CS 440/ECE448 Lecture 9: Game Theory

Slides by Svetlana Lazebnik, 9/2016

Modified by Mark Hasegawa-Johnson, 2/2019

Prisoner A \ Prisoner B	Prisoner B stays silent (*cooperates*)	Prisoner B betrays (*defects*)
Prisoner A stays silent (*cooperates*)	Each serves 1 year	Prisoner A: 3 years Prisoner B: goes free
Prisoner A betrays (*defects*)	Prisoner A: goes free Prisoner B: 3 years	Each serves 2 years

https://en.wikipedia.org/wiki/Prisoner's_dilemma

Game theory

- **Game theory** deals with systems of interacting agents where the outcome for an agent depends on the actions of all the other agents
  - Applied in sociology, politics, economics, biology, and, of course, AI
- **Agent design:** determining the best strategy for a rational agent in a given game
- **Mechanism design:** how to set the rules of the game to ensure a desirable outcome

Modelling behaviour

Game theory in practice

The Economist

Computing: Software that models human behaviour can make forecasts, outfox rivals and transform negotiations

Sep 3rd 2011 | from the print edition

http://www.economist.com/node/21527025

http://www.spliddit.org

WIRED

CADE METZ BUSINESS 09.21.15 7:00 AM

FACEBOOK DOESN'T MAKE AS MUCH MONEY AS IT COULD—ON PURPOSE

YOU CAN THINK of John Hegeman as Facebook's chief economist. He spends his days thinking about the economics of Facebook advertising.

That's an enormous thing. Facebook pulled in $4.04 billion in the second quarter of this year. And the overall economy of Facebook advertising, as Hegeman describes it, is far larger. Advertising, you see, is very much a part of everything else on the world's largest social network. Hegeman doesn't just think about ads. He thinks about how ads fit with the rest of Facebook.

When he joined Facebook in 2007, after getting a master's in economics at Stanford University, Hegeman helped build the online auction that drives the company's advertising system. Auctions are the standard way that online services accept ads from advertisers and place them on web pages and inside smartphone apps. That's what Google uses with AdWords, the system that serves up all those ads when you look for stuff on the company's Internet search engine. Advertisers bid (in dollars) for placement on the results page when you key in a particular word or group of words. But in building Facebook's advertising system, Hegeman and team took online auctions in a new direction.

http://www.wired.com/2015/09/facebook-doesnt-make-much-money-couldon-purpose/

Outline of today's lecture

- Nash equilibrium, Dominant strategy, and Pareto optimality
- Stag Hunt: Coordination Games
- Chicken: Anti-Coordination Games, Mixed Strategies
- The Ultimatum Game: Continuous and Repeated Games
- Mechanism Design: Inverse Game Theory

Nash Equilibria, Dominant Strategies, and Pareto Optimal Solutions

Recall: Multi-player, non-zero-sum game

- Players act in sequence.
- Each player makes the move that is best for them, when it's their turn to move.

