Multi Agent Systems, Game-theory/Auction

Makoto Yokoo
Kyushu University
lang.is.kyushu-u.ac.jp/~yokoo/

Outline

- Example
- Why game-theory/economics?
- Game-theory
- Auctions

Example: Standard (First-price) Sealed-bid Auction

Protocol: Each bidder (agent) privately bids the price it is willing to pay, and the highest bidder wins; pays the value of its bid.

Demerit: Spying other agents' bids is profitable.

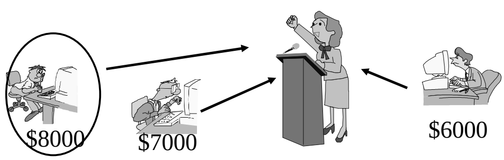

Second-price Sealed-bid Auction (Vickrey Auction)

Protocol : The highest bidder wins, but pays the value of the second highest bid.

Characteristic:

–Bidding its true evaluation value (the maximal value where it does not want to pay any more) is the optimal strategy (incentive compatibility).

–Spying other agents' bids is meaningless.

[redacted]

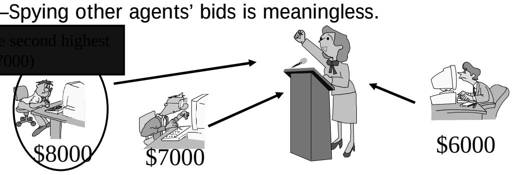

Spying is Useless Because...

Assume its own evaluation value is $8000, and found that others' highest bid is:

(I) less than $8000: its payment is the same to the case when it submits the true evaluation value (without spying).

(II) more than $8000: the agent cannot obtain positive utility anyway.

Its profit cannot be more than the case when it truthfully submits the true evaluation value.

[redacted]

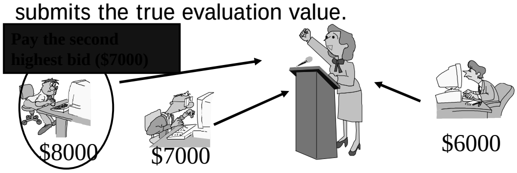

Quiz: The larger also serves for the smaller?

- If the highest bidder wins, and pays the value equals to the third highest bid, is honesty still the best policy?

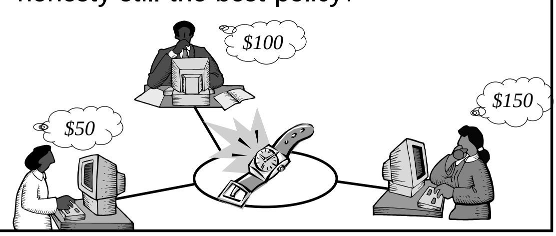

Answer

- Honesty is no longer the best policy.
  - For the first and second bidder, if he wins, he is going to pay only $50.
  - Can increase his bid to infinity to beat the opponent.
- If two identical units are sold, then honesty is the best policy.

Theory of Auctions

- For seller:
  - The behaviors of bidders change according to the auction protocol.
  - can give a protocol that can achieve socially desirable outcome or robustness against cheating.

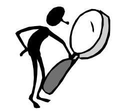

- For bidders:
  - The best strategy changes according to the auction protocol.
  - can give a method for finding the best strategy.

important technology for Electronic Commerce

Why Game-theory/Economics?

- Game-theory/Economics provide analytical tools for the design/analysis of multi-agent systems (MAS).
  - Game-theory/economics are specially useful when the MAS
    - are not centrally designed.
    - do not have a notion of global utility.
    - will not necessarily act “benevolently”.

Why Game-theory/Economics (contd.)?

- Underlying assumptions:
  - Each player (agent) has consistent preference/utility.
  - Each player is rational, i.e., tries to maximize his/her utility.
  - Each player chooses a strategy to play.

These assumptions only approximate human behavior, but are fully applicable to MAS.

Why Game-theory/Economics (contd.)?

John von Neumann

- Increasing intersectional topics:
  - Internet auctions, electronic commerce
  - resource allocation for automated agents

From the theory for describing human behavior to the theory for designing systems!

Outline

- Example
- Why game-theory/economics?
- Game-theory
  - Introduction
  - Games with Complete Information
  - Games with Incomplete Information
- Auctions

Selfish Agents

- Preference
  - A certain (social) state is preferred than another, e.g., having $100 is better than having $50.
  - Preference may vary among agents, e.g., some prefer having apples while others prefer oranges.
- Utility
  - A value of a (social) state given by each agent.
  - Preferred state derives higher utility.
- "Selfish" agents: Each agent behaves to maximize its own utility.

Pareto Efficiency (1)

Definition:

A state is said Pareto efficient, iff there exists no state that is

- better for one agent, and
- no worse for all the rest

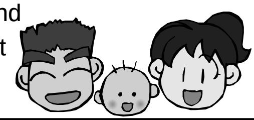

✕	Movie	2	2	2
	Shopping	2	2	5
	Zoo	2	3	1
✕	Home	1	1	1

Pareto Efficiency (2)

- Pareto efficiency can be considered a minimal requirement for social optimality.
  - If a state is not Pareto efficient, there exists another state that all members think it is better (or the same).
- However, a Pareto efficient state is *not always unique*.

  Example: Dividing $100 between two people. Both "$50 to each" and "$100 to one, $0 to the other" are Pareto efficient.

Outline

- Example
- Why game-theory/economics?
- Game-theory
  - Introduction
  - Games with Complete Information
  - Games with Incomplete Information
- Auctions

Outline: Games with Complete Information

- Definition
- Dominant Strategy Equilibrium
- Iterated Dominance Equilibrium
- Min-Max Strategy
- Mixed Strategy
- Nash Equilibrium

Game with Complete Information

- Each agent knows the possible actions/utilities of both of itself and opponents with certainty.
- Can be described as a matrix.

I \ II	F	S
F	4 \ 4	8 \ 2
S	2 \ 8	1 \ 1

Competing Newspaper

- There are two competing newspapers.
- Each has a choice whether to cover Finance news as a top news or cover Sports.
- 80% of people prefer Finance, while 20% people prefer Sports.

