Decision Methods and Models

Master's Degree in Computer Science

Roberto Cordone

DI - Università degli Studi di Milano

Schedule: Tuesday 15.30 - 18.30 in classroom V4

Friday 13.30 - 16.30 in classroom Beta

Office hours: on appointment

E-mail: roberto.cordone@unimi.it

Web page: https://homes.di.unimi.it/cordone/courses/2024-mmd/2024-mmd.html

Ariel site: https://myariel.unimi.it/course/view.php?id=7438

Game theory

We assume

- preference relations Π_d that are weak orders,
  possibly with a known consistent value function $u^{(d)}(f)$
- a certain environment: $|\Omega| = 1 \Rightarrow f(x, \bar{\omega})$ reduces to $f(x)$
- multiple decision-makers: $|D| > 1$

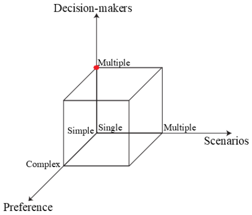

We consider a special class of games with strong properties that relate

- the worst-case strategies
- Nash equilibria

Zero-sum games

In a zero-sum game the payoffs sum to zero in all profile strategies

$$\sum_{d \in D} f^{(d)}(x^{(1)}, \dots, x^{(|D|)}) = 0 \text{ for all } x \in X$$

A uniform sum $\bar{f}$ is reduced to 0 by affine transformation $f' = f - \dfrac{\bar{f}}{|D|}$

They have the same properties

For two-player games, there is a conventional simplified strategic form that reports only the row payoffs

$$\begin{bmatrix} \text{(a,-a)} & \text{(c,-c)} \\ \text{(b,-b)} & \text{(d,-d)} \end{bmatrix} \rightarrow \begin{bmatrix} \text{a} & \text{c} \\ \text{b} & \text{d} \end{bmatrix}$$

Adaptations for the simplified form: dominance

Since the column payoffs are not explicitly provided, one must remember that the column player minimises costs, instead of maximising payoffs

Example

	1	2
1	-2	3
2	-1	1

- there is no explicit dominance on rows
- column 1 dominates column 2 (*smaller costs*!)
- now, row 2 dominates row 1

Strategy profile $(2, 1)$ is the rational choice for both players

Adaptations for the simplified form: dominance

Since the column payoffs are not explicitly provided, one must remember that the column player minimises costs instead of maximising payoffs

Example

	1	2	3
1	2	2	1
2	-1	0	-1
3	3	2	1

- row 3 dominates rows 1 and 2
  (*weakly on some strategies, but strictly overall*)
- column 3 dominates columns 1 and 2 (*smaller costs*!)

Strategy profile $(3, 3)$ is the rational choice for both players

Adaptations for the simplified form: dominance

Since the column payoffs are not explicitly provided, one must remember that the column player minimises costs instead of maximising payoffs

Example

	1	2	3
1	-1	0	-1
2	1	2	3
3	2	-1	-1

- row 2 dominates row 1
- now, column 2 dominates column 3 (*smaller costs*!)

The column domination was hidden in the original matrix

The game is not "solved", but reduced to rows $\{2, 3\}$ and columns $\{1, 2\}$

Adaptations for the simplified form: equilibria

A similar adaptation must be applied to compute Nash equilibria:

$$(r, c) \text{ is an equilibrium } \Leftrightarrow \begin{cases} f_{rc} \geq f_{ic} & \text{for all } i \in X^{(r)} \\ f_{rc} \leq f_{rj} & \text{for all } j \in X^{(c)} \end{cases}$$

Therefore, mark the maxima in each column and the minima in each row: the entries with two marks are equilibria

They are saddle points

Example

	1	2
1	$-2^{\mid}$	3^{-}
2	-1^{+}	1

Entry $(2, 1)$ is an equilibrium

Adaptations for the simplified form: equilibria

A similar adaptation must be applied to compute Nash equilibria:

$$(r, c) \text{ is an equilibrium } \Leftrightarrow \begin{cases} f_{rc} \geq f_{ic} & \text{for all } i \in X^{(r)} \\ f_{rc} \leq f_{rj} & \text{for all } j \in X^{(c)} \end{cases}$$

Therefore, mark the maxima in each column and the minima in each row: the entries with two marks are equilibria

Example

	1	2	3	4
1	2	1^{+}	1^{+}	3
2	-1^{-}	0	-1^{-}	$4^{\vert}$
3	$3^{\vert}$	1^{+}	1^{+}	2
4	-1^{-}	-1^{-}	1	0

There are 4 reciprocally indifferent equilibria: $(1, 2)$, $(1, 4)$, $(3, 2)$, $(3, 4)$

All equilibria have the same payoff value

Is it a general property?