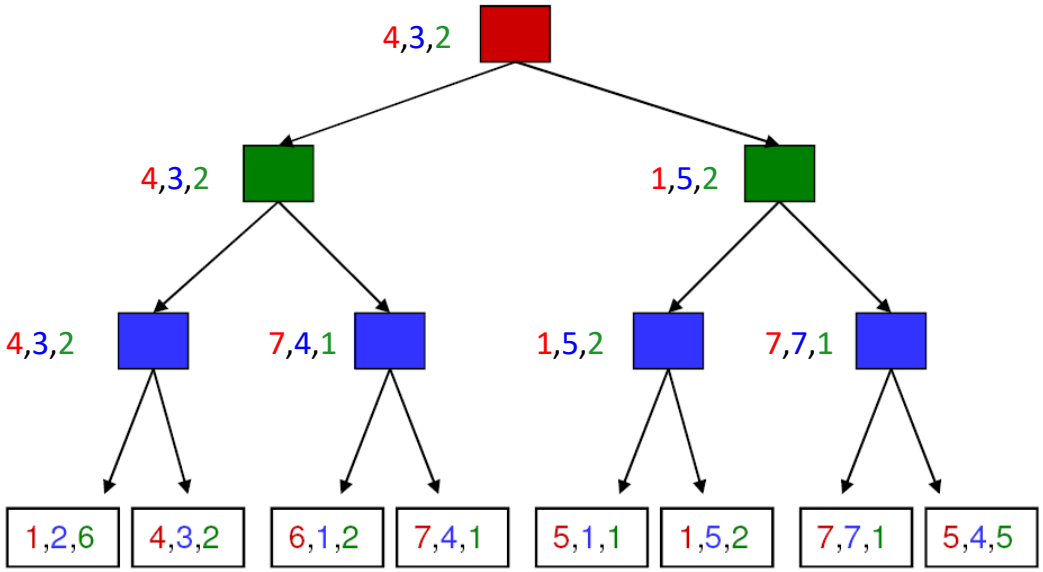

Simultaneous single-move games

- Players must choose their actions at the same time, without knowing what the others will do
  - Form of partial observability

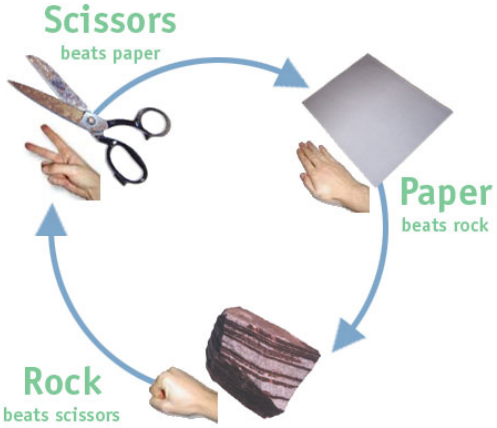

Normal form representation:

Player 2 \ Player 1			
	0,0	1,-1	-1,1
	-1,1	0,0	1,-1
	1,-1	-1,1	0,0

Payoff matrix
(Player 1's utility is listed first)

Is this a zero-sum game?

Prisoner's dilemma

- Two criminals have been arrested and the police visit them separately
- If one player testifies against the other and the other refuses, the one who testified goes free and the one who refused gets a 10-year sentence
- If both players testify against each other, they each get a 5-year sentence
- If both refuse to testify, they each get a 1-year sentence

	Alice: Testify	Alice: Refuse
Bob: Testify		
Bob: Refuse		

Prisoner's dilemma

- Alice's reasoning:
  - Suppose Bob testifies. Then I get 5 years if I testify and 10 years if I refuse. So I should testify.
  - Suppose Bob refuses. Then I go free if I testify, and get 1 year if I refuse. So I should testify.
- **Nash equilibrium:** A pair of strategies such that no player can get a bigger payoff by switching strategies, provided the other player sticks with the same strategy
  - (Testify, Testify) is a *Nash equilibrium*

	Alice: Testify	**Alice: Refuse**
Bob: Testify	-5,-5	-10,0
Bob: Refuse	0,-10	-1,-1

Prisoner's dilemma

- **Dominant strategy:** A strategy whose outcome is better for the player regardless of the strategy chosen by the other player.
  - **TESTIFY!**
- **Pareto optimal outcome:** It is impossible to make one of the players better off without making another one worse off.
  - (Testify, Refuse)
  - (Refuse, Refuse)
  - (Refuse, Testify)
- Other games can be constructed in which there is no dominant strategy – we'll see some later

	Alice: Testify	**Alice: Refuse**
Bob: Testify	-5,-5	-10,0
Bob: Refuse	0,-10	-1,-1

Prisoner's dilemma in real life

- Price war
- Arms race
- Steroid use
- Diner's dilemma
- Collective action in politics

	Defect	**Cooperate**
Defect	Lose – lose	Lose big – win big
Cooperate	Win big – lose big	Win – win

http://en.wikipedia.org/wiki/Prisoner's_dilemma

Is there any way to get a better answer?

