CUR 412: Game Theory and its Applications, Lecture 1

Prof. Ronaldo CARPIO

Feb. 28, 2020

Welcome to CUR412

- This course is an introduction to Game Theory, the study of *strategic* situations (i.e. situations with more than one decision-maker)
- This course will be taught entirely in English.
- Website: `http://rncarpio.com/teaching/CUR412`
- Announcements, slides, & homeworks will be posted on website

About Me: Ronaldo Carpio

- BS Electrical Engineering & CS, UC Berkeley
- Master's in Public Policy, UC Berkeley
- PhD Economics, UC Davis
- Joined School of Banking & Finance in 2012

Textbook

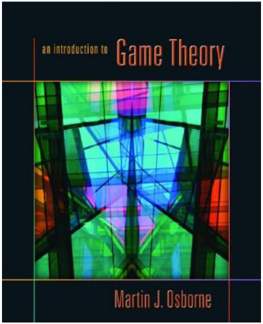

- *An Introduction to Game Theory* (2003) by Martin Osborne
- A Chinese translation is available, it is recommended that you get it if you have difficulty following the textbook
- The first 3 chapters are available for download on the textbook webpage (see course website for details)
- Prerequisites: you should be familiar with optimization using derivatives, and basic probability

Grading

- Homework 15%, Midterm Exam 35%, Final Exam 50%
- Last semester's midterm and final exams are posted on the website
- Homework:
  - There will be 5 homework assignments, posted on the website
  - Write-ups must be individual, you may discuss the concepts in small groups
  - Exam problems will be similar to homework. The best way to prepare is to do the homework problems yourself
- Exam dates: to be announced

Contacting Me

- Email: rncarpio@yahoo.com
- Wechat: rncarpio
- Office: 913 Boxue Bldg

Course Outline

- Introduction & Motivation (Ch. 1)
- Static Games, Nash Equilibrium (Ch. 2)
- Nash Equilibrium: Applications (Ch. 3)
- Mixed Strategies & Mixed Strategy Equilibrium (Ch. 4)
- Extensive Form Games (Ch. 5)
- Sequential Games & Backwards Induction (Ch. 6, 7)
- Games with Imperfect Information (Ch. 10)
- Repeated Games (Ch. 14, 15)

What is Game Theory, and Why do we Need It?

- *Game Theory* is the mathematical study of *strategic* situations, i.e. where there is more than one decision-maker, and each decision-maker can affect the outcome.
- Previously in microeconomics, you studied *single-person* problems. For example:
  - How much of each good to consume, in order to maximize my utility?
  - How much output should a firm produce, in order to maximize profits?
- Rational behavior: choose the level that maximizes utility (or profits, or payoffs).
- However, in multi-agent situations, my choice may change your problem.
- We need a method that takes everyone's choices into account.

Examples of Strategic Situations

- Business
  - Competition between firms: price, quality, location...
  - Market segmentation by firm: offer different levels of quality
  - Auctions
- Political Science
  - Voting Strategically: always vote for your candidate, or vote to ensure your least preferred candidate loses?
- Sports
  - Tennis Serving
  - Soccer Penalty Kicks
- Biology
  - Why do animals confront each other, but rarely fight? (Hawk-Dove game)
  - Why does the peacock have a huge, costly tail? (Signaling, Handicap Principle)

Mathematical Definition of a Strategic (or Normal-Form) Game

- Terminology:
  - The decision-makers are called **players**.
  - Each player has a set of possible **actions**. A list of actions (e.g. the list of what players choose) is called an *action profile*.
  - Each player has *preferences* over the outcome of the game. The outcome is determined by the actions that all players have chosen. Some outcomes are more desirable than others.
- A *strategic game* is a model of interaction in which each player chooses an action *without knowing* what other players choose
- We can think of this as players choosing their actions *simultaneously*.

Mathematical Definition of a Strategic (or Normal-Form) Game

- We need to specify:
  - who the players are
  - what they can do
  - their preferences over the possible outcomes
- Definition: A *strategic game* consists of:
  - a set of *players*
  - for each player, a set of *actions*
  - for each player, *preferences* (i.e. a ranking) over all possible action profiles
- We will usually use *payoff functions* that represent preferences, instead of using preferences directly.

A 2-Player Static Game: The Prisoner's Dilemma

- Let's consider a specific example. Imagine this situation:
  - There are two suspects in a crime.
  - Each suspect can be convicted of a minor offense, but can only be convicted of a major offense if the other suspect "finks" (i.e. gives information to the police).
  - Each suspect can choose to be *Quiet* (don't inform) or *Fink* (inform).
  - If both stay quiet, each gets 1 year in prison.
  - If only one suspect finks, he goes free while the other suspect gets 4 years.
  - If both suspects fink, they both get 3 years.

