Multiagent Systems

CS 486/686: Introduction to Artificial Intelligence
Fall 2013

Introduction

- So far almost everything we have looked at has been in a single-agent setting
  - Today - Multiagent Decision Making!
- For participants to act optimally, they must account for how others are going to act
- We want to
  - Understand the ways in which agents interact and behave
  - Design systems so that agents behave the way we would like them to

Hint for the final exam: MAS is my main research area. I like MAS problems. I even enjoy marking MAS questions. Two of the TAs for this course do MAS research. They also like marking MAS questions. There ***will*** be an MAS question on the final exam.

Introduction

- Multiagent systems can be
  - cooperative or self-interested
- Self-interested multiagent systems can be studied from different viewpoints
  - non-strategic and strategic
- We will look at strategic self-interested systems

Self-Interest

- Self-interested does not mean
  - Agents want to harm others
  - Agents only care about things that benefit themselves
- Self-interested means
  - Agents have their own description of states of the world
  - Agents take actions based on these descriptions

Tools for Studying MAS

- Game Theory
  - Describes how self-interested agents should behave
- Mechanism Design
  - Describes how we should design systems to encourage certain behaviours from self-interested agents

What is Game Theory?

- The study of games!
  - Bluffing in poker
  - What move to make in chess
  - How to play Rock-Paper-Scissors

Also auction design, strategic deterrence, election laws, coaching decisions, routing protocols,...

What is Game Theory?

- Game theory is a formal way to analyze interactions among a group of rational agents that behave strategically

What is Game Theory?

- Game theory is a formal way to analyze **interactions** among a **group** of **rational** agents that behave **strategically**
  - **Group:** Must have more than 1 decision maker
    - Otherwise, you have a decision problem, not a game

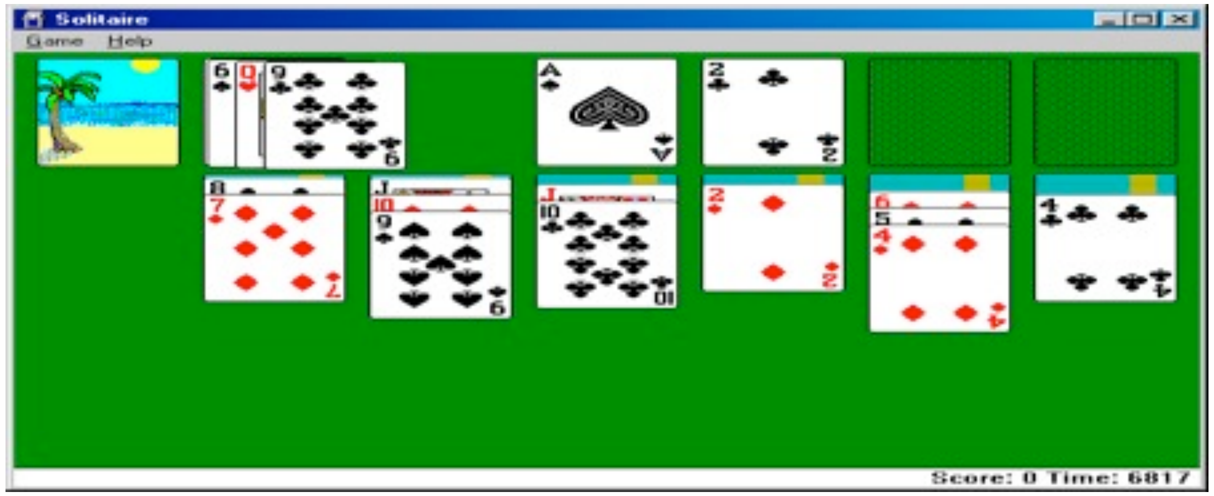

Solitaire is not a game!

What is Game Theory?

