Game Theory - Repeated Games

Stéphane

today :-)

Outline

Outline

What is a normal form game?

Definition

a n-player game can be represented by a mapping
$R : A_1 \times A_2 \times \ldots \times A_n \mapsto \mathbb{R}^n$

where A_i denotes the discrete set of action available to player i

- $a = (a_1, a_2, \ldots, a_n)$ is the joint action of the players
- $R(a)$ is the payoff for each player ($R_i(a)$ is the payoff of the i^{th} player, i.e. the i^{th} component of $R(a)$)

For a 2-player game, R can be represented by 2 matrices.

What is a strategy?

Definition

A pure strategy is a synonym for an action $a \in A_i$

Definition

A mixed strategy π_i is a probability distribution over the action space A_i

examples

Example (Battle of the sexes)

	D	C
D	2,2	4,3
C	3,4	1,1

Problem: Where to go on a date: Soccer or Opera?

Requirements:

1. avoid to be alone
2. be at the best place

Example (Prisoners' dilemma)

	D	C
D	2,2	4,1
C	1,4	3,3

- Problem: me and my buddy got busted!
- Cooperate: I shut my mouth
- Defect : I blame my buddy

Game Theory is a big field

other concepts

simultaneous or sequential: play simultaneously: each player makes a decision in turn (game tree).

perfect/imperfect information: ability to observe the actions of the opponent(s)

complete/incomplete information: complete information: knowledge of the structure of the games (payoffs matrices).

one stage/multistage game: the outcome of a joint action can be a new game

Game Theory is a big field

other concepts

simultaneous or sequential: play simultaneously: each player makes a decision in turn (game tree).

perfect/imperfect information: ability to observe the actions of the opponent(s)

complete/incomplete information: complete information: knowledge of the structure of the games (payoffs matrices).

one stage/multistage game: the outcome of a joint action can be a new game

Game Theory is a big field

other concepts

simultaneous or sequential: play simultaneously: each player makes a decision in turn (game tree).

perfect/imperfect information: ability to observe the actions of the opponent(s)

complete/incomplete information: complete information: knowledge of the structure of the games (payoffs matrices).

one stage/multistage game: the outcome of a joint action can be a new game

Properties of the payoffs

stochastic game: payoff can be stochastic

Bayesian game: incomplete information game: at the start of the game, some player have private information that others do not(example: bargaining game)

constant/general sum game: for each joint action $a \in \prod_i A_i$, the sum of the payoff $\sum_i R_i(a)$ can be constant. ex: Zero-sum game, purely adverserial game.

Team Game or Cooperative game: all the players receive the same payoff for a joint action.

Dominance

Definition

An outcome X strongly dominates another outcome B if all agents receive a higher utility in X compared to Y.

$$a > b \iff \forall i \in [1..n] R_i(a) > R_i(b)$$

An outcome X weakly dominates (or simply dominates) another outcome B if at least one agent receives a higher utility in X and no agent receives a lesser utility compared to outcome Y.

$$a \geq b \iff \exists j | R_j(a) > R_j(b) \text{ and } \forall i \in [1..n], i \neq j \; R_i(a) \geq R_i(b)$$

Pareto Optimality

Definition

A Pareto optimal outcome is one such that there is no other outcome where some players can increase their payoffs without decreasing the payoff of otherplayers. A non-dominated outcome is Pareto optimal.

Regret

measures how much worse an algorithm performs to the best static strategy.

Definition

the external regret is the difference that a player would receive if it were to play the pure strategy j instead of playing according to π.

Definition

the internal regret is the benefit that player i would get by switching all of its plays of action j to action k instead.

Definition

the total internal (external) regret is the max of the internal (external) regret.

Equilibrium

Definition

An equilibrium is a self-reinforcing distribution over strategy profile.

- Assumption: players are rational (issue with bounded rationality)
- Different natures of equilibrium.

Minimax equilibrium for constant-sum games

minimize the payoff of the opponent: If deviation from equilibrium, the opponent gets an advantage.

Minimax value of a game for player 1

$$\min_{y} \max_{x} R_1(x, y)$$

Properties

- There exists at least one minimax equilibrium in constant sum game.
- set of minimax equilibrium is convex, all have the same value

Nash equilibrium: rationality

mutual best response

if the strategy of the opponent remains fix, the player does not benefit by changing its strategy

Properties

- existence:
  - pure strategy Nash equilibrium may not always exist
  - but there always exists a mixed strategy Nash equilibrium
- complexity to find a Nash equilibrium: there exists exponential time algorithms to compute it, but nobody proved it is NP-Complete.