I \ II	F	S
F	4, 4	8, 2
S	2, 8	1, 1

Rational Agent/Player

- Each Agent/Player is rational:
  - Tries to maximize its own utility.
  - Does not care about other people's utility.
  - There is no felling such as pity or unfair.
  - Sometimes called selfish, but:
    - if the agent has feelings like sympathy, moral, or whatever, we assume they are already represented in the matrix.

I \ II	F	S
F	4, 4	8, 2
S	2, 8	1, 1

Assumption

- Each agent knows the possible actions/utilities of both of itself and opponents (i.e., the payoff matrix) with certainty
- Of course, the agent does not know which action his opponent will choose.

I \ II	F	S
F	4, 4	8, 2
S	2, 8	1, 1

Assumption (cont'd)

- Each agent chooses its action simultaneously, without negotiation.
  - They cannot negotiate, say, I'll chose F, so could you please chose S, then I'll pay you $1000, etc.
  - An agent cannot choose its action after observing the opponent's action.

I \ II	F	S
F	4, 4	8, 2
S	2, 8	1, 1

Competing Newspaper

- Which action should you choose if you are player I?
- Your best scenario is you choose F, and the opponent choose S, but you cannot control the opponent's action.
- Your opponent is not a fool (actually, very wise) and tries to maximize its own utility.

I \ II	F	S
F	4, 4	8, 2
S	2, 8	1, 1

Dominant Strategy

Strategy: the way for choosing an action

Dominant Strategy: the strategy that gives you higher (or equal) utility than any other strategy, no matter the action the opponent chooses.

- Clearly, a rational player will choose a dominant strategy if exists.
- We don't need to care whether your opponent is rational (even for a very weird player, or you have no idea of your opponent's utility, it is just fine).

I \ II	F	S
F	4, 4	8, 2
S	2, 8	1, 1

Dominant Strategy Equilibrium

If each player has a dominant strategy, the combination is called a dominant strategy equilibrium.

- If players are rational, and there exists a dominant strategy equilibrium, we can assume that the result will be that dominant strategy equilibrium.

I \ II	F	S
F	4, 4	8, 2
S	2, 8	1, 1

Dominant Strategy

Dominant Strategy does not necessarily exists.

- Paper-Rock-Scissors: no dominant strategy
- If only paper and rock are allowed, paper is the dominant strategy

- Most games (which a human enjoys to play) does not have a dominant strategy.
- In mechanism design (e.g., auctions), the goal is to design the rule so that a dominant strategy equilibrium exists.

Battle of the Bismarck Sea

- in the South Pacific, 1943.
- Rear Admiral Kimura needs to transport Japanese troops across the Bismarck Sea to New Guinea.
- He can choose either a short north route or a long south route.
- Admiral Kenny must decide where to send his bomber planes.
- If he chose wrong route, the time for bombing is reduced.

Ke \ Ki	N	S
N	2, -2	2, -2
S	1, -1	3, -3

Battle of the Bismarck Sea

- This is a zero-sum game.
- No dominant strategy for Kenny.
- Which route should Kenny choose?

K \ I	N	S
N	2, -2	2, -2
S	1, -1	3, -3

Iterated Dominance Equilibrium

- For Kimura, choosing North is a (weakly) dominant strategy.
- Assuming Kimura is rational, he would choose North.
- Then, Kenny should choose North.
- By iteratively removing dominated strategy, we can obtain an iterated dominance equilibrium.
- Your opponent needs to be rational.

Ke \ Ki	N	S
N	2, -2	2, -2
S	1, -1	3, -3

Boxed Pigs

- Example used in a psychological test for animals.
- A big pig and a small pig are in a (large) box.
- If a pig push a button, then some foods appear in a slightly far away place.
- If a small pig pushes the button, the big pig gets most of the foods.
- If a big pig pushes the button, then a small pig can get about a half.
- What kind of actions these pigs learn?

Boxed Pigs

- The payoff matrix is as follows.
- The bigger one does not always win.
- Burn one's bridge can be good!

Big Pig \ Small Pig	Push	Wait
Push	5, 1	4, 4
Wait	9, -1	0, 0

Outline: Games with Complete Information

- Definition
- Dominant Strategy Equilibrium
- Iterated Dominance Equilibrium
- Min-Max Strategy
- Mixed Strategy
- Nash Equilibrium

Min-Max Strategy

- What should we do if there is no dominant strategy/iterated dominance equilibrium?
- One possibility: let's avoid the worst-case!
  - Min-max strategy: for each action, consider the worst-case caused by the opponent action, then choose the best own action for each worst-case.

Min-Max Strategy

- The payoff matrix is as follows (zero-sum game, only player I's utilities are shown.
- Which action should player I choose?

I \ II				
	7	2	5	1
	2	2	3	4
	5	3	4	4
	5	2	1	6

Saddle Point

- What if the player II thinks in the similar way?
- II wants to minimize I's utility.
- The crossing point is minimum in the row, and maximum in the column.

I \ II				
	7	2	5	1
	2	2	3	4
	5	3	4	4
	5	2	1	6

Saddle Point

- In a zero-sum game, if there exists a saddle point, then the result of the game will be that point.
- Not necessarily exists.

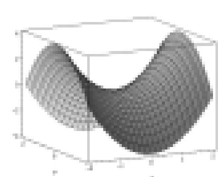

Case with no saddle point

- Assume a penalty kick in a soccer game.
- The goal keeper is very good for the right-side.
- If he is expecting the right-side:
  - the kick is actually in the right-side: he can stop 80%.
  - in the left-side: he cannot stop at all.
- expecting the left-side:
  - the kick is actually in the left-side: he can stop 30%.
  - in the right-side: he can stop 10%.

Keeper \ Kicker	R	L
R	8	0
L	1	3

Using Min-Max Strategy

- A timid keeper：I want to avoid the worst-case, let's wait for left-side, then I can stop at least 10%!
- A timid kicker:
  To be stopped 80% is terrible, let's kick left-side!
- Then, the keeper can stop 30%!

Keeper \ Kicker	R	L
R	8	0
L	1	3

Is this result reasonable?

- This result is too bad for the kicker.
  - Why I need to kick to the left while I'm quite sure that the keeper is expecting the left-side?
  - If I kick to the right-side, then the keeper can stop just 10%.
- On the other hand, if the keeper knows the kick is coming to the right-side, then he can do better.
  - He is good for the right-side!

