Game Theory Intro

CMPUT 654: Modelling Human Strategic Behaviour

S&LB §3.2-3.3.3

Lecture Outline

1. Recap
2. Noncooperative game Theory
3. Normal form games
4. Solution concept: Pareto Optimality
5. Solution concept: Nash equilibrium
6. Mixed strategies

Recap: Utility Theory

- **Rational preferences** are those that satisfy axioms
- Representation theorems:
  - von Neumann & Morgenstern: Any rational preferences over **outcomes** can be represented by the maximization of the **expected value** of some **scalar** utility function
  - Savage: Any rational preferences over **acts** can be represented by maximization of the expected value of some scalar utility function with respect to **some probability distribution**

(Noncooperative) Game Theory

- **Utility theory** studies rational **single-agent** behaviour
- **Game theory** is the mathematical study of interaction between multiple **rational**, self-interested agents
  - **Self-interested**: Agents pursue only their **own preferences**
  - *Not* the same as "agents are psychopaths"! Their preferences may include the well-being of other agents.
  - Rather, the agents are **autonomous**: they decide on their own priorities independently.

Fun Game: Prisoner's Dilemma

	Cooperate	Defect
Cooperate	-1,-1	-5,0
Defect	0,-5	-3,-3

Two suspects are being questioned separately by the police.

- If they both remain silent (**cooperate** -- i.e., with each other), then they will both be sentenced to **1 year** on a lesser charge
- If they both implicate each other (defect), then they will both receive a reduced sentence of **3 years**
- If one defects and the other cooperates, the defector is given immunity (0 years) and the cooperator serves a full sentence of **5 years**.

Play the game with someone near you. Then find a new partner and play again. Play 3 times in total, against someone new each time.

Normal Form Games

The Prisoner's Dilemma is an example of a **normal form game**.
Agents make a single decision **simultaneously**, and then receive a payoff depending on the profile of actions.

Definition: Finite, n-person normal form game

- N is a set of n ***players***, indexed by i
- $A = A_1 \times A_2 \times \ldots \times A_n$ is the set of **action profiles**
    - A_i is the **action set** for player i
- $u = (u_1, u_2, \ldots, u_n)$ is a **utility function** for each player
    - $u_i : A \to \mathbb{R}$

Normal Form Games as a Matrix

	Cooperate	Defect
Cooperate	-1, -1, 1	-5, 0, 5
Defect	0,-5, 5	-3,-3, 3

Truthful

	Cooperate	Defect
Cooperate	-1,-1	-5,0
Defect	0,-5	-3,-3

	Cooperate	Defect
Cooperate	-1,-1, 1	-5, -5, 7
Defect	-5,-5, 7	-5, -5, 7

Lying

- Two-player normal form games can be written as a matrix with a tuple of utilities in each cell
- By convention, row player is first utility, column player is second
- Three-player normal form games can be written as a set of matrices, where the third player chooses the matrix

Games of Pure Competition (Zero-Sum Games)

Players have **exactly opposed** interests

- There must be precisely **two** players
  - Otherwise their interests can't be exactly opposed
- For all action profiles $a \in A$, $u_1(a) + u_2(a) = c$
  - $c=0$ without loss of generality by affine invariance
- In a sense it's a one-player game
  - Only need to store a single number per cell
  - But also in a deeper sense, by the Minimax Theorem

Matching Pennies

Row player wants to match, column player wants to mismatch

	Heads	Tails
Heads	1,-1	-1,1
Tails	-1,1	1,-1

Play against someone near you. Repeat 3 times.

Games of Pure Cooperation

Players have **exactly the same** interests.

- For all $i,j \in N$ and $a \in A$, $u_i(a) = u_j(a)$
- Can also write these games with one payoff per cell

Question: In what sense are these games **non-cooperative**?

Coordination Game

Which side of the road should you drive on?

	Left	Right
Left	1	-1
Right	-1	1

Play against someone near you.
Play 3 times in total, playing against someone new each time.

General Game:
Battle of the Sexes

The most interesting games are simultaneously both **cooperative and competitive**!

	Ballet	Soccer
Ballet	2, 1	0, 0
Soccer	0, 0	1, 2

Play against someone near you.
Play 3 times in total, playing against someone new each time.

Optimal Decisions in Games

- In single-agent decision theory, the key notion is **optimal decision**: a decision that maximizes the agent's expected utility
- In a multiagent setting, the notion of optimal strategy is **incoherent**
    - The best strategy **depends** on the strategies of others

Solution Concepts

- From the viewpoint of an **outside observer**, can some outcomes of a game be labelled as **better** than others?
  - We have no way of saying one agent's interests are more important than another's
  - We can't even **compare** the agents' utilities to each other, because of affine invariance! We don't know what "**units**" the payoffs are being expressed in.
- Game theorists identify certain subsets of outcomes that are interesting in one sense or another. These are called **solution concepts**.

Pareto Optimality

- Sometimes, some outcome *o* is **at least as good** for any agent as outcome *o*', and there is some agent who **strictly prefers** *o* to *o*'.
    - In this case, *o* seems defensibly better than *o*'

Definition: *o* **Pareto dominates** *o*' in this case

Definition: An outcome *o** is **Pareto optimal** if no other outcome Pareto dominates it.

