types of deals are possible? What messages do agents have to exchange to agree on one such deal?
- **Strategies**: What strategies may an agent use for a given protocol? How can we give incentives to agents to behave in a certain way?
- **Algorithms**: How do we solve the computational problems faced by agents when engaged in negotiation?

Centralised vs. Distributed Allocation

Centralised case	Distributed case
• A **single entity decides** on the final allocation, possibly after having elicited the preferences of the other agents. • Example: combinatorial auctions	• **Allocations emerge** as the result of a sequence of local negotiation steps. • Such local steps may or may not be subject to structural restrictions (say, bilateral deals).

Which approach is appropriate under what circumstances?

Centralised vs. Distributed Comparison

Centralised	Distributed
• The **communication** protocols required are relatively **simple**. • Many **results** from **economics** and **game theory** , in particular on mechanism design, can be exploited. • **Powerful algorithms** for winner determination in combinatorial auctions. • Possible **trust** issues. • Difficult to deal with **unbounded problems**.	• Avoids **trust** issues. • Inherently **scalable**. • Can take an **initial allocation** into account. • More natural to model **step-wise improvements** over the status quo. • Can deal with **unbounded domains**. • More **complex** protocols significantly more **difficult** to analyse (convergence etc.)

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

Auctions Rules

Auction mechanism is **specified** by **auction rules**

- *rules of the game*

Bidding rules: How offers are made:

- by whom
- when
- what their content is

Clearing rules: Who gets which goods (**allocation**) and what money changes hands (**payment**).

Information rules: What information about the state of the negotiation is revealed to whom and when.

Valuation Models

Agent's payoff from participating in an auction

- if winner: payoff = item's valuation – price paid for the item
- if not winner: payoff = zero

Common value: the good has the same value to all agents

- a 100 dollar note

Private value: an agent A's valuation of the good is independent from other agent's valuation of the good

- a painting, John Lennon's last dollar bill

Correlated value: valuations of the good are related

- i.e. the more other agents are prepared to pay, the more agent A prepared to pay.
- i.e. purchase of items for later resale