- Game theory is a formal way to analyze **interactions** among a **group** of **rational** agents that behave **strategically**
  - **Interaction:** What one agent does directly affects at least one other
  - **Strategic**: Agents take into account that their actions influence the game
  - **Rational**: Agents chose their best actions

Example

- Decision Problem
  - Everyone pays their own bill
- Game
  - Before the meal, everyone decides to split the bill evenly

Strategic Game (Matrix Game, Normal Form Game)

- Set of agents I={1,2,.,.,N}
- Set of actions $A_i=\{a_i^1,\ldots,a_i^m\}$
- Outcome of a game is defined by a profile $a=(a_1,\ldots,a_n)$
- Agents have preferences over outcomes
  - Utility functions u_i:A->**R**

Examples

		Agent 2 One	Agent 2 Two
Agent 1	One	2,-2	-3,3
	Two	-3,3	4,-4

Zero-sum game.

$\Sigma_{i=1}^{n} u_i(o)=0$

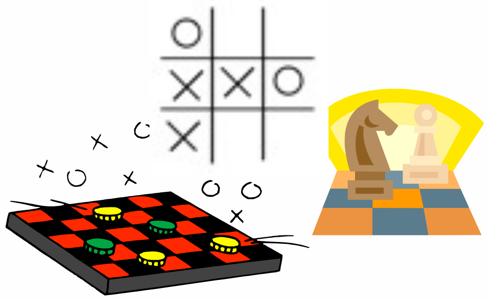

I={1,2}
Ai={One,Two}
An outcome is (One, Two)
U_1((One,Two))=-3 and U_2((One,Two))=3

Examples

BoS

	B	S
B	2,1	0,0
S	0,0	1,2


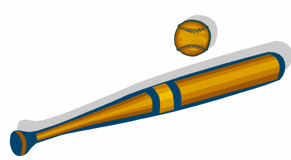

Coordination Game

Chicken

	T	C
T	-1,-1	10,0
C	0,10	5,5

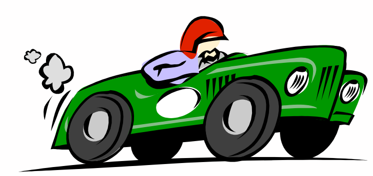


Anti-Coordination Game

Example: Prisoners' Dilemma

	Confess	Don't Confess
Confess	-5,-5	0,-10
Don't Confess	-10,0	-1,-1

Playing a Game

- Recall, agents are rational
  - Let p_i be agent i's belief about what its opponents will do
  - Agent i will try to maximize its expected utility given its belief over the others
  - $a_i = \text{argmax} \Sigma_{a_{-i}} u_i(a_i, a_{-i}) p_i(a_{-i})$

Notation Break: $a_{-i} = (a_1, \ldots, a_{i-1}, a_{i+1}, \ldots, a_n)$

Dominated Strategies

- A strategy a_i is strictly dominated if

$$u_a(a_i', a_{-i}) > u_i(a_i, a_{-i}) \forall a_i' \neq a_i$$

- A rational agent will never play a dominated strategy!

Example

	Confess	Don't Confess
Confess	-5,-5	0,-10
Don't Confess	-10,0	-1,-1

Example

	Confess	Don't Confess
Confess	-5,-5	0,-10
Don't Confess	~~-10,0~~	~~-1,-1~~

	Confess	Don't Confess
Confess	-5,-5	0,-10

Example

	Confess	Don’t Confess
Confess	-5,-5	0,-10
Don’t Confess	~~-10,0~~	~~-1,-1~~

	Confess	Don’t Confess
Confess	-5,-5	~~0,-10~~

	Confess
Confess	-5,-5

Equilibrium Outcome

Strict Dominance Does Not Capture the Whole Picture

	A	B	C
A	0,4	4,0	5,3
B	4,0	0,4	5,3
C	3,5	3,5	6,6

What strict domination eliminations can we do?

What would you predict the players of this game would do?