Keeper \ Kicker	R	L
R	8	0
L	1	3

Mixed Strategy

- The kicker probabilistically mixes left/right.
- The keeper also changes whether to wait left/right.
- Such a strategy is called a mixed strategy.
  - Choosing a single action is called a pure strategy.

Keeper \ Kicker	R	L
R	8	0
L	1	3

Saddle point in a mixed strategy

- If the kick is tend to be right, then the keeper waits for the right-side more often, so that he can stop more.
- If the keeper tends to wait right-side more, then the kicker tries to kick to the left-side more often.
- Where is a stable point?

Keeper \ Kicker	R	L
R	8	0
L	1	3

Saddle point in a mixed strategy

- The probability that the keeper is expecting the right-side: x
- The probability that the kicker kicks to the right-side: y
- stopping probability：$0.1[xy*8 + x(1-y)*0 + (1-x)y*1 + (1-x)(1-y)*3]$
  $= 0.1[10xy - 3x - 2y + 3]$
- partially differentiate by x: $10y - 3$
  - $y=0.3$, i.e., if the kicker kicks to the right for 30%, then regardless of the keeper's strategy, the stopping probability is 24%.
- partially differentiate by y: $10x - 2$
  - if $x=0.2$, i.e., if the keeper waits for the right for 20%, then regardless of the kicker's strategy, the stopping probability is 24%.

Keeper \ Kicker	R	L
R	8	0
L	1	3

Is it really stable?

- Assume each player gradually adapt to the opponent's strategy…

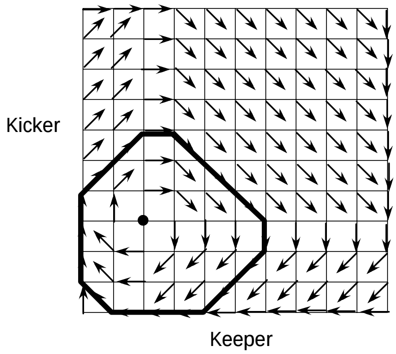

Pareto Efficiency

Definition:

A state is said Pareto efficient, iff there exists no state that is

- better for one agent, and
- no worse for all the rest

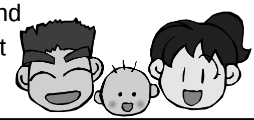

✕	Movie	2	2	2
	Shopping	2	2	5
	Zoo	2	3	1
✕	Home	1	1	1

Pareto Efficiency (cont'd)

- Divide $1000 between you and me.
- We can throw away some.
- (0, 1000), (x, 1000-x), (1000, 0), all are Pareto efficient.
- We can agree on (250+x, 750-x) is better than throw away $500 and get (250, 250).

Pareto Efficiency (cont'd)

- In principle, it is not clear whether we can compare utilities of different people.
  - Do we really have a common measure (e.g., money)?
- The definition of Pareto efficiency can be applied even if we cannot compare utilities.
- We can consider Pareto efficiency is a minimal requirement of a desirable social choice.
- If x is not Pareto efficient, then there is another choice x', where everybody prefers x' to x (or at least the same).

Rational Player

- A rational player tries to maximize its utility by all means, using possibly unlimited computational power.
- In a zero-sum or constant sum game, all outcomes are Pareto efficient.
- If a game is not zero/constant-sum, there is a possibility that both can be happy at the same time (to some extent).
- If players are irrational/incompetent, then both might become unhappy.
- However, if players are rational, then the result should be Pareto efficient.
  - It is unlikely that very clever players throw out some utilities.

Prisoners' Dilemma

- The police arrests two suspects (prisoners).
  - If both of them do not confess, both will be released.
  - If one confesses, while the other does not confess, the one who confessed will get a reward, while the other will get a severe punishment.
  - If both confess, both receives a normal punishment.

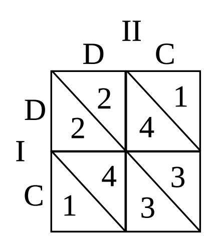

Prisoners' Dilemma

- (C, C): (3, 3) is Pareto efficient; both player can agree it is better than (D,D): (2, 2)
- It sounds irrational that a dominant strategy equilibrium is not Pareto efficient!
- Although players are very clever, they cannot reach a Pareto efficient situation.

If this game is repeated...

- Assume the game is repeated three times.
- Let's find an Iterated Dominance Equilibrium
- Let's consider the last (third) game.
- Your rational opponent will deceive you anyway.
- Then, it is meaningless to give a favor at the second game.
- Then, both player will deceive at the second game.
- Similarly, in the first game, both rational player will deceive.
- This is the same even if the game is repeated 1000 times...
- Note: if your opponent is irrational, then you might be better off by cooperating.

Prisoners' Dilemma Tournament

- Held by Robert Axelrod (Univ. Michigan, political scientist)
- Computer programs repeatedly play Prisoners' Dilemma.
- The program that obtains the highest total score wins.

Prisoners' Dilemma Tournament

- Surprisingly, a very simple program wins (called Tit-for-Tat)
- Cooperate in the first game.
- Then imitate the opponent's play in the previous round.
- If the opponent deceives, then it retaliate with D.
- As long as the opponent cooperate, it keeps on cooperate.
- Characteristic: Not persistent.

Prisoners' Dilemma Tournament

- There was a second tournament.
- Many programs try to beat Tit-for-Tat.
- However, Tit-for-Tat won again!
  - Actually, a new program beats Tit-for-Tat.
  - But when two new programs face each other, they often deceive and their obtained scores are low.
  - Tit-for-Tat does not defeat anybody, but plays relatively well for everybody.
  - As a result, the total score was highest.
- Lesson: the goal is not to defeat your opponent, but to receive the high score in total.
  - You need a strategy that can draw cooperation from the opponent.

Outline: Games with Complete Information

- Definition
- Dominant Strategy Equilibrium
- Iterated Dominance Equilibrium
- Min-Max Strategy
- Mixed Strategy
- Nash Equilibrium

Equilibrium of a Game

- What should we do if there exists no dominant strategy equilibrium or iterated dominance equilibrium?
- Let’s consider a weaker notion of equilibrium.
  – Nash equilibrium

Nash Equilibrium

- A set of strategies (s,t) is in Nash equilibrium if they are the best reply to each other.
- A dominant strategy equilibrium is a Nash equilibrium, but not vice versa.
- A saddle point in a zero-sum game is a Nash equilibrium.
- If there exists a unique Nash equilibrium, then the result of a game played by rational players would be that Nash equilibrium.
  – Other results are unstable.