Single Good Auctions

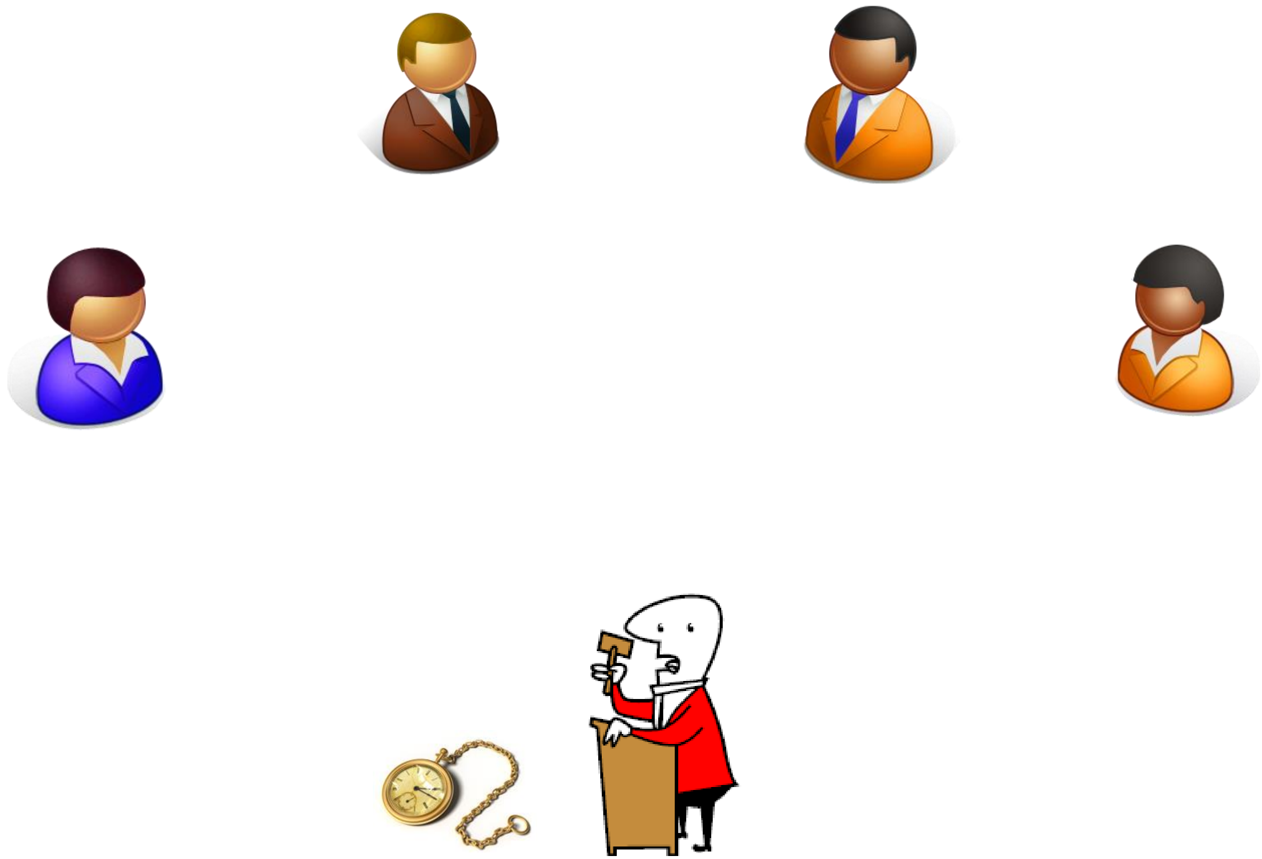

Multi-Unit Auctions




single-unit vs. multiple-unit demand

Indistinguishable items

Multi-Item Auctions

Reverse Auctions

FRONT SUSPENSION, FRONT WHEEL BEARING ACQUISITION

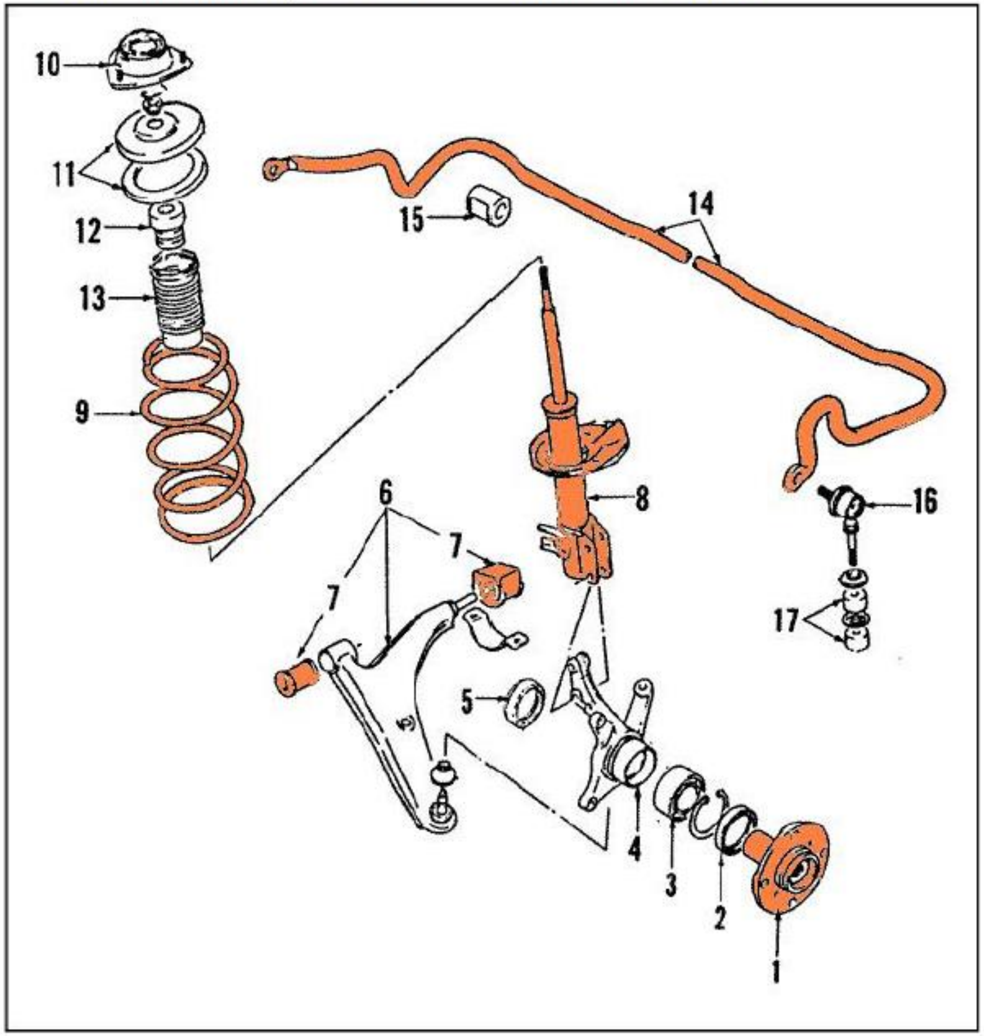

- Goal: Buy parts to produce a front suspension.
- The buyer issues a request for bids to his providers.

PART #	DESCRIPTION
1	FRONT HUB
7	LOWER CONTROL ARM BUSHINGS
8	STRUT
9	COIL SPRING
14	STABILIZER BAR

Multi-Attribute Auctions

Negotiation over other **attributes in addition to price**

- e.g. color, weight, or delivery time

For instance: Provider John Doe offers to deliver a stainless-steel stabilizer bar that weighs 500 g at the cost of 200 EUR by July 18th 2011.

Promise **higher market efficiency** through a more **effective information exchange** of buyer's preferences and supplier's offerings.

Least understood type of auctions.