Nash Equilibrium

- An agent's best-response depends on the actions of other agents
- An action profile a* is a **Nash equilibrium** if no agent has incentive to deviate given that others do not deviate

$$\forall i u_i(a_i^*, a_{-i}^*) \geq u_i(a_i', a_{-i}^*) \forall a_i'$$

Nash Equilibrium

- Equivalently, a* is a N.E. iff

$$\forall i a_i^* = \arg\max_{a_i} u_i(a_i, a_{-i}^*)$$

	A	B	C
A	0,4	4,0	5,3
B	4,0	0,4	5,3
C	3,5	3,5	6,6

(C,C) is a N.E. because

$$u_1(C,C) = \max \begin{bmatrix} u_1(A,C) \\ u_1(B,C) \\ u_1(C,C) \end{bmatrix}$$

AND

$$u_2(C,C) = \max \begin{bmatrix} u_2(C,A) \\ u_2(C,B) \\ u_2(C,C) \end{bmatrix}$$

Nash Equilibrium

- If (a_1^*, a_2^*) is a N.E. then player 1 won't want to change its action given player 2 is playing a_2^*
- If (a_1^*, a_2^*) is a N.E. then player 2 won't want to change its action given player 1 is playing a_1^*

-5,-5	0,-10
-10,0	-1,-1

	A	B	C
A	0,4	4,0	5,3
B	4,0	0,4	5,3
C	3,5	3,5	6,6

Another Example

	B	S
B	2,1	0,0
S	0,0	1,2

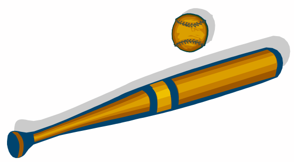

2 Nash Equilibria

Coordination Game

Yet Another Example

		Agent 2 One	Agent 2 Two
Agent 1	One	2,-2	-3,3
	Two	-3,3	4,-4

(Mixed) Nash Equilibria

- **(Mixed) Strategy**: s_i is a probability distribution of Ai
- **Strategy profile**: $s=(s_1,...,s_n)$
- **Expected utility**: $u_i(s)=\Sigma_a \Pi_j s(a_j) u_i(a)$
- **Nash equilibrium**: s* is a (mixed) Nash equilibrium if

$$u_i(s_i^*, s_{-i}^*) \geq u_i(s_i', s_{-i}^*) \forall s_i'$$

Yet Another Example

		q One	Two
p	One	2,-2	-3,3
	Two	-3,3	4,-4

How do we determine p and q?