Examples

Example (Battle of the sexes)

	D	C
D	2,2	4,3
C	3,4	1,1

Example (Prisoners' dilemma)

	D	C
D	2,2	4,1
C	1,4	3,3

Examples

Example (Battle of the sexes)

	D	C
D	2,2	4,3
C	3,4	1,1

Example (Prisoners' dilemma)

	D	C
D	2,2	4,1
C	1,4	3,3

Examples

Example (Battle of the sexes)

	D	C
D	2,2	4,3
C	3,4	1,1

Example (Prisoners' dilemma)

	D	C
D	2,2	4,1
C	1,4	3,3

Examples

Example (Battle of the sexes)

	D	C
D	2,2	4,3
C	3,4	1,1

Example (Prisoners' dilemma)

	D	C
D	2,2	4,1
C	1,4	3,3

Examples

Example (Battle of the sexes)

	D	C
D	2,2	4,3
C	3,4	1,1

Example (Prisoners' dilemma)

	D	C
D	2,2	4,1
C	1,4	3,3

Examples

Example (Battle of the sexes)

	D	C
D	2,2	4,3
C	3,4	1,1

Nash equilibrium D,C and C,D and one mixed strategy $(\frac{3}{4}, \frac{1}{4})$

Pareto Opimal D,C and C,D

Example (Prisoners' dilemma)

	D	C
D	2,2	4,1
C	1,4	3,3

Examples

Example (Battle of the sexes)

	D	C
D	2,2	4,3
C	3,4	1,1

Nash equilibrium D,C and C,D and one mixed strategy $(\frac{3}{4}, \frac{1}{4})$

Pareto Opimal D,C and C,D

Example (Prisoners' dilemma)

	D	C
D	2,2	4,1
C	1,4	3,3

Examples

Example (Battle of the sexes)

	D	C
D	2,2	4,3
C	3,4	1,1

Nash equilibrium D,C and C,D and one mixed strategy $(\frac{3}{4}, \frac{1}{4})$

Pareto Opimal D,C and C,D

Example (Prisoners' dilemma)

	D	C
D	2,2	4,1
C	1,4	3,3

Examples

Example (Battle of the sexes)

	D	C
D	2,2	4,3
C	3,4	1,1

Nash equilibrium D,C and C,D and one mixed strategy $(\frac{3}{4}, \frac{1}{4})$

Pareto Opimal D,C and C,D

Example (Prisoners' dilemma)

	D	C
D	2,2	4,1
C	1,4	3,3

Examples

Example (Battle of the sexes)

	D	C
D	2,2	4,3
C	3,4	1,1

Nash equilibrium D,C and C,D and one mixed strategy $(\frac{3}{4}, \frac{1}{4})$

Pareto Opimal D,C and C,D

Example (Prisoners' dilemma)

	D	C
D	2,2	4,1
C	1,4	3,3

Examples

Example (Battle of the sexes)

	D	C
D	2,2	4,3
C	3,4	1,1

Nash equilibrium D,C and C,D and one mixed strategy $(\frac{3}{4}, \frac{1}{4})$

Pareto Opimal D,C and C,D

Example (Prisoners' dilemma)

	D	C
D	2,2	4,1
C	1,4	3,3

Examples

Example (Battle of the sexes)

	D	C
D	2,2	4,3
C	3,4	1,1

Nash equilibrium D,C and C,D and one mixed strategy $(\frac{3}{4}, \frac{1}{4})$

Pareto Opimal D,C and C,D

Example (Prisoners' dilemma)

	D	C
D	2,2	4,1
C	1,4	3,3

Nash equilibrium (D, D) is the only Nash equilibria of the game.

Pareto Optimal (D, C), (C, D) and (C, C)

N.B. A Nash equilibrium may not be Pareto Optimal

Correlated equilibrium

Example (Battle of the sexes)

	D	C
D	2,2	4,3
C	3,4	1,1

- both agents play mixed strategy $(\frac{1}{2}, \frac{1}{2})$: average payoff is 2.5
- how to avoid bad outcome?

Correlated equilibrium

Example (Battle of the sexes)

	D	C
D	2,2	4,3
C	3,4	1,1

- both agents play mixed strategy $(\frac{1}{2}, \frac{1}{2})$: average payoff is 2.5
- how to avoid bad outcome?

Correlated equilibrium

Players can observe a public random variable and make their decision based on that observation. Player's distribution may no longer be independant. solved by linear program

Examples

Example (Battle of the sexes)

	D	C
D	2,2	4,3
C	3,4	1,1

- flip a (fair?) coin
- head: husband cooperates
- tail: wife cooperates

Example (Traffic light)

- 2 actions Stop or Go
- model the light as being randomly Green or Red. It is the public random variable
- choose life

Outline

1. **Basic Game Theoretic Concept**
   - Basic Concepts
   - Properties
   - Equilibrium concepts

2. Repeated Game

Outline

Repeated Game

Definition

In the repeated game a game M (called stage game) is played over and over again

- one shot game: there is no tomorrow
- repeated game: model a likelyhood of playing the game again with the same opponent
- finitely/infinitely repeated game

Strategy

What is a strategy in a repeated game?

Example

Tit for Tat strategy

- Play the action played by the opponent the last round
- Tit for tat strategy can be an equilibrium strategy in PD or Chicken.

Strategy

What is a strategy in a repeated game?