II

I				
	7	2	5	1
	2	2	3	4
	5	3	4	4
	5	2	1	6

Multiple Nash Equilibria

- “The game of chicken” has two Nash equilibrium
  – (D, C)
  – (C, D)
- Not sure which one will occur.
- If a third-party player (who does not have any power, groundlessly) says “(C, D) will occur”, then it might be come true.

II

I	D	C
D	1 \ 1	4 \ 2
C	2 \ 4	3 \ 3

Nash equilibrium in mixed strategies

- Theorem：Any game has at least one Nash equilibrium in mixed strategies (Nash 1951) .
- In P-R-S, choosing each for probability 1/3 is a Nash equilibrium.

Quiz: Nash Equilibrium?

- Play P-R-S at a stairway.
- win by rock: advance 3 steps.
  win by paper: advance 6 steps.
  win by scissors: advance 6 steps.
- The one who reaches the top first wins.

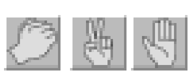

	0 \ 0	1 \ -1	-2 \ 2
	-1 \ 1	0 \ 0	2 \ -2
	2 \ -2	-2 \ 2	0 \ 0

Games with Incomplete Information

- Several sources of uncertainty
  – The utilities of opponents (types) are not known.
  – The result can be probabilistic (the choice of the nature)
  – In a game where plays are interleaved: the play of opponents cannot be observed.

Modified Game of Chicken

- There can be different types of players.
  - Bull: losing is as bad as dying
  - Chicken: be scared to death for not hitting brakes

Bull

	D	C
D	2, 1	4, 2
C	2, 4	3, 3

Chicken

	D	C
D	1, 1	3, 2
C	2, 4	3, 3

Bayesian Nash Equilibrium

- Assume the probability of each type normal/bull/chicken is 1/3.
- Assume this probability distribution is common knowledge.
- In the following strategy profile, each strategy maximizes the expected utility (given that other player uses this strategy).
  - A normal player chooses D/C for 0.5
  - A bull chooses D.
  - A chicken chooses C.
- Such a strategy profile is called Bayesian Nash equilibrium.

Review of Yesterday's Talk

- Dominant Strategy/Equilibrium
- Iterated Dominance Equilibrium
- Min-Max Strategy
- Mixed Strategy
- Nash Equilibrium
- Bayesian Nash Equilibrium

Signaling

- It is not clear whether the education at universities really improve the productivity of workers; then why people go to universities and companies hire university graduates with high salary?
- One possible answer: the university education works as a “signal” to distinguish high-quality workers and low-quality workers.

Problem Settings

- Worker:
  - There are high-quality worker (High) and low-quality worker (Low); the probability is ½ for each.
  - High can produce 6, while Low can produce 3.
  - The cost for graduating from a university is 0 for High, 3 for Low.
- Company: can choose whether to offer a worker a high salary (4) or low salary (1). It’s utility is the difference of the productivity and salary.
- High can obtain 3 by himself if he decided not to work for a company, while Low cannot make money alone.
- A worker can choose whether to go to a university.
- A company can set a salary according to the education level (or just ignore it).

What kinds of Bayesian Nash equilibrium exist?

Characteristics

- The type of a worker (High/Low) cannot be observed by a company.
- Getting the university education does not increase the productivity at all (in a sense, it is just a waste of efforts).

67

Separating Equilibrium

- The company offers a high salary (4) to university graduates, and a low salary (1) to others.
- High goes to the university: his utility is 4.
- Low does not go to the university: his utility is 1.
- The utility of a company is 2, i.e., (6-4)*0.5+(3-1)*0.5.
- Low cannot increase his utility if he goes to the university.
- For workers, the university education works as a signal to show his ability.

68

Pooling Equilibrium

- The company offers a high salary (4) to university graduates, and a low salary (1) to others.
- Everybody goes to the university.
- The utility of the company is: (6-4)*0.5 + (1-4)*1/2=0.5
- The company cannot increase the utility if it decides to hire everybody with low salary.
- Low cannot increase his utility if he decides not to go to the university.
- The signaling does not work; the university education is totally a waste in this case.
- Solution: make the university education more difficult for Low.

69

Quiz: Signaling

- Point out instances that seem to be "signaling".
  - Something that has no real value itself, but it works to distinguish people, company, product, etc.
  - Something that is relatively easy for a good guy, difficult for a bad guy.

70

Outline

- Example
- Why game-theory/economics?
- Game-theory
- Auctions
  - Assumptions/Preliminaries
  - Single-item, single-unit auctions
  - Combinatorial auctions
  - False-name bids

71

Characteristics of Player

risk neutral/averse

- risk neutral: only cares the expected utility
  - e.g., indifferent between two lotteries:
    1) he/she obtains 0 for the head and $100 for the tail,
    2) he/she obtains $50 for sure.
- risk averse: prefers that is more certain (even with less expected utility)
  - e.g., prefers getting $45 for sure to 1.

72

Assumption for Simplicity

Quasi-linear utility: an agent's utility is defined as the difference between its evaluation value of the allocated good and its payment.

- If an agent wins a good whose evaluation value is $100 by paying $90, its utility is $100-$90=$10.
- If it does not win, its utility is $0.
- If it does win, paying $100, its utility is 0.

Social Surplus

- Assuming agents' utilities are quasi-linear, if a state is Pareto efficient, the social surplus (sum of all agents' utilities) must be maximized.

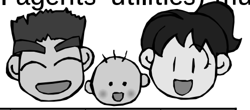

×	Movie	6	2	2	2
	Shopping	9	2	2 →3	5 →4
×	Zoo	6	2	3	1
×	Home	3	1	1	1

The St. Petersburg Paradox

Assume the following lottery...

- Flip a coin, if it comes up tail, you get $2.
- If head, flip a coin again, if it comes up tail, you get $4.
- If head, flip a coin again, ...
- If it comes up tail the first time at n-th trial, you get $\$2^n$.

How much are you willing to pay to participate this lottery?

Incomplete Information in Auctions

- Types of players
  - If the evaluation values of opponents are known, an auction becomes trivial
- His/her own evaluation value
  - There exists some uncertainty in the value of the auctioned good.