A 2-Player Static Game: The Prisoner's Dilemma

- Each suspect has 2 choices: *Quiet* or *Fink*.
- Let's see what would happen in each possible case:

Choice		Prison Sentence	
Suspect 1	Suspect 2	Suspect 1	Suspect 2
Quiet	Fink	4 years	goes free
Quiet	Quiet	1 year	1 year
Fink	Fink	3 years	3 years
Fink	Quiet	goes free	4 years

- Consider the first row, where Suspect 1 receives 4 years.
- Suppose the police offers Suspect 1 the option of changing his mind, and chooses *Fink* instead.
- Then his sentence would go down from 4 years to 3 years, while Suspect 2's sentence would go up from 0 to 3 years.

Choice		Prison Sentence	
Suspect 1	Suspect 2	Suspect 1	Suspect 2
Quiet	Fink	4 years	goes free
Quiet	Quiet	1 year	1 year
Fink	Fink	3 years	3 years
Fink	Quiet	goes free	4 years

- Consider the third row, where both suspects choose *Fink* and get 3 years.
- Suppose the police offered either suspect the option of changing his mind and choosing *Quiet* instead.
- That suspect would raise his sentence from 3 to 4 years.

Modeling the Prisoner's Dilemma

- Players: The two suspects.
- Actions: Each player's set of actions is Q, F.
- Preferences: We'll write down the action profile as: (Suspect 1's choice, Suspect 2's choice).
- Suspect 1's preferences, from best to worst:
  - $(F, Q) > (Q, Q) > (F, F) > (Q, F)$
- Suspect 2's preferences, from best to worst:
  - $(Q, F) > (Q, Q) > (F, F) > (F, Q)$
- Instead of using preferences directly, we will use a payoff function that assigns a utility to each outcome:
  - Suspect 1:
    $u_1(F, Q) = 3, u_1(Q, Q) = 2, u_1(F, F) = 1, u_1(Q, F) = 0$
  - Suspect 2:
    $u_2(F, Q) = 0, u_2(Q, Q) = 2, u_2(F, F) = 1, u_2(Q, F) = 3$

Bi-Matrix Form of Prisoner's Dilemma

		Player 2	
		Q	F
Player 1	Q	2,2	0,3
	F	3,0	1,1

- We can collect the payoff values into a *payoff matrix*:
- The two rows are the two possible actions of Player 1.
- The two columns are the two possible actions of Player 2.
- In each cell, the first number is the payoff of Player 1; the second is the payoff of Player 2.

Let's Play the Prisoner's Dilemma

- Everyone should have two cards: one Black and one Red card.
- How to play:
  - Start with two players, each with a Black and Red card.
  - Each player chooses to play Black or Red, and puts the card facedown.
  - Reveal both cards at the same time (why?)
- Suppose you are Player 1. If you play Red, then you get $+2$ and Player 2 gets $+0$.
- If you play Black, you get $+0$ and the other player gets $+3$.
- So, Red is beneficial to *you*, while Black benefits the *other* player.

	Black	*Red*
Black	3,3	0,5
Red	5,0	2,2

Modeling Other Situations as a Prisoner's Dilemma

- Suppose you are working with a friend on a joint project.
- Each of you can choose to *Work hard* or *Goof off* (be lazy).
- If the other person *Works hard*, each of you prefers to *Goof off*.
- Project would be better if both work hard, but not worth the extra effort.

		Player 2	
		Work hard	*Goof off*
Player 1	*Work hard*	2,2	0,3
	Goof off	3,0	1,1

Duopoly

- Two firms produce the same good.
- Each firm can charge a *High* price or a *Low* price.
- If both firms charge a high price, both get profit of 1000.
- If only one firm charges a high price, it loses customers, makes loss of 200. Other firm charges low price, gets profit of 1200
- If both firms charge low price, both get profit of 600.

	High	*Low*
High	1000, 1000	-200,1200
Low	1200,-200	600,600

Similarities to Prisoner's Dilemma

- Names of actions and payoffs are different, but *relative* payoffs are the same
- Preferences (i.e. ranking) over outcomes are the same as in Prisoner's Dilemma
- If both players cooperate, both get an outcome with good payoffs
- But if only one player chooses to defect, he gets an even better payoff (and cooperating player gets low payoff)

Applications of Prisoner's Dilemma

- Arms Race
  - Players: Countries
  - Actions: Arm, Disarm
- Provision of a Public Good
  - Players: Citizens
  - Actions: Contribute, Free-Ride
- Managing a Common Resource (Tragedy of the Commons)
  - Players: Animal Herders
  - Actions: Reduce Grazing, Overgraze

Bach or Stravinsky? (also known as Battle of the Sexes)

- Two people want to go to a concert by either *Bach* or *Stravinsky*.
- They prefer to go to the same concert, but one person prefers Bach while the other prefers Stravinsky.
- If they go to different concerts, each is equally unhappy.

