

# Computer modeling of physical phenomena



Lecture XI Game Theory


# Game theory

Game theory studies the ways in which strategic interactions among rational players produce outcomes with respect to the players' preferences (The outcomes might not have been intended by any of them)

# First book on game theory




since war is a game...



Sun Tzu, 544–496 BC

# Modern game theory



Neumann, J., and Morgenstern, O., (1947). The Theory of Games and Economic Behavior.



Nash, J., (1950 ). Equilibrium Points in n-Person Games, PNAS, 36, 48-49

Nash, J., (1950 ). Non-Cooperative Games, PhD Thesis, Princeton, 1950 (27pp)

# John and Alicia Nash



# Rules, Strategies, Payoffs, and Equilibrium

- The **rules** of the game state who can do what, and when they can do it
- A player's **strategy** is a plan for actions in each possible situation in the game
- A player's **payoff** is the amount that the player wins or loses in a particular situation in a game
- A player has a **dominant strategy** if his best strategy doesn't depend on what other players do

# Zero sum or not

**Zero-sum** describes a situation in which a participant's gain (or loss) is exactly balanced by the losses (or gains) of the other participant(s) The total gains of the participants minus the total losses always equals 0

Poker is a zero sum game

- The money won = the money lost

Trade is not a zero sum game

- If a country with an excess of bananas trades with another for their excess of apples, both may benefit from the transaction

Non-zero sum games are more complex to analyze **and more interesting!**

# Nash Equilibrium

- a situation when each player's strategy is optimal, given the strategies of the other players (no player can benefit from unilaterally changing his strategy, while all other players stay fixed)
- Every finite game has at least one Nash equilibrium in either pure (deterministic) or mixed (probabilistic) strategies,

(Nash, 1950; Nobel Prize in economics, 1994)



# Prisoner's Dilemma



Merrill Flood and Melvin Dresher (RAND, 1950)

# Intro: Buick sale



In June 1949, Flood wanted to buy a used Buick from a RAND employee who was moving back East. Buyer and seller were friends. They weren't looking to cheat each other, just to agree on a fair price for the car. How should they set a price?

- Dealer: buy for 500\$, sell for 800\$, should they split an extra \$300 profit?
- What if one defects (I will sell for 750\$!) - the other is still better off to accept it than to go the dealer

**The party who is more unreasonable is apt to get the better of the deal!**

based on: William Poundstone, *Prisoner's Dilemma*, 1992

# Other social experiments by Flood

- Two secretaries: either the prize 100\$ for one or 150\$ for both. How would they split?
- Reverse auction for the babysitting among Flood's three teenage children: the child agreeing to baby-sit for the lowest pay would get the job
- The opening bid was set at \$4.00, ended with 90c offer (no cooperation!)


Flood: “This is probably an extreme example, although not really so extreme when the magnitude of the children's error is compared with that made by mature nations at war because of inability to split-the-difference. I have noticed very similar “irrational” behaviour in many other real life situations since August 1949 and find it to be commonplace rather than rare.”

# Prisoner's dilemma

Bonnie and Clyde are arrested by the police and charged with various crimes. They are questioned in separate cells, unable to communicate with each other. They know how it works:

- If both resist interrogation (cooperating with each other) and proclaim mutual innocence, they will get a two year sentence for robbery
- If one confesses (defecting) and the other doesn't (cooperating), the confesser goes free and the other will get a severe 8-year sentence
- If they both confess (defecting), then the judge will sentence both to a moderate four years in prison

What should Bonnie do? What should Clyde do?




# The payoff matrix

		CLYDE	
		Confess	Not Confess
BONNIE	Confess	4 years each	0 years for Bonnie and 5 for Clyde
	Not Confess	5 years for Bonnie and 0 for Clyde	2 years each