Questions:

1. Can a game have more than one Pareto-optimal outcome?
2. Does every game have at least one Pareto-optimal outcome?

Pareto Optimality of Examples

	Coop.	Defect
Coop.	-1,-1	-5,0
Defect	0,-5	-3,-3

	Left	Right
Left	1	-1
Right	-1	1

	Ballet	Soccer
Ballet	2, 1	0, 0
Soccer	0, 0	1, 2

	Heads	Tails
Heads	1,-1	-1,1
Tails	-1,1	1,-1

Best Response

- Which **actions** are better from an **individual agent's** viewpoint?
- That depends on what the other agents are doing!

Notation:

$$a_{-i} \doteq (a_1, a_2, \ldots, a_{i-1}, a_{i+1}, \ldots, a_n)$$
$$a = (a_i, a_{-i})$$

Definition: **Best response**

$$BR_i(a_{-i}) \doteq \{a_i^* \in A_i \mid u_i(a^*, a_{-i}) \geq u_i(a_i, a_{-i}) \;\; \forall a_i \in A_i\}$$

Nash Equilibrium

- Best response is not, in itself, a solution concept
  - In general, agents won't know what the other agents will do
  - But we can use it to define a solution concept
- A Nash equilibrium is a **stable** outcome: one where no agent regrets their actions

Definition:
An action profile $a \in A$ is a (pure strategy) **Nash equilibrium** iff

$$\forall i \in N,\ a_i \in BR_{-i}(a_{-i})$$

Questions:

1. Can a game have more than one pure strategy Nash equilibrium?
2. Does every game have at least one pure strategy Nash equilibrium?

Nash Equilibria of Examples

The only equilibrium of Prisoner's Dilemma is also the *only* outcome that is Pareto-dominated!

	Coop.	Defect
Coop.	-1,-1	-5,0
Defect	0,-5	-3,-3

	Left	Right
Left	1	-1
Right	-1	1

	Ballet	Soccer
Ballet	2, 1	0, 0
Soccer	0, 0	1, 2

	Heads	Tails
Heads	1,-1	-1,1
Tails	-1,1	1,-1

Mixed Strategies

- So far, we have been assuming that agents play a single action **deterministically**
    - But that's a pretty bad idea in, e.g., Matching Pennies

Definition:

- A **strategy** s_i for agent i is any probability distribution over the set A_i, where each action a_i is played with probability $s_i(a_i)$.
    - **Pure strategy**: only a single action is played
    - **Mixed strategy**: randomize over multiple actions
- Set of i's strategies: $S_i \doteq \Delta(A_i)$
- Set of **strategy profiles**: $S \doteq S_1 \times \ldots \times S_n$

Utility Under Mixed Strategies

- The utility under a mixed strategy profile is **expected utility**
  - Because we assume agents are decision-theoretically rational
  - We assume that the agents randomize **independently**

Definition:

$$u_i(s) = \sum_{a \in A} u_i(a) \Pr(a \mid s)$$

$$\Pr(a \mid s) = \prod_{j \in N} s_j(a_j)$$

Best Response and Nash Equilibrium

Definition:
The set of *i*'s **best responses** to a strategy profile $s \in S$ is

$$BR_i(s_{-i}) \doteq \{s_i^* \in S \mid u_i(s_i^*, s_{-i}) \geq u_i(s_i, s_{-i}) \ \ \forall s_i \in S_i\}$$

Definition:
A strategy profile $s \in S$ is a **Nash equilibrium** iff

$$\forall i \in N, \ \ s_i \in BR_{-i}(s_{-i})$$

- When at least one s_i is mixed, s is a **mixed strategy Nash equilibrium**

Nash's Theorem

Theorem: [Nash 1951]
Every game with a finite number of players and action profiles has at least one Nash equilibrium.

Proof idea:

1. Brouwer's fixed-point theorem guarantees that any continuous function from a simpletope to itself has a fixed point.

2. Construct a continuous function $f : S \to S$ whose fixed points are all Nash equilibria.

   - NB: S is a simpletope, because it is the product of simplices

Interpreting
Mixed Strategy Nash Equilibrium

What does it even mean to say that agents are playing a mixed strategy Nash equilibrium?

- They truly are **sampling a distribution** in their heads, perhaps to **confuse** their opponents (e.g., soccer, other zero-sum games)
- The distribution represents the **other agents' uncertainty** about what the agent will do
- The distribution is the **empirical frequency** of actions in repeated play
- The distribution is the frequency of a pure strategy in a **population** of pure strategies (i.e., every individual plays a pure strategy)

Summary

- Game theory studies the **interactions of rational agents**
  - Canonical representation is the **normal form game**
- Game theory uses **solution concepts** rather than optimal behaviour
  - "Optimal behaviour" is not clear-cut in multiagent settings
  - **Pareto optimal**: no agent can be made better off without making some other agent worse off
  - **Nash equilibrium**: no agent regrets their strategy given the choice of the other agents' strategies