- Superrationality
  - Assume that the answer to a symmetric problem will be the same for both players
  - Maximize the payoff to each player while considering only identical strategies
  - Not a conventional model in game theory
  - … same thing as the Categorical Imperative?
- Repeated games
  - If the number of rounds is fixed and known in advance, the equilibrium strategy is still to defect
  - If the number of rounds is unknown, cooperation may become an equilibrium strategy

The Stag Hunt: Coordination Games

	Hunter 1: Stag	Hunter 1: Hare
Hunter 2: Stag	2,2	1,0
Hunter 2: Hare	0,1	1,1

- Both hunters cooperate in hunting for the stag → each gets to take home half a stag
- Both hunters defect, and hunt for rabbit instead → each gets to take home a rabbit
- One cooperates, one defects → the defector gets a bunny, the cooperator gets nothing at all

	Hunter 1: Stag	Hunter 1: Hare
Hunter 2: Stag	2,2	1,0
Hunter 2: Hare	0,1	1,1

- What is the Pareto Optimal solution?
- Is there a Nash Equilibrium?
- Is there a Dominant Strategy for either player?
- Model for cooperative activity under conditions of incomplete information (the issue: trust)

Prisoner's dilemma vs. stag hunt

Prisoner' dilemma

	Cooperate	Defect
Cooperate	Win – win	Win big – lose big
Defect	Lose big – win big	Lose – lose

Players improve their winnings by defecting unilaterally

Stag hunt

	Cooperate	Defect
Cooperate	Win big – win big	Win – lose
Defect	Lose – win	Win – win

Players reduce their winnings by defecting unilaterally

Chicken: Anti-Coordination Games, Mixed Strategies

Game of Chicken

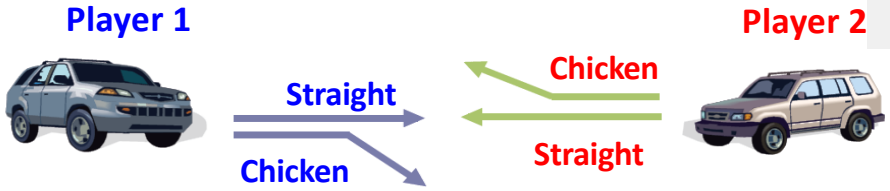

	S	C
S	-10, -10	-1, 1
C	1, -1	0, 0

- Two players each bet $1000 that the other player will chicken out
- Outcomes:
  - If one player chickens out, the other wins $1000
  - If both players chicken out, neither wins anything
  - If neither player chickens out, they both lose $10,000 (the cost of the car)

http://en.wikipedia.org/wiki/Game_of_chicken

Prisoner's dilemma vs. Chicken

Prisoner' dilemma

	Cooperate	Defect
Cooperate	Win – win	Win big – Lose big
Defect	Lose big – Win big	Lose – Lose

Players can't improve their winnings by unilaterally cooperating

Chicken

	Chicken	Straight
Chicken	Nil – Nil	Win – Lose
Straight	Lose – Win	Lose big – Lose big

The best strategy is always the opposite of what the other player does

Game of Chicken

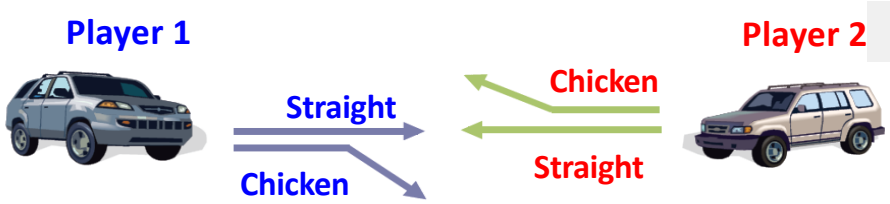

	S	C
S	-10, -10	-1, 1
C	1, -1	0, 0

- Is there a dominant strategy for either player?
- Is there a Nash equilibrium?

  (straight, chicken) or (chicken, straight)
- *Anti-coordination* game: it is mutually beneficial for the two players to choose different strategies
  - Model of escalated conflict in humans and animals (hawk-dove game)
- How are the players to decide what to do?
  - Pre-commitment or threats
  - Different roles: the "hawk" is the territory owner and the "dove" is the intruder, or vice versa