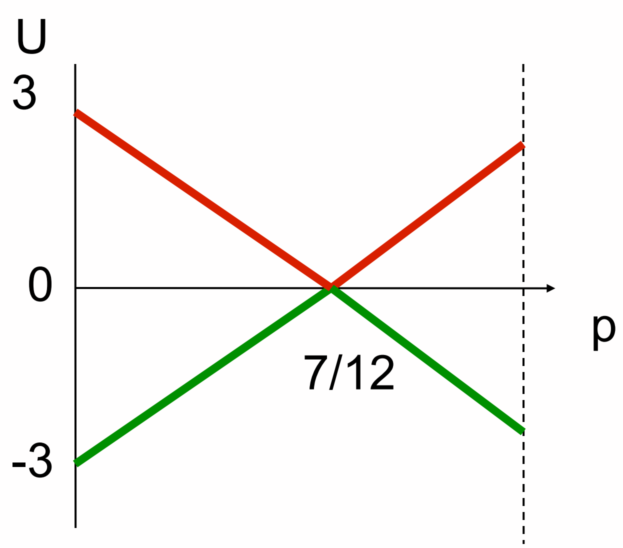

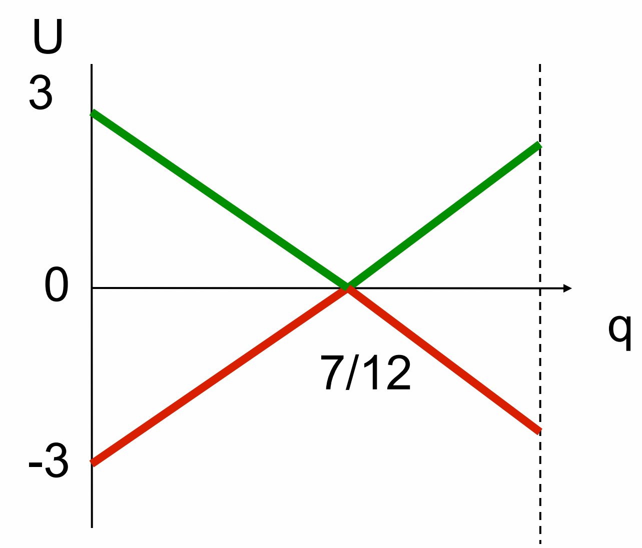

Yet Another Example

		q	One	Two
p	One		2,-2	-3,3
	Two		-3,3	4,-4

How do we determine p and q?

Exercise

	B	S
B	2,1	0,0
S	0,0	1,2

This game has 3 Nash Equilibrium (2 pure strategy NE and 1 mixed strategy NE). Find them.

Mixed Nash Equilibrium

- Theorem (Nash 1950): Every game in which the action sets are finite, has a mixed strategy equilibrium.

John Nash
Nobel Prize in Economics (1994)

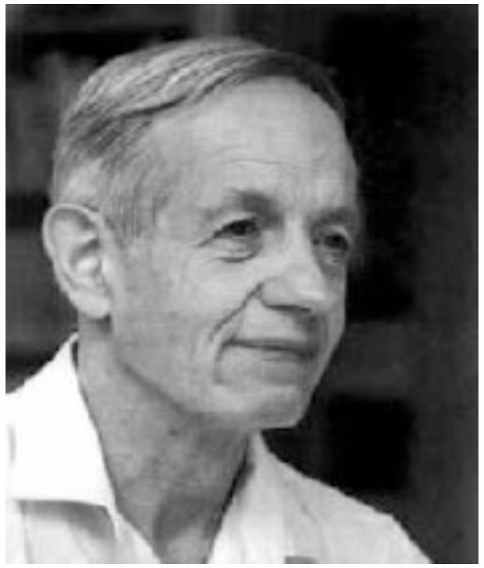

Finding NE

- Existence proof is non-constructive
- Finding equilibria?
  - 2 player zero-sum games can be represented as a linear program (Polynomial)
  - For arbitrary games, the problem is in PPAD
  - Finding equilibria with certain properties is often NP-hard

Mechanism Design

- Game Theory asks
  - Given a game, what should rational agents do?
- Mechanism Design asks
  - Given rational agents, what sort of games should we design?
  - Can we guarantee that agents will reach an outcome with properties **we** want

Scenario

- Network routing problem to allocate resources to minimize the total cost of delay over all agents

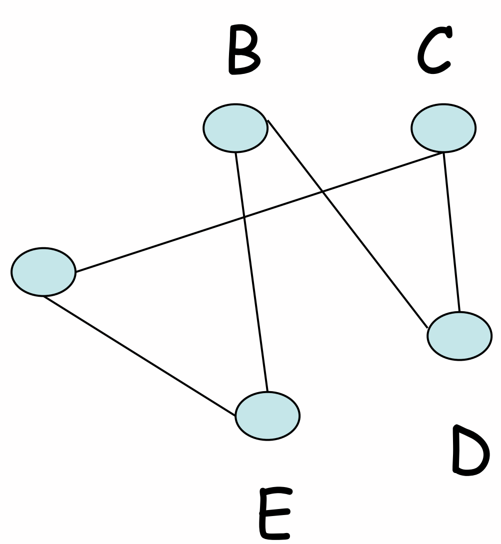

My unit cost of delay for sending messages from A to D is $1

My unit cost of delay for sending messages between E and D is $5

From Our Perspective

- As the system designer, we want to reach some desirable social outcome
- Social choice function $f{:}T_1 x \ldots x T_n \text{->} O$ maps every possible **type** profile to some outcome