In the repeated game, a pure strategy depends also on the history of play thus far.

Example

Tit for Tat strategy

- Play the action played by the opponent the last round
- Tit for tat strategy can be an equilibrium strategy in PD or Chicken.

Strategy

What is a strategy in a repeated game?

In the repeated game, a pure strategy depends also on the history of play thus far.

Example

Tit for Tat strategy

- Play the action played by the opponent the last round
- Tit for tat strategy can be an equilibrium strategy in PD or Chicken.

Payoff criterion

Average criterion

Average payoff received throughout the game by player i:

$$\sum_{t=0}^{\infty} M_i(a^t)$$

where a^t is the joint action of iteration t.

Discounted-sum criterion

Discounted sum of the payoff received throughout the game by player i:

$$\sum_{t=0}^{\infty} \gamma^t M_i(a^t)$$

Payoff Space for a two-player game

- $n \times n$ two-player game
- $\mathcal{R}$ and $\mathcal{C}$ are the matrices of the row and column player.
- $\mathcal{V} = \{(\mathcal{R}(i,j), \mathcal{C}(j,i)) | (i,j) \in [1..n]^2\}$
- the payoff space is the Convex Hull $\mathcal{H}$ with vertices in $\mathcal{V}$

Proof.

$\forall (x,y) \in \mathcal{H}, \exists \lambda \in \mathbb{R}^{n^2} \mid x = \sum_{i=1}^{n} \lambda_i \mathcal{R}(i)$ and $y = \sum_{i=1}^{n} \lambda_i \mathcal{C}(j)$ with $\sum_{i=1}^{n} \lambda_i = 1$.
Play the joint action i with the proportion λ_i. □

Example and payoff with independent distribution

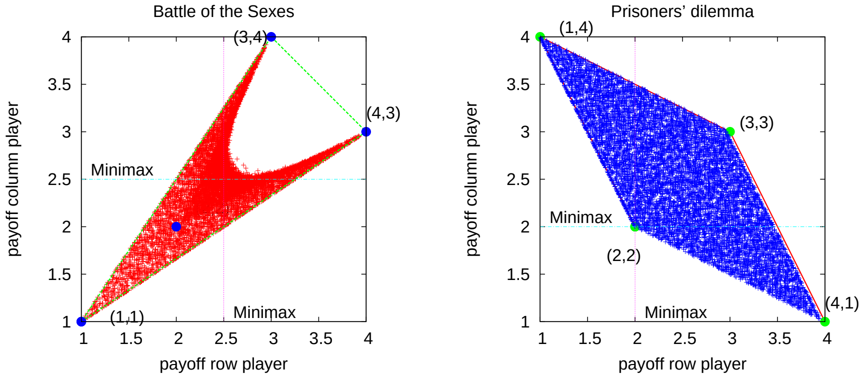

Minimax Value

Feasible region for equilibrium

Minimax value for row and column player:

$$v_r = \min_y \max_x R(x,y)$$

$$v_c = \min_x \max_y C(x,y)$$

The minimax value security value
It defines a feasible region (for an equilibrium)

$$\mathcal{F} = \{(x,y) \in \mathcal{H} | x \geq v_r, y \geq v_c\}.$$

Feasible region for Battle of Sexes and Prisoners' dilemma

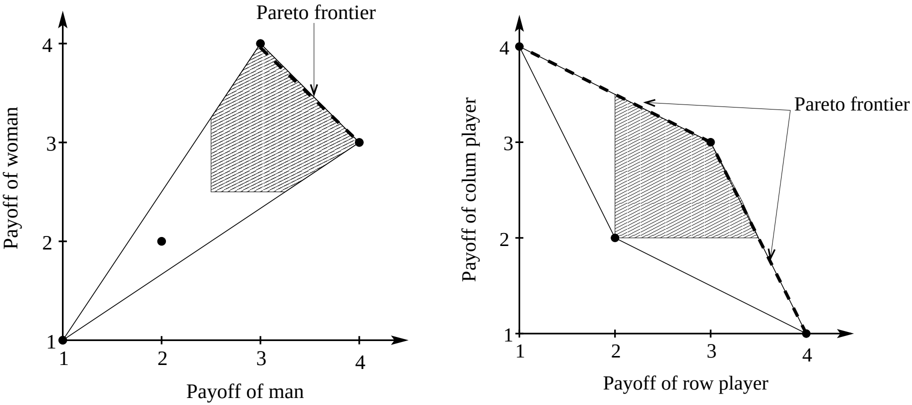

Folk Theorem

Theorem

Any payoff $r \in \mathcal{F}$ can be sustained by a Nash equilibrium.

Proof.

Build strategies that converge to the desired payoff and that make it non-rational to deviate from the strategy. □

Learning in Games

Desirable Properties

Convergence: a learning algorithm should converge

Rationality: play optimally against a stationary opponent

no regret: avoid regrets

Or are they?

Is it possible to find equilibrium that can be good for both players?

Questions