http://en.wikipedia.org/wiki/Game_of_chicken

Mixed strategy equilibria

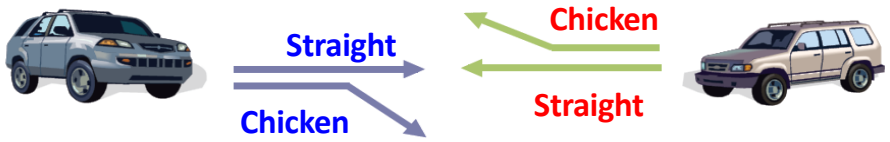

	S	C
S	-10, -10	-1, 1
C	1, -1	0, 0

- **Mixed strategy:** a player chooses between the moves according to a probability distribution
- Suppose each player chooses S with probability 1/10.
  Is that a Nash equilibrium?
- Consider payoffs to P1 while keeping P2's strategy fixed
  - The payoff of P1 choosing S is (1/10)(–10) + (9/10)1 = –1/10
  - The payoff of P1 choosing C is (1/10)(–1) + (9/10)0 = –1/10
  - Can P1 change their strategy to get a better payoff?
  - Same reasoning applies to P2

Finding mixed strategy equilibria

	P1: Choose S with prob. p	**P1:** Choose C with prob. $1-p$
P2: Choose S with prob. q	-10, -10	-1, 1
P2: Choose C with prob. $1-q$	1, -1	0, 0

- Expected payoffs for P1 given P2's strategy:
  - P1 chooses S: $q(-10) + (1-q)1 = -11q + 1$
  - P1 chooses C: $q(-1) + (1-q)0 = -q$
- In order for P2's strategy to be part of a Nash equilibrium, P1 has to be indifferent between its two actions:
  - $-11q + 1 = -q$ or $q = 1/10$
  - Similarly, $p = 1/10$

Existence of Nash equilibria

- Any game with a finite set of actions has at least one Nash equilibrium (which may be a mixed-strategy equilibrium)
- If a player has a dominant strategy, there exists a Nash equilibrium in which the player plays that strategy and the other player plays the *best response* to that strategy
- If both players have *strictly dominant* strategies, there exists a Nash equilibrium in which they play those strategies

Computing Nash equilibria

- For a two-player zero-sum game, simple linear programming problem
- For non-zero-sum games, the algorithm has worst-case running time that is exponential in the number of actions
- For more than two players, and for sequential games, things get pretty hairy

Nash equilibria and rational decisions

- If a game has a *unique* Nash equilibrium, it will be adopted if each player
  - is rational and the payoff matrix is accurate
  - doesn't make mistakes in execution
  - is capable of computing the Nash equilibrium
  - believes that a deviation in strategy on their part will not cause the other players to deviate
  - there is *common knowledge* that all players meet these conditions

http://en.wikipedia.org/wiki/Nash_equilibrium

The Ultimatum Game: Continuous and Repeated Games

Continuous actions: Ultimatum game

- Alice and Bob are given a sum of money S to divide
  - Alice picks A, the amount she wants to keep for herself
  - Bob picks B, the smallest amount of money he is willing to accept
  - If $S - A \geq B$, Alice gets A and Bob gets $S - A$
  - If $S - A < B$, both players get nothing
- What is the Nash equilibrium?
  - Alice offers Bob the smallest amount of money he will accept: $S - A = B$
  - Alice and Bob both want to keep the full amount: $A = S$, $B = S$ (both players get nothing)
- How would humans behave in this game?
  - If Bob perceives Alice's offer as unfair, Bob will be likely to refuse
  - Is this rational?
    - Maybe Bob gets some positive utility for "punishing" Alice?

http://en.wikipedia.org/wiki/Ultimatum_game

Sequential/repeated games and threats: Chain store paradox

- A monopolist has branches in 20 towns and faces 20 competitors successively
  - Threat: respond to “in” with “aggressive”