Private/Common/Correlated Values

Private Value: each agent knows its value with certainty, which is independent from other agents' evaluation values (e.g., antiques which are not resold).

Common value： the evaluation values for all agents are the same, but agents do not know the exact value and have different estimated values (e.g., US Treasury bills, mining right of oil fields).

Correlated Value: Something between above two extremes.

Desirable Properties of Auction Protocols

- For a bidder, there exists a dominant strategy.
- The protocol is robust against various frauds (e.g., spying).
- A Pareto efficient allocation can be achieved.

Pareto Efficiency in Auctions

- The social surplus must be maximized; the good must be allocated to the agent who has the highest evaluation value.
  - The agent whose evaluation value is $8000 wins and pays $7000.
  - Utility of this agent: $8000 - $7000=$1000
  - Utility of the seller: $7000
  - Social Surplus: $8000

$8000 $7000 $6000

Protocol/Mechanism Design

- Designing a protocol is determining the rules of a game.
- The designer cannot control the actions of each agent.
  - No way to force an agent to be honest or to refrain from doing frauds.

Protocol/Mechanism Design

How can a designer achieve a certain desirable property (e.g., Pareto efficiency)?

- Design rules so that:
  - For each agent, there exists a dominant strategy.
  - In the dominant strategy equilibrium, the desirable property is achieved.

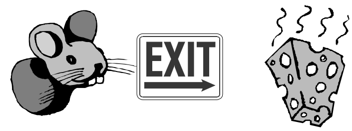

Incentive Compatibility

Direct revelation mechanism: directly ask types/evaluation values for each agents

Incentive compatibility: A direct revelation mechanism is (dominant-strategy) incentive compatible if truth-telling is a dominant strategy for each agent.

Revelation Principle: If a certain property (e.g., Pareto efficiency) can be achieved in a dominant strategy equilibrium using an indirect mechanism, that property can be achieved using an incentive compatible direct revelation mechanism.

We can restrict our attention only to (incentive compatible) direct revelation mechanism!

Revelation Principle

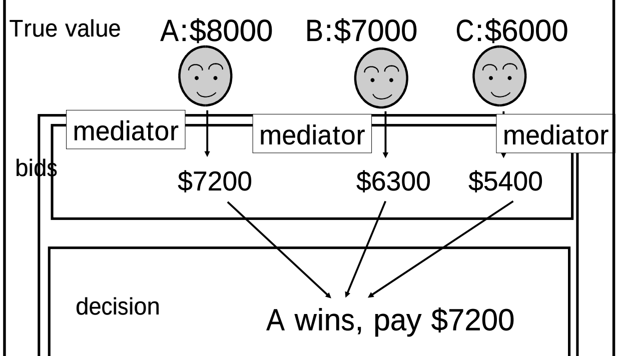

Outline

- Example
- Why game-theory/economics?
- Game-theory
- Auctions
  - Assumptions/Preliminaries
  - Single-item, single-unit auctions
  - Combinatorial auctions
  - False-name bids
- Market-based systems

English (open cry)

Protocol: Each agent is free to revise its bid upwards. When nobody wishes to revise its bid further, the highest bidder wins the good and pays its own price.

Dominant strategy (in private value): keep bidding some small amount more than the previous highest bid until the price reaches its evaluation value, then quit.

English (open cry)

- In the dominant-strategy equilibrium, the agent with the highest evaluation value wins and pays the second highest evaluation value + ε .
- The obtained allocation is Pareto efficient.

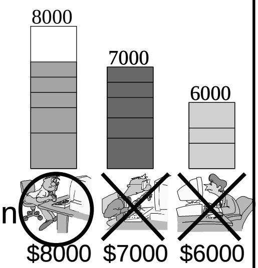

First-price Sealed-bid

Protocol：Each agent submits its bid without knowing other agents' bids. The agent with the highest bid wins and pays its own price.

Dominant strategy： does not exist in general.

Example: First-price

- Assume you attend an auction on behalf of your uncle.
- There are 10 goods to be auctioned.
- Your uncle specifies the maximal price you can bid for each good.
- The generous uncle will give you the difference if you can buy it less than the maximal price.
- If you failed to buy a good, you receive nothing for the good.

Example: First-price

- Assume there is only one opponent.
- Your opponent is also a proxy.
- You don't know how much your opponent will bid, but you know his maximal price is uniformly distributed among [0, 200].

Strategy for Bidding

- For good 1, assume the maximal price is $100.
- Then, bidding more than $100 is meaningless (you must pay the difference).
- If you bid $90 and win, your profit is $10.
- If you bid $1 and (by any chance) win, your profit is $99.
- How much should you bid?

Let's Play with Computer!

- There are 10 goods.
- You see your maximal price, which is chosen from a uniform distribution [0, 200].
- You know the maximal price of your opponent (computer) is also chosen from [0, 200].
- The computer player chooses an optimal strategy (in some sense).

Answer

- We cannot tell the expected utility unless knowing the strategy of the opponents.
- Let us assume the bidding price of the opponents is uniformly distributed between [0, 100]
- If you bid low, the probability of winning is low, but the utility when wins is large.
  ⇒high-risk, high-return
- If you bid high, the probability of winning is high, but the utility when wins is low.
  ⇒low-risk, low-return
- The best point is in between, bidding the half (50), where the expected utility is 25.

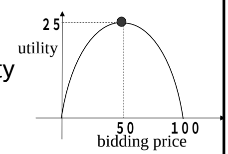

Answer (contd.)

- If the opponent uses the same strategy, his bid is uniformly distributed between [0, 100], since his maximal price is uniformly distributed between [0, 200].
- This pair of strategies is Bayesian Nash equilibrium.
- Even if you are not a proxy, i.e., the maximal price is your own evaluation value, you should use the same strategy.

Dutch (descending)

Protocol: The seller announces a very high price, then continuously lowers the price until some agent says “stop”, then the agent wins the good and pays the current price.

Dominant strategy： does not exist in general.