Adaptations for the simplified form: equilibria

A similar adaptation must be applied to compute Nash equilibria:

$$(r, c) \text{ is an equilibrium } \Leftrightarrow \begin{cases} f_{rc} \geq f_{ic} & \text{for all } i \in X^{(r)} \\ f_{rc} \leq f_{rj} & \text{for all } j \in X^{(c)} \end{cases}$$

Therefore, mark the maxima in each column and the minima in each row: the entries with two marks are equilibria

Example

	1	2	3
1	-1^{-}	0	-1^{-}
2	1^{-}	$2^{\mid}$	$3^{\mid}$
3	$2^{\mid}$	-1^{-}	-1^{-}

There is no equilibrium

The value of the game in zero-sum games

The value of the game for the row player is as usual

$$u^{(r)} = \max_{i \in X^{(r)}} \min_{j \in X^{(c)}} f_{ij}$$

For the column player, it derives from the same function f

$$u^{(c)} = \max_{j \in X^{(c)}} \min_{i \in X^{(r)}} -f_{ij} = \max_{j \in X^{(c)}} \left(- \max_{i \in X^{(r)}} f_{ij} \right) = - \min_{j \in X^{(c)}} \max_{i \in X^{(r)}} f_{ij}$$

and we usually redefine it as a cost of the game reversing its sign

$$u^{(c)} = \min_{j \in X^{(c)}} \max_{i \in X^{(r)}} f_{ij}$$

Example

	1	2	3	
1	-1	0	-1	-1
2	1	2	3	1
3	2	-1	-1	-1
	2	2	3	

- $u^{(r)} = \max(-1, 1, -1) = 1$
- $u^{(c)} = \min(2, 2, 3) = 2$

The worst-case strategy in zero-sum games

By definition

$$(r, c) \text{ is an equilibrium } \Leftrightarrow \begin{cases} f_{rc} \geq f_{ic} & \text{for all } i \in X^{(r)} \\ f_{rc} \leq f_{rj} & \text{for all } j \in X^{(c)} \end{cases}$$

Value and cost of a zero-sum game are strongly related

$$\begin{cases} u^{(r)} = \max\limits_{i \in X^{(r)}} \min\limits_{j \in X^{(c)}} f_{ij} \\ u^{(c)} = \min\limits_{j \in X^{(c)}} \max\limits_{i \in X^{(r)}} f_{ij} \end{cases}$$

Both refer to maximum gains and minimum costs

In general games, they refer to different functions

Equilibria and value of zero-sum games

Theorem (Von Neumann and Morgenstern)

For any two-player zero-sum game

- $u^{(r)} \leq u^{(c)}$
- strategy profile (i, j) is an equilibrium if and only if $u^{(r)} = u^{(c)} = f_{ij}$

The guaranteed gain of the row player is limited
by the guaranteed loss of the column player

This sounds natural, because one gains what the other loses

Example

	1	2	3	
1	-1^{-}	0	-1^{-}	-1
2	1^{-}	$2^{\vert}$	$3^{\vert}$	1
3	$2^{\vert}$	-1^{-}	-1^{-}	-1
	2	2	3	

No equilibrium and $u^{(r)} = \max(-1, 1, -1) = 1 < 2 = \max(2, 2, 3) = u^{(c)}$

Equilibria and value of zero-sum games

Theorem (Von Neumann and Morgenstern)

For any two-player zero-sum game

- $u^{(r)} \leq u^{(c)}$
- strategy profile (i, j) is an equilibrium if and only if $u^{(r)} = u^{(c)} = f_{ij}$

Example

	1	2	
1	-2^{-}	$3^{|}$	-2
2	-1^{+}	1	-1
	-1	3	

Entry $(2, 1)$ is an equilibrium and $u^{(r)} = u^{(c)} = -1 = f_{ij}$

Equilibria and value of zero-sum games

Theorem (Von Neumann and Morgenstern)

