Nash Equilibrium

Game Theoretic Analysis

Kevin Leyton-Brown

University of British Columbia
Canada CIFAR AI Chair, Amii

James R. Wright

University of Alberta
Canada CIFAR AI Chair, Amii

Lecture Overview

Recap

Solution Concepts: Pareto Optimality

Solution Concepts: Nash Equilibrium

Mixed Strategies

Recap: Normal Form Games

In a normal form game:

- Agents **simultaneously** make a **single decision**
- They then receive an outcome that depends on the **profile of actions**

Definition: n-player normal form game

A normal form game is a tuple $G = (N, A, u)$, where

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

Recap: Normal Form Games as a Matrix

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

- 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 utility

Lecture Overview

Optimal Decisions in Games

- In single-agent decision theory, the key notion is the **optimal decision**: The decision that maximizes the agent's expected utility:

$$a^* = \arg\max_{a \in A} \mathbb{E}[u(a)]$$

Optimal Decisions in Games

- In single-agent decision theory, the key notion is the **optimal decision**: The decision that maximizes the agent's expected utility:

$$a^* = \arg\max_{a \in A} \mathbb{E}[u(a)]$$

- In a multiagent setting, the notion of an optimal strategy is ill-defined:

$$a_i^* = \arg\max_{a_i \in A_i} \mathbb{E}[u_i(a_i, a_{-i})]$$

Optimal Decisions in Games

- In single-agent decision theory, the key notion is the **optimal decision**: The decision that maximizes the agent's expected utility:

$$a^* = \arg\max_{a \in A} \mathbb{E}[u(a)]$$

- In a multiagent setting, the notion of an optimal strategy is ill-defined:

$$a_i^* = \arg\max_{a_i \in A_i} \mathbb{E}[u_i(a_i, a_{-i})]$$

- The best strategy depends on the strategies of the **other agents**
- But the other agents are simultaneously solving the same problem!

Solution Concepts

- From the viewpoint of an **outside observer**, can some outcomes of a game be considered **better** than others?
  - We have no justification for saying that one agent's interests are more important than another's
  - We cannot even **compare** the agents' utilitys 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 desirable and/or interesting
- These are called **solution concepts**

Pareto Optimality

Suppose outcome o is **at least as good** as o' for every agent i

- Further, there is *some* agent who **strictly prefers** o to o'
- E.g., $o' =$ "Everyone gets pie", and
  $o =$ "Everyone gets pie *and also* Alice gets cake"

Pareto Optimality

Suppose outcome o is **at least as good** as o' for every agent i

- Further, there is *some* agent who **strictly prefers** o to o'
- E.g., o' = "Everyone gets pie", and
  o = "Everyone gets pie *and also* Alice gets cake"
- In this situation, o seems defensibly better than o'

Pareto Optimality

Suppose outcome o is **at least as good** as o' for every agent i

- Further, there is *some* agent who **strictly prefers** o to o'
- E.g., $o' =$ "Everyone gets pie", and
  $o =$ "Everyone gets pie *and also* Alice gets cake"
- In this situation, o seems defensibly better than o'

Definition

o **Pareto dominates** o' whenever $o \succeq_i o'$ for **all** $i \in N$, and $o \succ_i o'$ for **some** $i \in N$

Definition

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

Pareto Optimality

Suppose outcome o is **at least as good** as o' for every agent i

- Further, there is *some* agent who **strictly prefers** o to o'
- E.g., $o' =$ "Everyone gets pie", and
  $o =$ "Everyone gets pie *and also* Alice gets cake"
- In this situation, o seems defensibly better than o'

Definition

o **Pareto dominates** o' whenever $o \succeq_i o'$ for **all** $i \in N$, and $o \succ_i o'$ for **some** $i \in N$

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?

Pareto Optimality

Suppose outcome o is **at least as good** as o' for every agent i

- Further, there is *some* agent who **strictly prefers** o to o'
- E.g., $o' =$ "Everyone gets pie", and $o =$ "Everyone gets pie *and also* Alice gets cake"
- In this situation, o seems defensibly better than o'

Definition

o **Pareto dominates** o' whenever $o \succeq_i o'$ for **all** $i \in N$, and $o \succ_i o'$ for **some** $i \in N$

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

Which outcomes are Pareto optimal in our running examples?

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

Pareto Optimality of Examples

Which outcomes are Pareto optimal in our running examples?

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

	Left	Right
Left	$1, 1$	$-1, -1$
Right	$-1, -1$	$1, 1$

Pareto Optimality of Examples

Which outcomes are Pareto optimal in our running examples?

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

	Left	Right
Left	$1, 1$	$-1, -1$
Right	$-1, -1$	$1, 1$

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

Pareto Optimality of Examples

Which outcomes are Pareto optimal in our running examples?

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

	Left	Right
Left	$1, 1$	$-1, -1$
Right	$-1, -1$	$1, 1$

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

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

Lecture Overview

Best Response

We can also ask: Which **actions** are better from an **individual agent's** viewpoint?

- That depends on what the other agents are doing!

Best Response

We can also ask: Which **actions** are better from an **individual agent's** viewpoint?

- That depends on what the other agents are doing!

Notation

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

Best Response

We can also ask: Which **actions** are better from an **individual agent's** viewpoint?

- That depends on what the other agents are doing!

Notation

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

Definition: Best response

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

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 called Nash equilibrium

A Nash equilibrium is a **stable** outcome: one where no agent regrets their action.

Definition

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

$$\forall i \in N : a_i \in BR_i(a_{-i})$$

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 called Nash equilibrium

A Nash equilibrium is a **stable** outcome: one where no agent regrets their action.