Dutch (descending)

- Strategically equivalent to the first-price sealed-bid auction
  – there is one-to-one mapping between the strategy sets in two auctions.
- Example:
  – Dutch flower market
  – Ontario tobacco auction
  – bargain sale

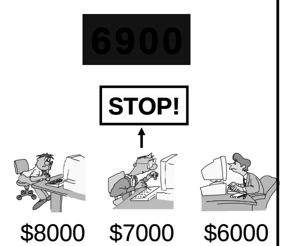

Vickrey (Second-price Sealed-bid)

Protocol：Each agent submits its bid without knowing other agents’ bids. The agent with the highest bid wins and pays the value of the second highest bid.

Dominant strategy (in Private value)： Bidding its true evaluation value is the dominant strategy (honesty is the best policy, incentive compatibility)

- The obtained allocation is Pareto efficient.
- The obtained result is identical to English in the dominant-strategy equilibrium.

Characteristics of Protocols (in Private Value Auctions)

- Dutch＝First-price Sealed-bid
- English＝Vickrey
- Under several assumptions, the expected revenue of the seller is the same in all 4 auction protocols (revenue equivalence theorem, Vickrey 1961）.
  – If there exists a Bayesian Nash equilibrium, the expected revenue would be the same at the equilibrium.

Quiz: Expected Utility (for Vickrey auction)

- Assume you are facing one opponent in a Vickrey auction.
- Your evaluation value is 100.
- The evaluation value of the opponent is uniformly distributed from 0 to 200.
- What is your best bid and how much is your expected utility?

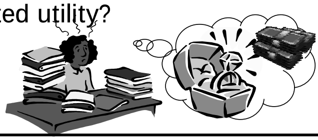

Answer

- Honesty is the best policy.
- By bidding the true evaluation value 100, the probability of winning is 0.5. When you win, your price/utility is uniformly distributed between 0 to 100 (in average, 50).
- Therefore, your expected utility is 25.
- We can see the revenue equivalence theorem holds: social surplus = revenue of the seller + utilities of buyers

Difficulties for Using Vickrey Auction

- Hard to understand!
- Do not know/aware of the true evaluation value (even in private value auctions).
- Cannot trust the seller.
- Do not want to reveal the private/sensitive information.

Common Value Auctions

- English and Vickrey can be different.
  - Agents can obtain more information in English.
  - Can revise the estimation using the obtained information.
- Dutch and First-price sealed-bid are still equivalent.

Quiz: Common Value Auction

- Assume there are two bidders.
- Each bidder has an estimated value, which is randomly chosen from [v-10, v+10], where v is the common value.
- Does the revenue equivalence theorem holds for English and Vickrey?
- In other words, can we create a strategy, that is possible for English, but impossible for Vickrey?

Answer

- The revenue equivalence theorem holds!
- We cannot create a strategy that is possible only for English.
- The only meaningful information obtained in English is that your opponent quits at a certain price, but the auction is already closed (and you won).

Winner's Curse

- In a common value auction, each agent does not know the real value of the auctioned good (which is common to all agents).
- Each agent has a different estimated value.
- Unless the agent has an especially good piece of information, the winner tends to be the agent that has the largest estimation error.
- If an agent increases its bid too close to its estimated value, the expected utility can be negative.

Winner's Curse (Example)

Settings: there are two bidders, the estimated value of the bidder can be either under-estimated (v-100) or over-estimated (v+100), where v is the real common value. Both are equally probable, i.e., each probability is 1/2.

- In a first-price sealed-bid auction, a bidder might (wrongly) think as follows:

  My estimated value is right in average. Thus, if I bid my estimated value - 40, my expected utility would be +40.

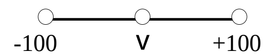

Winner's Curse (Example)

- Assume ties are broken by tossing a coin.
- Assume another bidder uses the same strategy:
  - There are 4 possible combinations of bids.
  - (v+60, v-140): with probability 1/4, the agent wins, the utility is –60.
  - (v+60, v+60): by tossing a coin, with probability 1/8, the agent wins,
    the utility is -60.
  - (v-140, v-140): by tossing a coin, with probability 1/8, the agent wins,
    the utility is 140.
- The expected utility is:
  -40/8 = -5

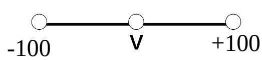

Quiz: Winner's Curse

- Assume you are considering buying company A.
- The value of A (i.e., v_A) is uniformly distributed between [0, 100].
- If you buy A, you can increase its value to 50% and sell.
- If you offer b, the owner of A will sell if $b > v_A$.
- If you buy at price b, your profit is $1.5\ v_A - b$.
- How much should you offer?

Answer: Winner's Curse

- If you offer b, the owner sells only if $b > v_{A.}$
- Assume this is true, then v_A is uniformly distributed between [0, b], thus in average, it's value is 0.5b.
- You can increase the value to 0.75b.
- Your profit is -0.25b!

Outline

- Example
- Why game-theory/economics?
- Game-theory
- Auctions
  - Assumptions/Preliminaries
  - Single-item, single-unit auctions
  - Combinatorial auctions
  - False-name bids

Combinatorial Auction

- Multiple different goods with correlated values are auctioned simultaneously.
  - Complementary: PC and memory
  - Substitutable: Dell or Gateway
- By allowing bids on any combinations of goods, the obtained social surplus/revenue of the seller can increase.
- e.g., FCC spectrum right auctions

Research Issues in Combinatorial Auctions (I)

- Finding the best combination of bids is a complicated combinatorial optimization problem
  - winner determination problem, one instance of a set packing problem
  - NP-complete
  - Various search techniques are introduced

Research Issues in Combinatorial Auctions (II)

- How to describe the preference of an agent is also a research issue --- 2^m subsets for m goods
  - Developing Bidding Languages
    - Compact, expressive, and allow efficient winner determination
  - Preference Elicitation

The Generalized Vickrey Auction Protocol (GVA)

- Each agent declares its evaluation values for subsets of goods.
- The goods are allocated so that the social surplus is maximized.
- The payment of agent 1 is equal to the decrease of the social surplus except agent 1, caused by the participation of agent 1.
- Satisfies incentive compatibility and Pareto efficiency.

An Example of the GVA

Setting: three agents (agent 1, 2, 3) are bidding for two goods.

	coffee	cake	both
Agent 1	$6	$0	$6
Agent 2	$0	$0	$8
Agent 3	$0	$5	$5

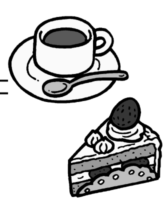

Result:

- Agent 1 gets the coffee, 3 gets the cake.
- Agent 1 pays \$8 − \$5 = \$3.
- Agent 3 pays \$8 − \$6 = \$2.