Auction Mechanism Taxonomy

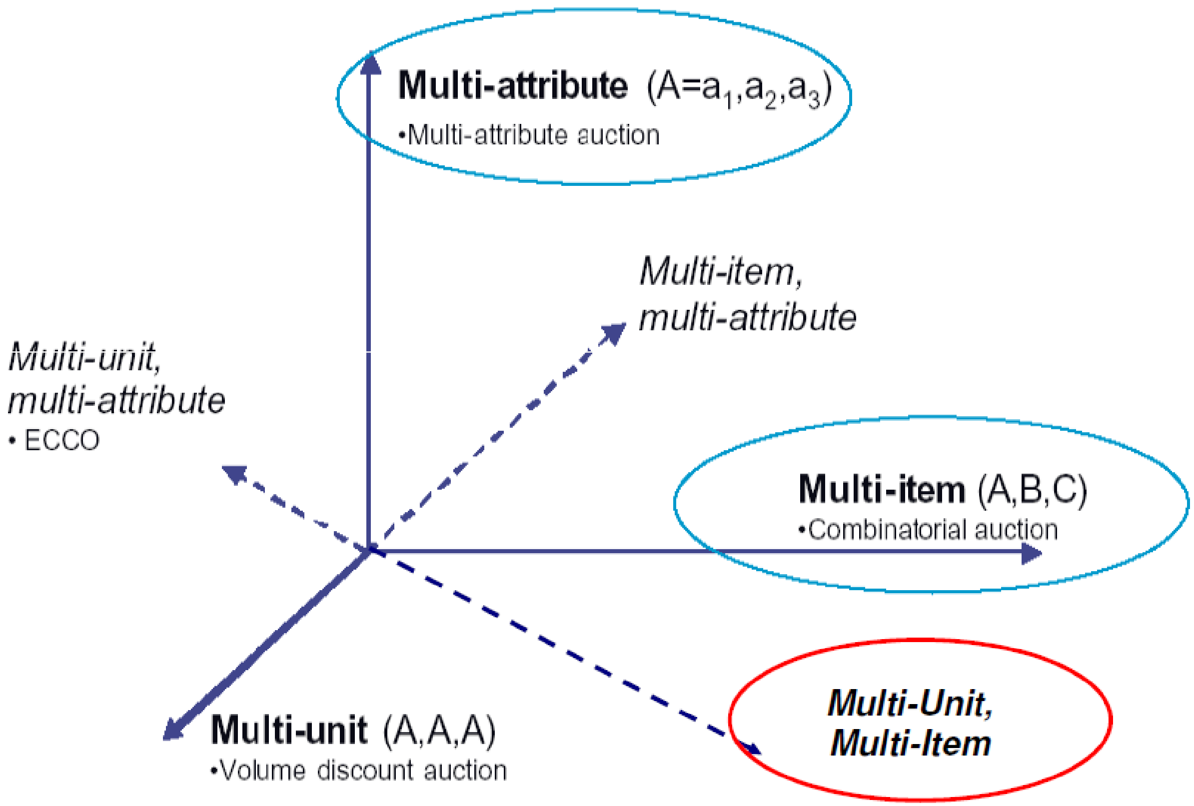

Other: First-price vs. *k*-th price, open cry vs. sealed bid, single. vs. double-sided, sell-side vs. buy-side

Single-Item Auctions

Basic Auction Mechanisms

English

Japanese

Dutch

First-Price

Second-Price

English Auction

Auctioneer starts the bidding at some **reservation price**

Bidders then shout out **ascending** prices

- minimum increments

Once bidders stop shouting, the *high bidder* gets the good at that price

Japanese Auctions

Same as an English auction except that the auctioneer calls out the prices

All bidders start out **standing**

When the price reaches a level that a bidder is not willing to pay, that bidder **sits down**

Once a bidder sits down, they **can't get back up** the **last** person **standing** gets the good

Dutch Auction

The auctioneer starts a clock at some high value; it descends

At some point, a bidder shouts “mine!" and gets the good at the price shown on the clock

Good when items need to be sold quickly (similar to Japanese)

No information is given away during auction

First-, Second-Price Sealed Bid Auctions

First-price sealed bid auction

- bidders write down bids on pieces of paper
- auctioneer **awards** the good to the bidder with the **highest bid**
- that bidder pays the amount of **his bid**

Second-price sealed bid auction (Vickerey auction)

- bidders write down bids on pieces of paper
- auctioneer **awards** the good to the bidder with the **highest bid**
- that bidder pays the amount bid by **the second-highest** bidder

Intuitive Comparison

	English	Dutch	Japanese	1st-Price	2nd-Price
Duration	#bidders, increment	starting price, clock speed	#bidders, increment	fixed	fixed
Info Revealed	2nd-highest val; bounds on others	winner's bid	all val's but winner's	none	none
Jump bids	yes	n/a	no	n/a	n/a
Price Discovery	yes	no	yes	no	no