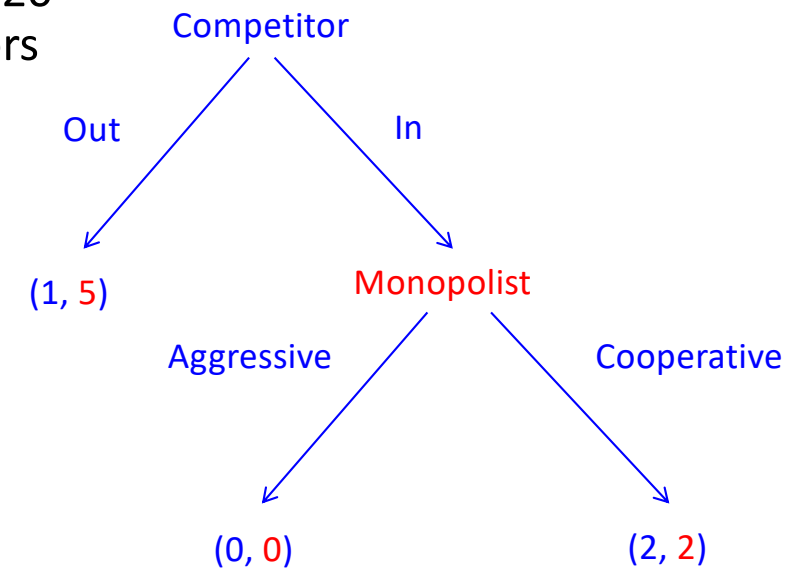

https://en.wikipedia.org/wiki/Chainstore_paradox

Mechanism Design: Inverse Game Theory

Mechanism design (inverse game theory)

- Assuming that agents pick rational strategies, how should we design the game to achieve a socially desirable outcome?
- We have multiple agents and a **center** that collects their choices and determines the outcome

Auctions

- Goals
  - Maximize revenue to the seller
  - Efficiency: make sure the buyer who values the goods the most gets them
  - Minimize transaction costs for buyer and sellers

Ascending-bid auction

- What's the optimal strategy for a buyer?
  - Bid until the current bid value exceeds your *private value*
- Usually revenue-maximizing and efficient, unless the reserve price is set too low or too high
- Disadvantages
  - Collusion
  - Lack of competition
  - Has high communication costs

Sealed-bid auction

- Each buyer makes a single bid and communicates it to the auctioneer, but not to the other bidders
  - Simpler communication
  - More complicated decision-making: the strategy of a buyer depends on what they believe about the other buyers
  - Not necessarily efficient
- **Sealed-bid second-price auction:** the winner pays the price of the second-highest bid
  - Let V be your private value and B be the highest bid by any other buyer
  - If $V > B$, your optimal strategy is to bid above B – in particular, bid V
  - If $V < B$, your optimal strategy is to bid below B – in particular, bid V
  - Therefore, your dominant strategy is to bid V
  - This is a **truth revealing** mechanism

Dollar auction

A malevolent twist on the second-price auction:

- Highest bidder gets to buy the object, and pays whatever they bid
- Second-highest bidder is required to pay whatever they bid, but gets nothing at all in return

- Dramatization: https://www.youtube.com/watch?v=pA-SNscNADk

Dollar auction

- A dollar bill is auctioned off to the highest bidder, but the second-highest bidder has to pay the amount of his last bid
  - Player 1 bids 1 cent
  - Player 2 bids 2 cents
  - …
  - Player 2 bids 98 cents
  - Player 1 bids 99 cents
    - If Player 2 passes, he loses 98 cents, if he bids $1, he might still come out even
  - So Player 2 bids $1
    - Now, if Player 1 passes, he loses 99 cents, if he bids $1.01, he only loses 1 cent
  - …
- What went wrong?
  - When figuring out the expected utility of a bid, a rational player should take into account the future course of the game
- What if Player 1 starts by bidding 99 cents?

Regulatory mechanism design: Tragedy of the commons

- States want to set their policies for controlling emissions
  - Each state can reduce their emissions at a cost of -10 or continue to pollute at a cost of -5
  - If a state decides to pollute, -1 is added to the utility of every other state
- What is the dominant strategy for each state?
  - Continue to pollute
  - Each state incurs cost of -5-49 = -54
  - If they all decided to deal with emissions, they would incur a cost of only -10 each
- Mechanism for fixing the problem:
  - Tax each state by the total amount by which they reduce the global utility (**externality cost**)
  - This way, continuing to pollute would now cost -54

Review: Game theory

- Normal form representation of a game
- Dominant strategies
- Nash equilibria
- Pareto optimal outcomes
- Pure strategies and mixed strategies
- Examples of games
- Mechanism design
  - Auctions: ascending bid, sealed bid, sealed bid second-price, “dollar auction”