Incentive Compatibility of GVA

- Goods are allocated so that social surplus is maximized.
- An agent can maximize its utility when the social surplus is maximized.

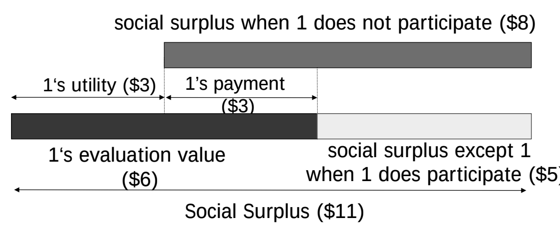

Quiz: GVA

- Assume there are 4 bidders.
- Who are going to win what, and how much to pay?

	coffee	cake	both
Agent 1	$70	$0	$70
Agent 2	$0	$0	$100
Agent 3	$0	$50	$50
Agent 4	$60	$20	$80

Quiz: Clarke Tax

- The GVA is one instance of the Clarke mechanism (a.k.a. Vickrey-Clarke-Groves mechanism, Clarke Tax).
- Can be used for more general setting in group decision making
  - Example: determine whether to extend this class 30 minutes.
    - You declare your monetary value for the choice (e.g., for extending, $20, -$10, etc., assuming not change is $0)
    - How can we guarantee that you will declare your true preference?

Answer

- Calculate the sum.
- If the sum is positive, we extend the class.
- If your vote changes the outcome, you are charged the minimum amount needed to change the outcome.
  - agent 1: $20, agent 2: -$10, agent 3: -$20, agent 4: $30
  - The result is to extend, the payments are:
    - agent 1: $0, agent 2: $0, agent 3: $0, agent 4: $10

Outline

- Example
- Why game-theory/economics?
- Game-theory
- Auctions
  - Assumptions/Preliminaries
  - Single-item, single-unit auctions
  - Combinatorial auctions
  - False-name bids

Internet Auction

- Internet auctions have become a particularly popular part of Electronic Commerce.
  - There exist many auction sites.

Merits:
  - can execute large-scale auctions with many more sellers and buyers from all over the world
  - can utilize software agents

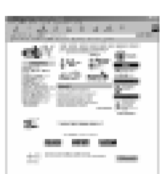

Demerit:
  - A new type of cheating using the anonymity available on the Internet is possible (false-name bids).

False-name Bids

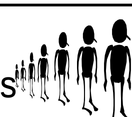

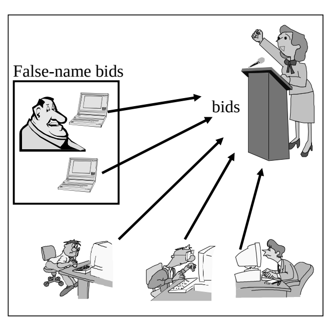

- An agent submits several bids under fictitious names.
- Detecting false-name bids is virtually impossible, since identifying each participant on the Internet is very difficult.

A Case where the GVA is Vulnerable

Setting: two agents (Agent 1, 2)

	coffee	cake	both
Agent 1	$6	$5	$11
Agent 2	$0	$0	$8

	coffee	cake	both
Agent 1	$6	$0	$6
Agent 2	$0	$0	$8
Agent 3	$0	$5	$5

When telling the truth:

- •Agent 1 gets both goods.
- •payment: $8 －$0 ＝$8

When agent 1 uses a false-name 3 and splits its bid:

- •Agent 1 gets both goods.
- •payment: $3+$2=$5

Main Research Results

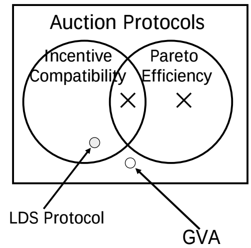

- ▮ found that the GVA is not robust against false-name bids
- ▮ proved that there exists no auction protocol that satisfies incentive compatibility and Pareto efficiency at the same time
- ▮ proved that revelation principle still holds when agents can submit false-name bids
- ▮ developed a new protocol (LDS protocol) that satisfies incentive compatibility, and can achieve a semi-optimal outcome

Non-existence Theorem

- No auction protocol exists that simultaneously satisfies incentive compatibility and Pareto efficiency at the same time for all cases, if agents can submit false-name bids.

Strategy of the Proof

- It is sufficient to show one instance where no auction protocol satisfies the prerequisites.
- By using the prerequisites, we clarify the bound of the payments and derive a contradiction.

Proof (Case 1)

two goods A and B, the evaluation value of an agent is represented as: (A only, B only, both)

- agent 1: (a, 0, a)
- agent 2: (0, 0, a+b)
- agent 3: (0, a, a)
- a > b
- By Pareto efficiency, each of agent 1 gets A and agent 3 gets B.
- By incentive compatibility, each pays b + ε (no incentive for under-bidding).

Proof (Case 2)

two agents

- agent 1: (a, a, 2a)
- agent 2: (0, 0, a+b)
- By Pareto efficiency, agent 1 gets both goods.
- By incentive compatibility, its payment is 2 (b + ε).
  - –Agent 1 can create the situation identical to case 1 using false-name bids.

Proof (Case 3)

- agent 1: (c, c, 2c)
- agent 2: (0, 0, a+b)
- $b + \varepsilon < c$, $2c < a + b$
- By Pareto efficiency, agent 2 gets both goods.
- If agent 1 lies and submit (a, a, 2a), it can create the situation identical to case 2.
- Agent 1 obtains both goods, its payment is $2(b + \varepsilon) < 2c$

Cannot satisfy incentive compatibility.

Explanation of the Proof

- To avoid free-riders problem, we must set the payment as low as possible in case 1.
- To avoid false-names, we also set the payment as low as possible in case 2.
- However, the payment becomes too low; thus an agent has an incentive for over-bidding (case 3)

Trivial Protocol (Set Protocol)

Protocol: always sell all goods in a bundle, and use the Vickrey auction protocol

- robust against false-name bids
- wasteful if the goods are substitutable for some agents

We need to develop a protocol that can sell goods separately in some cases.

Necessary Condition

- If goods A and B are sold separately, the sum of the payments must be larger than the highest evaluation value of the set.