For any two-player zero-sum game

- $u^{(r)} \leq u^{(c)}$
- strategy profile (i, j) is an equilibrium if and only if $u^{(r)} = u^{(c)} = f_{ij}$

Example

	1	2	3	4		row min
1	2	1^+	1^+	3		1
2	-1^-	0	-1^-	$4^{\vert}$		-1
3	$3^{\vert}$	1^+	1^+	2		1
4	-1^-	-1^-	1	0		-1
	3	1	1	4		

Entries $(1, 2)$, $(1, 3)$, $(3, 2)$ and $(3, 3)$ are equilibria with payoff $f_{ij} = 1$ and $u^{(r)} = \max(1, -1, 1, -1) = 1 = \min(3, 1, 1, 4) = u^{(c)}$

Mixed strategies versus pure strategies

A mixed strategy $\xi^{(d)}$ for player $d \in D$
is a probability vector defined on $X^{(d)}$

$$\xi^{(d)} = \left[\xi_1^{(d)} \ \dots \ \xi_{|X^{(d)}|}^{(d)} \right]^T \in \Xi^{(d)}$$

where

$$\Xi^{(d)} = \left\{ \xi^{(d)} \in [0,1]^{n_d} : \sum_{i \in X^{(d)}} \xi_i^{(d)} = 1 \right\}$$

It can be interpreted in many ways, such as

- in a single game the probability with which d chooses strategy $x^{(d)}$
- in a repeated game the frequency with which d chooses strategy $x^{(d)}$
- in a team game the fraction of players who choose strategy $x^{(d)}$

A pure strategy $x_i^{(d)}$ is the deterministic special case when $\xi^{(d)} \in \{0,1\}^{n_d}$

$$x_i^{(d)} \quad \leftrightarrow \quad \begin{cases} \xi_i^{(d)} = 1 \\ \xi_{i'}^{(d)} = 0 \text{ for all } i' \in X^{(d)} \setminus \{i\} \end{cases}$$

Solving a zero-sum game in mixed strategies

In this situation the payoff obtained becomes a random variable and we can pursue the aim to maximise its expected value

$$\max E\left[f^{(d)}(\xi^{(1)}, \ldots, \xi^{(|D|)})\right]$$

$$\xi^{(d)} \in \Xi^{(d)} \qquad \forall \xi^{(d')} \in \Xi^{(d')} \text{ and } d' \in D \setminus \{d\}$$

How to deal with the other players' strategies $\xi^{(d')}$?

- the worst-case criterium can be reformulated for mixed strategies
- in mixed strategies all zero-sum games have equilibria
- worst-case and equilibria combine nicely: in mixed strategies all zero-sum games have a worst-case equilibrium solution

The expected value of the game

The players aim to optimise the expected value of the game, that is

$$v^{(r)} = \max_{\xi^{(r)} \in \Xi^{(r)}} \min_{\xi^{(c)} \in \Xi^{(c)}} E[f(\xi)]$$

$$v^{(c)} = \min_{\xi^{(c)} \in \Xi^{(c)}} \max_{\xi^{(r)} \in \Xi^{(r)}} E[f(\xi)]$$

where

$$E[f(\xi)] = \sum_{i \in X^{(r)}, j \in X^{(c)}} \xi^{(r)} \xi^{(c)} f_{ij}$$

is the expected gain of the row player and loss of the column player

This looks like a complex problem:

- the row player has infinite strategies to increase the gain
- the column player has infinite strategies to reduce the loss

Who is going to win? Actually, both

This is why mixed strategies are useful in practice

A useful lemma

Lemma

In a two-player game, the worst case for any strategy of a player corresponds to one of the pure strategies of the adversary

- $v^{(r)} = \max_{\xi^{(r)} \in \Xi^{(r)}} \min_{\xi^{c} \in \Xi^{(c)}} E[f(\xi^{(r)}, \xi^{(c)})] = \max_{\xi^{(r)} \in \Xi^{(r)}} \min_{j \in X^{(c)}} E[f(\xi^{(r)}, j)]$
- $v^{(c)} = \min_{\xi^{(c)} \in \Xi^{(c)}} \max_{\xi^{r} \in \Xi^{(r)}} E[f(\xi^{(r)}, \xi^{(c)})] = \min_{\xi^{(c)} \in \Xi^{(c)}} \max_{i \in X^{(r)}} E[f(i, \xi^{(c)})]$