A Potential Problem

- Agents' types are not public, and agents are acting in their own self-interest

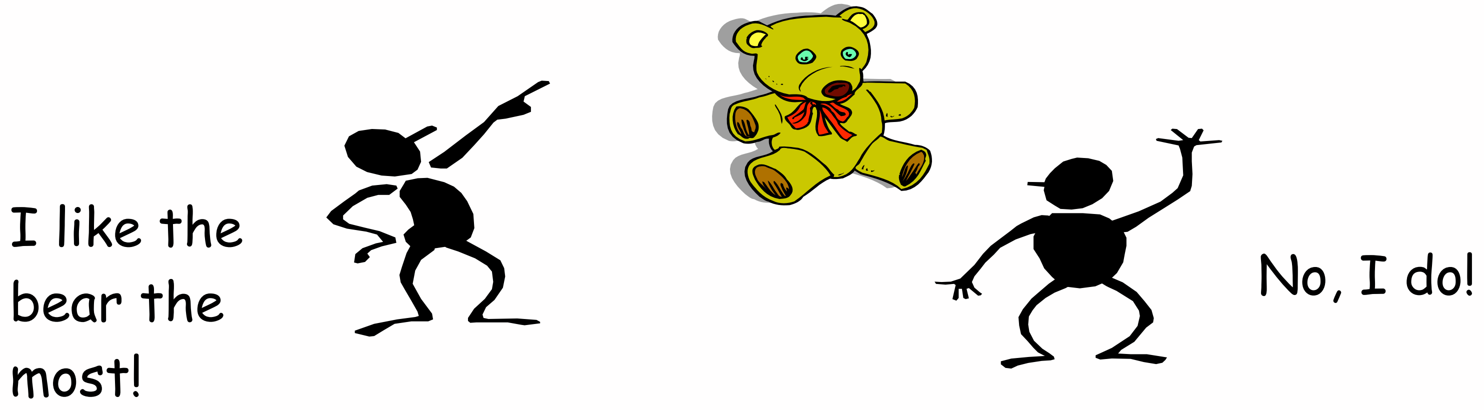

Mechanism Design Problem

- The mechanism design problem:
  - Design “rules of the game” so that the solution of a social choice function is implemented, **despite agents’ self-interest**
  - Mechanism

$$M=(S_1,\ldots,S_n,g)$$

Strategy spaces   Outcome function

  - M **implements** SCF f if for equilibrium $s^*=(s_1^*(t_1),\ldots,s_n^*(t_n))$, $f(t)=g(s^*(t))$, for all $t=(t_1,\ldots,t_n)$

Example: Allocation Problem

- Social choice function: Maximize social welfare
- Agents' utility functions: $u_i = v_i(o) - p_i$
  - Type of agent i is v_i
- Mechanism: Vickrey Auction
  - S_i=set of legal bids
    - Any non-negative real number
  - Outcome function g
    - Give the item to the agent with the highest bid
    - The winner pays an amount equal to the second highest bid, everyone else pays nothing