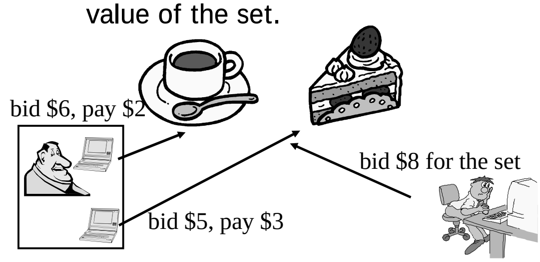

Vulnerable Protocol

Protocol: use the GVA to determine the (tentative) winners and payments. If the goods are sold separately and the payments does not satisfy the necessary condition, then sell the goods in a bundle; otherwise, use the result of the GVA.

Dilemma: satisfy the condition on payments without using the actual values of the payments

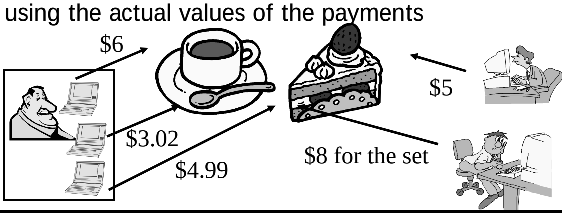

Solution: Utilize Reservation Prices

Reservation Price: the payment for good A (or B) must be larger than a predefined price r_A (or r_B) .

Protocol: if some agent values the set at more than r_A+r_B, then sell the goods in a bundle; otherwise, sell goods separately; never sell a good less than the reservation price.

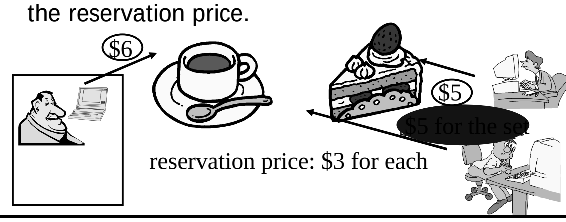

Leveled Division Set (LDS) Protocol

- The seller determines and announces a series of bundles (leveled division set) and reservation prices.

Protocol:

1. If the goods can be sold using level 1 division (there exists a bid that is larger than the sum of the reservation prices) then use the GVA considering only the current bundles.
2. If not, apply level 2, and so on.

The leveled division set is defined so that any union of bundles must appear in an smaller (earlier) level; give priority to agents that is willing to buy larger bundles.

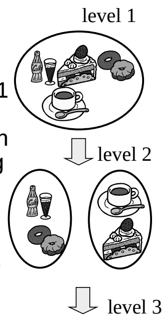

Return

Proof of Incentive Compatibility

Using false-name bids is useless:

- Assume agent x uses two identifiers x' and x'', and obtains S' and S''.
- If agent x uses a single identifier, it can obtain S'+S'' at the sum of the reservation prices in an earlier level.

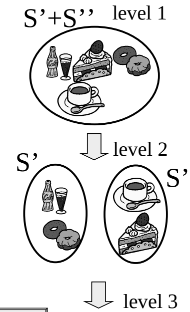

Return

Discussion (LDS Protocol)

- If the leveled division set and reservation prices are determined appropriately, the social surplus (or the revenue) can be larger than the set protocol.
- Communication/ computation costs are smaller than that for the GVA.

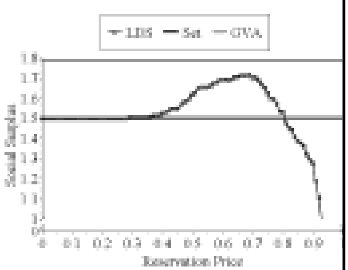

Return

Characterization of Strategy/False-name-proof Protocol

Let us think about a skeleton of a protocol called Price-oriented, Rationing Free (PORF) Protocol, which might look quite different from a standard auction protocol description…

- For each agent x, for each bundle B, a price p(x,B) is given.
- P(x,B) is set independently from the valuations of x (it depends on other agents' valuations)
- Each agent is guaranteed to obtain a bundle (or an empty bundle) that maximizes its utility based on the prices, regardless of the allocations of other agents (rationing-free).

Example of PORF Protocol

Single item, single unit auction:

The price of bidder x is the highest valuation except x.

–identical to Vickrey auction

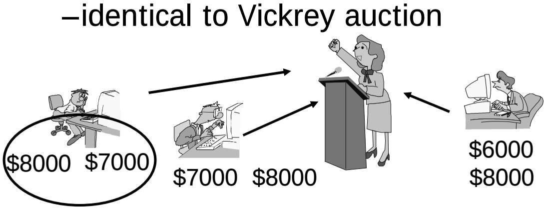

Characteristics of PORF Protocol

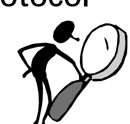

- Obviously, any PORF protocol is strategy-proof.
- Surprisingly, any strategy-proof protocol can be described as a PORF protocol.
- A PORF protocol that satisfies some additional conditions is false-name-proof, and vice versa.
  - Price of bundle (B1∪B2) ≦ Price of bundle B1 + Price of bundle B2

We can use PORF protocols as a guideline/standardized method to develop strategy/false-name proof protocols.

Another False-name-proof PORF Protocol

Price of bundle B (for bidder x): max valuation of all bundles B’, where B’∩B is non-empty and B’ is minimal (does not contain any useless item).

	coffee	cake	both
Bidder 1	$6 $8	$0 $8	$6 $8
Bidder 2	$0 $6	$0 $5	$8 $6
Bidder 3	$0 $8	$5 $8	$5 $8

Price of bundle (B1∪B2)

= max(Price of bundle B1, Price of bundle B2)

≦ Price of bundle B1 + Price of bundle B2

Further Readings

- Introductory-level textbooks:
  - Eric Rasmusen, Games & Information (3rd Edition), Blackwell Publishers, 2001.
  - Hal R. Varian, Intermediate Microeconomics, A Modern Approach (5th Edition), W.W. Norton & Company, 1999.
- Mid-level textbooks on economics in general:
  - Andreu Mas-Colell, Michael D. Whinston and Jerry R. Green, Microeconomic Theory, Oxford University Press, 1995.
- Extensive textbook on Auction
  - Vijay Krishna, Auction Theory, Academic Press, 2002.