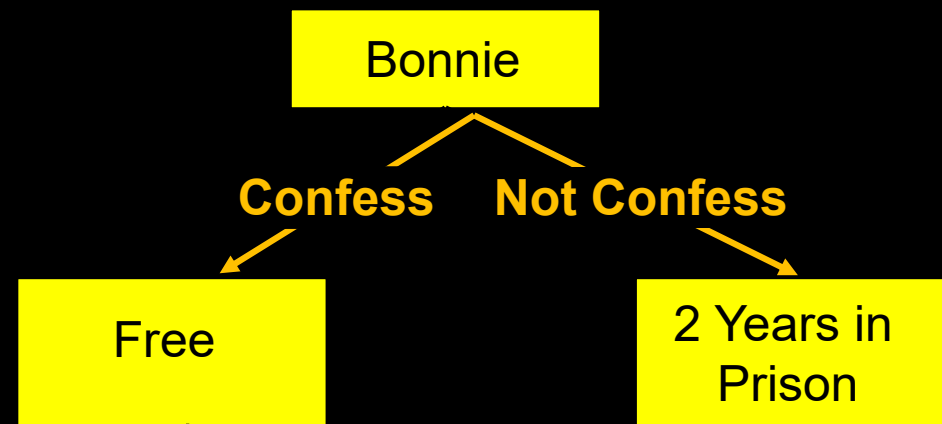
# Bonnie's Decision Tree

**If Clyde Confesses**



Best Strategy

**If Clyde Does Not Confess**




Best Strategy

The dominant strategy for Bonnie is to confess (defect) because no matter what Clyde does she is better off confessing.

# The dismal science

It seems we should always defect and never cooperate....






[Economics is] not a 'gay science,' I should say, like some we have heard of; no, a dreary, desolate, and, indeed, quite abject and distressing one: what we might call, by way of eminence, the dismal science.

(Thomas Carlyle)

# Examples...

- Cheating on a cartel:
  - Cartel members can charge the monopoly price or a lower price
  - The best strategy is to charge the low price
- Trade wars between countries
  - Free trade benefits both trading countries
  - Tariffs can benefit one trading country
  - Imposing tariffs can be a dominant strategy and establish a Nash equilibrium even though it may be inefficient
- Advertising
  - All firms advertising tends to equalize the effects
  - Everyone would gain if no one advertised
- Class team projects

# Examples...



Copyright © 1996 United Feature Syndicate, Inc.  
Redistribution in whole or in part prohibited

# Payoff matrix for the generic symmetric two person dilemma game

(A's payoff, B's payoff)

		Player B	
		cooperate	defect
Player A	cooperate	(CC,CC) reward for mutual cooperation	(CD,DC) sucker's payoff and temptation to defect
	defect	(DC,CD) temptation to defect and sucker's payoff	(DD,DD) punishment for mutual defection

# Payoffs

Four payoffs involved:

- CC: Both players cooperate
- CD: You cooperate but other defects (aka “sucker’s payoff”)
- DC: You defect and other cooperates (aka “temptation to defect”)
- DD: Both players defect


Some classical dilemma games:

Prisoner’s dilemma:  $DC > CC > DD > CD$

Chicken:  $DC > CC > CD > DD$

Stag Hunt / Trust dilemma:  $CC > DC > DD > CD$

# Chicken




Rebel Without  
a Cause (1955)

# Chicken

DC > CC > CD > DD

- Cooperation: swerving
- Defecting: not swerving



- The worst thing that can happen is for both players *not* to swerve
- The best thing that can happen, is to show your machismo by not swerving and letting the other driver swerve.
- Being chicken is the next to worst outcome, but still better than dying.
- There is a cooperative outcome in chicken. It's not so bad if both players swerve. Both come out alive, and no one can call the other a chicken.

# Example payoff matrix

		Player B	
		swerve	drive straight
Player A	swerve	2, 2	1, 3
	drive straight	3, 1	0, 0

- No dominant strategy: the best is to do opposite to what the other player does
- Playing safe? Swerving has the maximum minimum
- **Chicken vs the prisoner's dilemma:** Mutual defection (the crash when both players drive straight) is the most feared outcome in chicken. In the prisoner's dilemma, cooperation while the other player defects (being the sucker) is the worst outcome.