Analysing Auctions

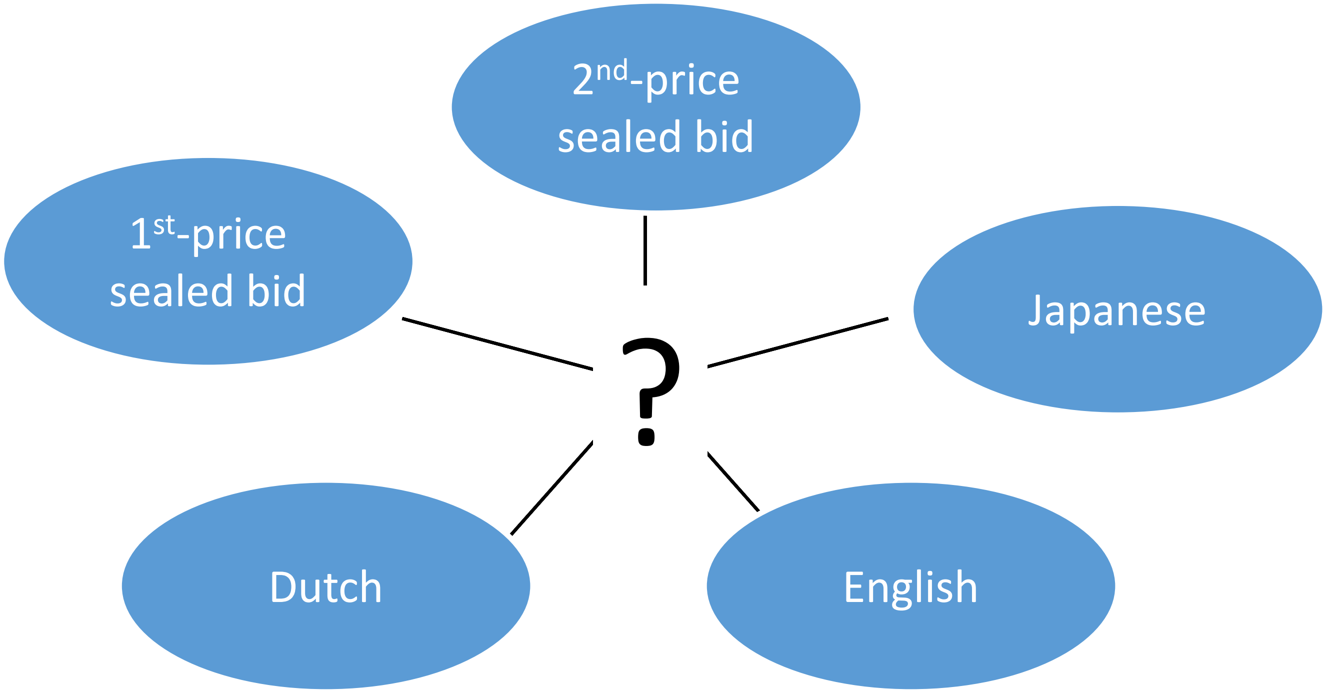

Are there fundamental similarities / differences between mechanisms described?

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 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