	Bach	*Stravinsky*
Bach	2, 1	0, 0
Stravinsky	0, 0	1, 2

Let's Play Bach or Stravinsky

	Black (Bach)	*Red (Stravinsky)*
Black (Bach)	2, 1	0, 0
Red (Stravinsky)	0, 0	1, 2

Matching Pennies

- Prisoner's Dilemma and BoS have both conflict and cooperation. Matching Pennies is purely conflict.
- Each of two people chooses either *Head* or *Tail*.
- If the choices differ, Player 1 pays Player 2 \$1.
- If they are the same, Player 2 pays Player 1 \$1.
- Each person cares only about the money he receives.

	Head	*Tail*
Head	1,-1	-1,1
Tail	-1,1	1,-1

Let's Play Matching Pennies

	Black (Head)	*Red (Tail)*
Black (Head)	1,-1	-1,1
Red (Tail)	-1,1	1,-1

Solution Concept

- We've defined the game. What outcomes are more likely to occur?
- A *solution concept* (or *solution theory*) is a way of saying certain outcomes are less reasonable than others.
- A solution concept has two parts:
  - An assumption about the *behavior* of the players. We will assume rational behavior, i.e. choosing the action with the highest payoff
  - An assumption about the *beliefs* of the players.

Beliefs

- Suppose you are Player 1. In order to choose the best action, you need to have some idea of what Player 2 will choose
- This is called a *belief* about Player 2. Includes: the rules of the game, Player 2's payoff function, but also...
- Player 2 will also have a belief about you, which includes your beliefs about him, etc...
- Reasoning about what other players know (and what they know you know...) is called *higher-order knowledge*.
- We'll make a (very strong!) simplifying assumption: beliefs of all players are *correct*

Terminology

- subscript i denotes player i or an action of player i
- subscript $-i$ denotes all other players except i, or their actions
- Action profile (i.e. a list of all actions chosen by all players) a^* is composed of a_i^* and a_{-i}^*:

$$a^* = (a_i^*, a_{-i}^*)$$

- a_i^* is the action chosen by player i
- a_{-i}^* is the set of actions chosen by everyone except player i

Nash Equilibrium

- This solution concept assumes that:
  - Players are rational (i.e. choose the highest payoff), given beliefs about other players
  - Beliefs of all players are correct
- We want to find an outcome that is a *steady state*, that is, starting from that outcome, no player wants to deviate.
- If an action profile a^* is a steady state, then all the players must *not* have other actions that they could play, that are *more preferable* to their current action in a^*.
- **Definition**: The action profile a^* in a strategic game is a **Nash Equilibrium** if, for every player i and every action b_i of player i, a^* is *at least* as preferable for player i as the action profile (b_i, a^*_{-i}):

$$u_i(a^*) \geq u_i(b_i, a^*_{-i}) \quad \text{for every action } b_i \text{ of player } i$$

Nash Equilibrium

- Note that this definition does not guarantee that a game has a Nash equilibrium.
- Some games may have one, more than one, or zero Nash equilibria.

Prisoner's Dilemma

		Player 2	
		Q	F
Player 1	Q	2,2	0,3
	F	3,0	1,1

- (F, F) is the unique Nash equilibrium. No other action profile satisfies the conditions:
- (Q, Q) does not satisfy conditions, since $u_1(Q, Q) < u_1(F, Q)$
- $(F, Q) : u_2(F, Q) < u_2(F, F)$
- $(Q, F) : u_1(Q, F) < u_1(F, F)$
- Thus, the outcome predicted by the Nash equilibrium solution concept is that both players will *defect*.
  - Joint project: both will *Goof off*
  - Duopoly: both will charge a *Low* price (this is bad for the firms, but good for consumers)

Prisoner's Dilemma

	Q	F
Q	2,2	0,3
F	3,0	1,1

- Note that F is the best action for each player, *regardless* of what the other player does. (This is not the case in other games).
- However, (Q, Q) is a better outcome for both players than (F, F).
- Individual rationality can lead to a socially inefficient outcome.
- How might players reach the better outcome, while still behaving rationally (payoff-maximizing)?
- Need to change the structure of the game, e.g.:
  - External: laws, contracts, reputation
  - Internal: emotions, social norms

BoS

	Bach	*Stravinsky*
Bach	2, 1	0, 0
Stravinsky	0, 0	1, 2

- Starting from $(Bach, Bach)$, no player can get a higher payoff by changing his action.
- Same for $(Stravinsky, Stravinsky)$.
- For $(Bach, Stravinsky)$ or $(Stravinsky, Bach)$, at least one player has an incentive to deviate.
- Two Nash equilibria: $(Bach, Bach)$ and $(Stravinsky, Stravinsky)$.
- Both outcomes are compatible with a steady state.

Matching Pennies

	Head	*Tail*
Head	1,-1	-1,1
Tail	-1,1	1,-1

- There is *no* Nash equilibrium.
- For every action profile, at least one player has an incentive to deviate.
- There will never be a steady state in this situation.

Next Week

- Next week, we will look at ways of finding Nash equilibria (if they exist).
- For next week's lecture, please read Chapter 1 and Chapters 2.1-2.5 in Osborne.
- If you don't have the textbook, check the course website for information on how to get the textbook chapters.