# Trust Dillema



CC > DC > DD > CD

		Player B	
		haircut	no haircut
Player A	haircut	5, 5	0, 4
	no haircut	4, 0	2, 2

- the best possible outcome would be for both of you to get the haircut
- it would be awful to be the *only* one to show up with the haircut
- If you *didn't* get the haircut, but the friend did, and looked like a real jerk, that would almost be as good as if both of you got the haircut.
- it wouldn't really be so bad if *no one* got the haircut

# Trust Dillema



CC > DC > DD > CD

		Player B	
		haircut	no haircut
Player A	haircut	5, 5	0, 4
	no haircut	4, 0	2, 2

**risk dominant strategy** - no haircut to minimize the risk

**payoff dominant strategy** - haircut to get the maximum reward

The more uncertainty players have about the actions of the other players, the more likely they will choose the risk dominant strategy

# Morality?



# How to save morality? Iterated Prisoner's Dilemma

Game theory shows that a rational player should always defect when engaged in a prisoner's dilemma situation

We know that in real situations, people don't always do this

Why not? Possible explanations:

- People aren't rational
- Morality
- Social pressure
- Fear of consequences
- Evolution of species-favoring genes

In frames of game theory, some of these ideas can be explored by playing the iterated version of the games

# Iterated Prisoner's Dilemma

- **Key idea:** In many situations, we play more than one “game” with a given player.
- Players have complete knowledge of the past games, including their choices and the other player's choices.
- Your choice in future games when playing against a given player can be partially based on whether he has been cooperative in the past.
  
- Repeated encounters with another player you can recognize.
- Car garages that depend on repeat business versus those on busy highways.
- Regular vs irregular clients on the markets
- If you're unlikely to ever see someone again, you're back into a non-iterated dilemma.

# Axelrod & Hamilton tournament

- round-robin computer tournament between strategies submitted by colleagues
- each game between competing strategies consisted of 200 moves, yielding possible scores from 0 to 1000, so that a "good performance" score was judged to be 600 points (complete mutual cooperation) and a "poor performance" score would be 200 points (complete mutual defection).
- the overall winner with an average of 504 points per game, submitted by game theorist Anatol Rapoport, was so called Tit-for-Tat: a player cooperates on the first move and then does exactly what the other player did on the previous move.
- While never actually attaining a higher score in any particular game, this strategy can lead to many rounds of cooperation between opponents with its commensurate rewards, whereas the more defecting strategies do poorly.

Robert Axelrod, and William D. Hamilton. "The evolution of cooperation *Science* 211.4489 (1981): 1390-1396 – more than 28K cites

# Axelrod on Tit-for-Tat



"What accounts for Tit-for-Tat's robust success is its combination of being nice [not being the first to defect] , retaliatory, forgiving, and clear. Its niceness prevents it from getting into unnecessary trouble. Its retaliation discourages the other side from persisting whenever defection is tried. Its forgiveness helps restore mutual cooperation. And its clarity makes it intelligible to the other player, thereby eliciting long-term cooperation" (Axelrod [1984]).

# GA for Prisoner's Dillema

You can try to do better than TfT by emplying genetic algorithms.


Some hints can be found here:

<http://www-personal.umich.edu/~axe/research/Evolving.pdf>

# Spatial version of PD

- cooperation is easier to maintain in a sedentary population: defectors can thrive in an anonymous crowd, but mutual aid is frequent among neighbors.
- as a model of the above, one can consider PD played on a lattice with either Moore or von Neumann neighbourhood
- various strategies are allowed to compete against each other.
- Initially, strategies are randomly assigned to the cells
- Each cell then competes for a fixed number of rounds with each of its neighbors
- At the end of each session a cell will adopt the strategy of its most successful neighbor in terms of highest total score

# Spatial version of PD



The evolution of the spatial iterated prisoner's dilemma with five competing strategies from an initial random configuration. Each cell's strategy plays off five times against its eight neighbors and then adopts the strategy of its most successful neighbor. Images are taken at the initial state, after 2, 5, 10, and 50 iterations.

 Always Cooperate

 Tit-for-Tat

 Pavlov


 Always Defect

 Random

Pavlov - repeats its former choice whenever it wins a round (C/C or D/C) and switches otherwise (If winning, hold fast; if not, change course)

TfT seems to win, but this depends on initial conditions!

# The same with noise



A certain level of noise introduced, that is, a cell will make a random move in a single round, thus increasing the dynamics of the evolution.

-  Always Cooperate
-  Tit-for-Tat
-  Pavlov
-  Always Defect
-  Random

Can you devise a better strategy for spatial PD using GA?