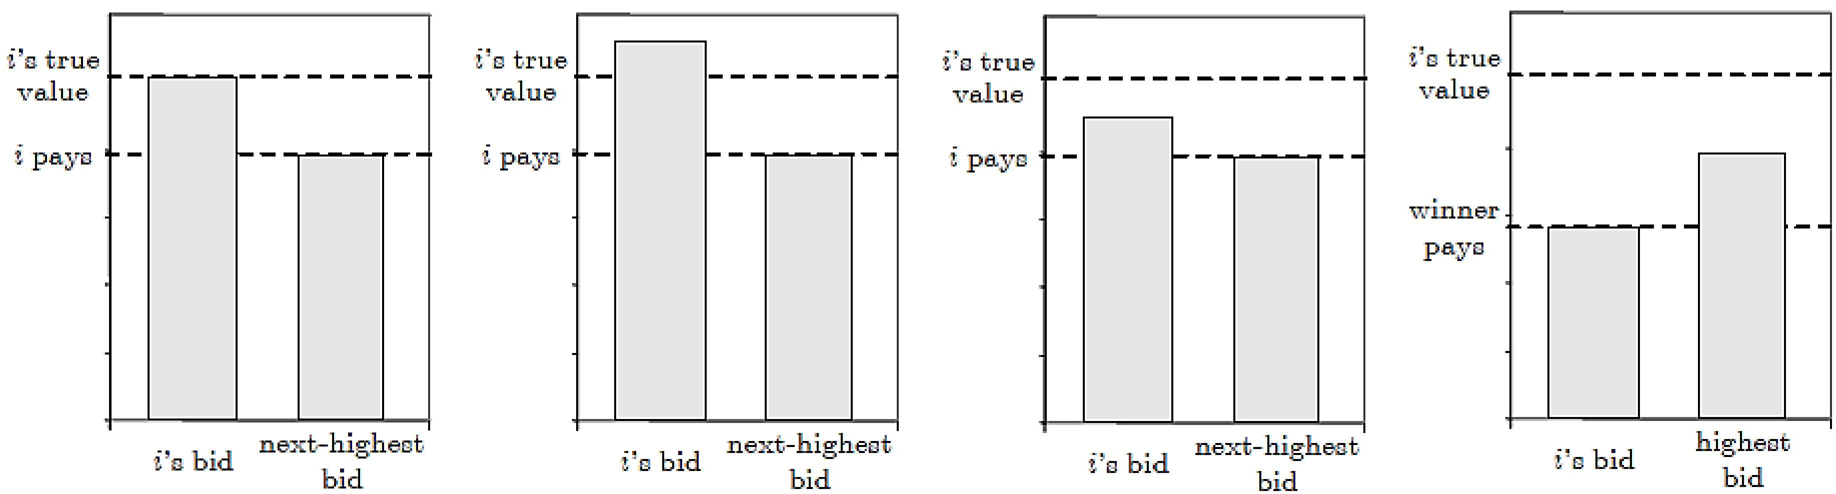

Bidding honestly, i is the winner

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

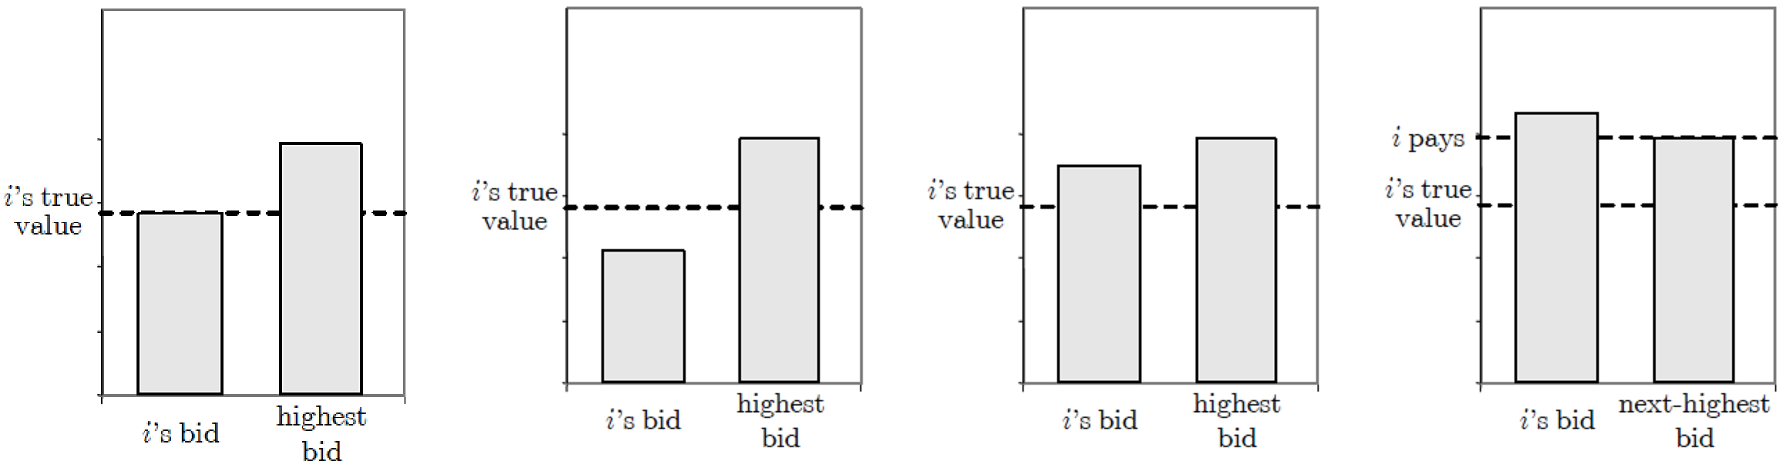

Bidding honestly, i is not the winner

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 $\Rightarrow$ **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** $(\frac{n-1}{n}v_1,\dots,\frac{n-1}{n}v_n)$.

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.

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