Vickrey Auction

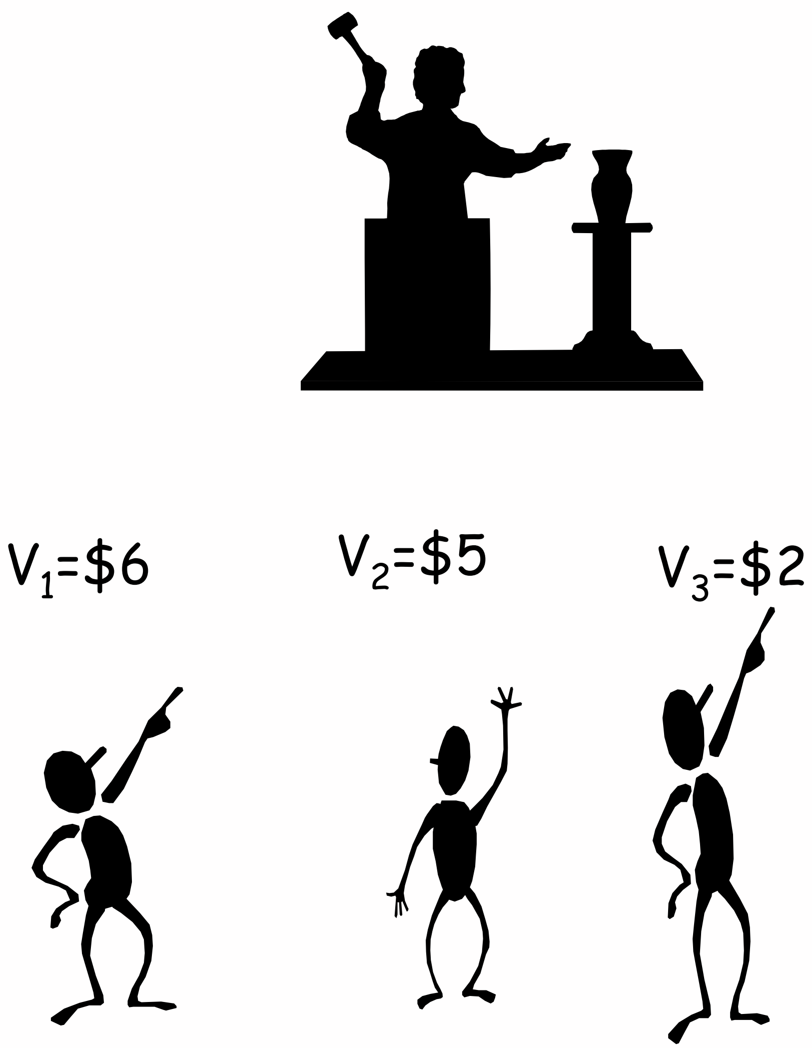

If agents bid truthfully then
Agent 1 wins
Pays \$5

$U_1 = \$6 - \$5 = \$1$

$U_2 = U_3 = 0$

Vickrey Auction

- Case 1: Bidding truthfully and you are the highest bidder

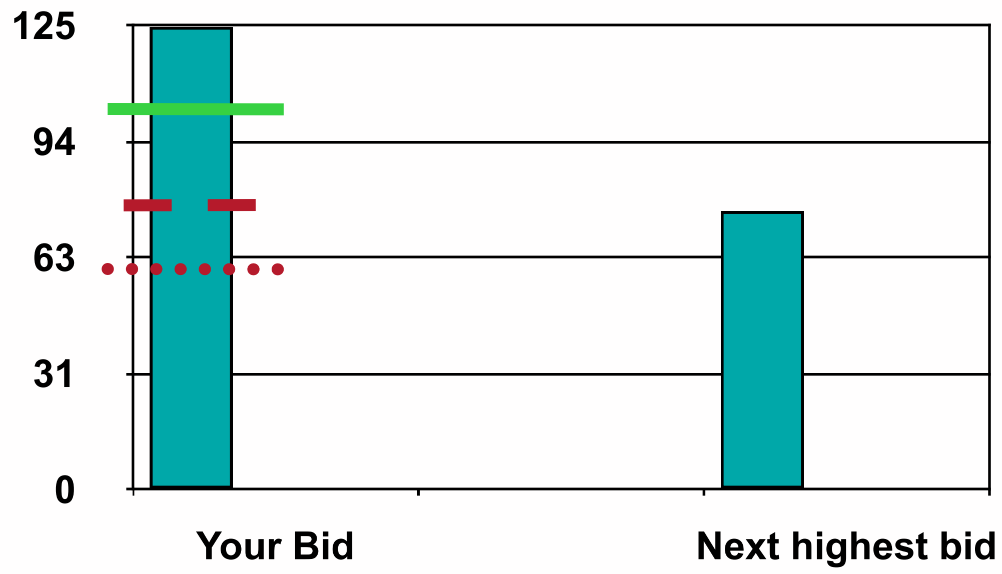

Bid more:
    No difference
    Still pay the same

Bid less:
    No difference
    Lose the auction

Vickrey Auction

- Case 2: Bidding truthfully and you are not the highest bidder

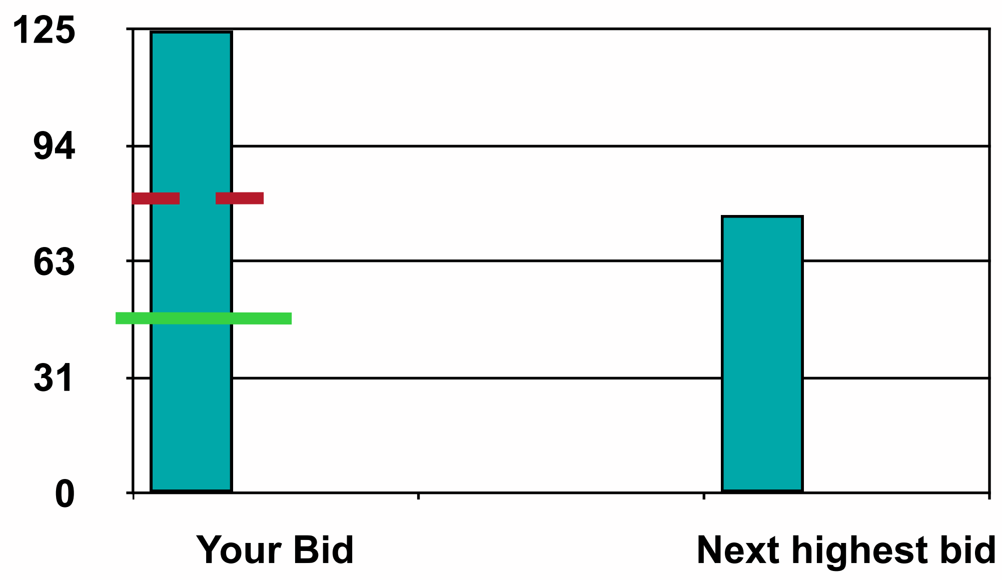

Bid less:
No difference

Bid more:
No difference
Win the auction and pay too much

Other Application: Sponsored Search

Slot 1
Slot 2
Slot 3
Slot 4
Slot 5

<Keyword>

Bid Bid

1. Advertisers are ranked and assigned slots based on the ranking.
2. If an ad is clicked on, only then does the advertiser pay.

Ranking

- Rank-by-relevance
  - Assign slots of order of (quality score)*(bid)

Bidder	Bid	Quality Score
A	1.50	0.5
B	1.00	0.9
C	0.75	1.5

→

Ranking
C (1.25)
B (0.9)
A (0.75)

Pricing

- An advertiser only pays when its ad is clicked on
- How much does it pay?
  - The lowest price it could have bid and still been in the same position

Example

Bidder	Bid	Quality Score
A	1.50	0.5
B	1.00	0.9
C	0.75	1.5

→

Ranking
C (1.25)
B (0.9)
A (0.75)

C will pay p=0.9/1.5=0.6
B will pay p=0.75/0.9 = 0.83

How much will A pay?

Sponsored Search

- How would you design a bidding agent for sponsored search?
- Different from the Vickrey auction
  - There is no single best strategy
  - It depends on the strategies of others

Summary: What you Should Know

- What a game is
- What a (Nash) Equilibrium is
- What a mechanism is
- Some uses of mechanisms