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?

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 called Nash equilibrium

A Nash equilibrium is a **stable** outcome: one where no agent regrets their action.

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

Which action profiles are pure strategy Nash equilibria in our running examples?

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

Nash Equilibria of Examples

Which action profiles are pure strategy Nash equilibria in our running examples?

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

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

Nash Equilibria of Examples

Which action profiles are pure strategy Nash equilibria in our running examples?

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

	Left	Right
Left	$1, 1$	$-1, -1$
Right	$-1, -1$	$1, 1$

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

Nash Equilibria of Examples

Which action profiles are pure strategy Nash equilibria in our running examples?

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

	Left	Right
Left	$1, 1$	$-1, -1$
Right	$-1, -1$	$1, 1$

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

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

Nash Equilibria of Examples

Which action profiles are pure strategy Nash equilibria in our running examples?

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

	Left	Right
Left	$1, 1$	$-1, -1$
Right	$-1, -1$	$1, 1$

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

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

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

Lecture Overview

Mixed Strategies

So far, we have been assuming that agents play a single action **deterministically**

- But we have seen that that is a pretty bad idea!
- E.g., Matching Pennies, security games

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 $s_i(a_i)$.

Mixed Strategies

So far, we have been assuming that agents play a single action **deterministically**

- But we have seen that that is a pretty bad idea!
- E.g., Matching Pennies, security games

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: $s_i(a_i) = 1$ for some a_i *Only one action played*
- Mixed strategy: $s_i(a_i) < 1$ for all a_i *Randomize over multiple actions*

Mixed Strategies

So far, we have been assuming that agents play a single action **deterministically**

- But we have seen that that is a pretty bad idea!
- E.g., Matching Pennies, security games

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: $s_i(a_i) = 1$ for some a_i *Only one action played*
- Mixed strategy: $s_i(a_i) < 1$ for all a_i *Randomize over multiple actions*
- Set of i's strategies: $S_i = \Delta(A_i)$
- Strategy profiles: $S = S_1 \times \cdots \times S_n$

Utility Under Mixed Strategies

The utility of a mixed strategy profile is its **expected utility** (**why?**)

Utility Under Mixed Strategies

The utility of a mixed strategy profile is its **expected utility** (**why?**)

1. We assume agents are decision theoretically rational
2. We assume that agents randomize **independently**
   $\Delta(A_i) \times \cdots \Delta(A_n)$, **not** $\Delta\left(A_i \times \cdots A_n\right)$

Utility Under Mixed Strategies

The utility of a mixed strategy profile is its **expected utility** (**why?**)

1. We assume agents are decision theoretically rational
2. We assume that agents randomize **independently** $\Delta(A_i) \times \cdots \Delta(A_n)$, **not** $\Delta\left(A_i \times \cdots A_n\right)$

Definition

$$u_i(s) = \sum_{a \in A} \Pr(a \mid s) u_i(a)$$
$$= \sum_{a \in A} \left(\prod_{j \in N} s_j(a_j) \right) u_i(a)$$

Best Response and Nash Equilibrium

Definition

The set of i's best responses to a strategy profile $s_{-i} \in S_{-i}$ is

$$BR_i(s_{-i}) = \{a_i^* \in A_i \mid u_i(a_i^*, s_{-i}) \geq u_i(a_i, s_{-i}) \quad \forall a_i \in A_i\}$$

Best Response and Nash Equilibrium

Definition

The set of i's best responses to a strategy profile $s_{-i} \in S_{-i}$ is

$$BR_i(s_{-i}) = \{a_i^* \in A_i \mid u_i(a_i^*, s_{-i}) \geq u_i(a_i, s_{-i}) \quad \forall a_i \in A_i\}$$

Definition

A strategy profile s is a **Nash equilibrium** iff

$$\forall i \in N, s_i' \in S_i : u_i(s) \geq u_i(s_i', s_{-i})$$

Equivalently,

$$\forall i \in N, a_i \in A_i : s_i(a_i) > 0 \iff a_i \in BR_i(s_{-i}).$$

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.

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 at least one fixed point.
   - A simpletope is a cross product of simplices, so S is a simpletope
2. Construct a continuous function $f : S \to S$ whose fixed points are all Nash equilibria

Interpreting Nash Equilibrium

Question: Is it ever rational for an agent to play any strategy other than a Nash equilibrium strategy?

Interpreting Nash Equilibrium

Question: Is it ever rational for an agent to play any strategy other than a Nash equilibrium strategy? *Yes!*

- Even if the agent is perfectly rational, playing a Nash equilibrium strategy is only optimal if they believe that the other agents will play their parts of the same Nash equilibrium
- Even in a zero-sum game, if you think the other agent will play in a particular sub-optimal way, a non-equilibrium strategy might be the best way to exploit them

Interpreting Nash Equilibrium

Question: Is it ever rational for an agent to play any strategy other than a Nash equilibrium strategy? *Yes!*

- Even if the agent is perfectly rational, playing a Nash equilibrium strategy is only optimal if they believe that the other agents will play their parts of the same Nash equilibrium
- Even in a zero-sum game, if you think the other agent will play in a particular sub-optimal way, a non-equilibrium strategy might be the best way to exploit them

Example

Lisa: Poor, predictable Bart. Always takes Rock.

Bart: Good ol' Rock! Nothing beats Rock!

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., zero-sum games)
- The distribution represents the **other agents' uncertaintly** 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, but individuals are sampled

Summary

Game theory studies **solution concepts** rather than simply optimal behavior

- "Optimal behavior" is not unconditionally defined 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 strategies of the other agents