The intuitive idea is to assume that

- the row player adopts a mixed strategy $\xi^{(r)}$; so, $E[f(\xi^{(r)}, j)] = \xi_i^{(r)} f_{ij}$
- one of the column strategies $j \in X^{(c)}$ is the strongest against $\xi^{(r)}$
- then, systematically playing j hits harder than watering it down in a convex combination with weaker strategies

As well, the worst-case for the column player is a pure row strategy

This simplifies the problem, requiring just a parametric expression for

- $\min_{j \in X^{(c)}} E[f(\xi^{(r)}, j)] = \min_{j \in X^{(c)}} \left(\sum_{i \in X^{(r)}} \xi_i^{(r)} f_{ij} \right)$
- $\max_{i \in X^{(r)}} E[f(i, \xi^{(c)})] = \max_{i \in X^{(r)}} \left(\sum_{j \in X^{(c)}} \xi_j^{(c)} f_{ij} \right)$

Example

	1	2
1	2	−3
2	−1	1

The game has

- no dominated strategy
- no Nash equilibrium
- values $u^{(r)} = -1$ and $u^{(c)} = 1$

Assume a general mixed strategy $\xi^{(r)} = (\alpha, 1-\alpha)$ for the row player

Consider the $|X^{(c)}| = 2$ pure strategies available to the column player

1. the expected value of the payoff is

$$E[f(\xi)] = \alpha \cdot 2 + (1-\alpha) \cdot (-1) = 3\alpha - 1$$

2. the expected value of the payoff is

$$E[f(\xi)] = \alpha \cdot (-3) + (1-\alpha) \cdot (1) = 1 - 4\alpha$$

Example

The expected payoff for the pure column strategies is a linear function

1. $j = 1 \Rightarrow E[f(\xi)] = 3\alpha - 1$
2. $j = 2 \Rightarrow E[f(\xi)] = 1 - 4\alpha$

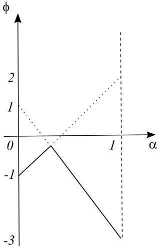

- for mixed strategies it is a linear function in the intermediate region
- the worst case corresponds to the lower envelope (minimum gain)
- it changes for different mixed row strategies α,
  but is always one of the two pure strategies
- the best mixed row strategy α^* can be computed graphically

$$3\alpha^* - 1 = 1 - 4\alpha^* \;\Rightarrow\; \alpha^* = \frac{2}{7} \;\Rightarrow\; E[f(\xi^*)] = -\frac{1}{7}$$

Example

	1	2
1	2	−3
2	−1	1

For mixed column strategy $\xi^{(c)} = (\beta, 1-\beta)$, the pure row strategies give

1. $i = 1 \Rightarrow E[f(\xi)] = \beta \cdot 2 + (1-\beta)\cdot(-3) = 5\beta - 3$
2. $i = 2 \Rightarrow E[f(\xi)] = \beta \cdot (-1) + (1-\beta)\cdot 1 = 1 - 2\beta$

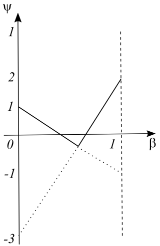

- the worst case corresponds to the upper envelope (maximum loss)
- the best mixed column strategy β^* can be computed graphically

$$5\beta^* - 3 = 1 - 2\beta^* \;\Rightarrow\; \beta^* = \frac{4}{7} \;\Rightarrow\; E[f(\xi^*)] = -\frac{1}{7}$$

Odds and evens

	Odd	Even
Odd	1	-1
Even	-1	1

Once again

- there is no dominated strategy
- there is no Nash equilibrium
- the values of the game are $u^{(r)} = -1$ and $u^{(c)} = 1$ (*useless*)

If the row player applies mixed strategy $\xi^{(r)} = (\alpha, 1-\alpha)$

- when the column player chooses Odd,
  $E[f(\xi)] = \alpha \cdot 1 + (1-\alpha) \cdot (-1) = 2\alpha - 1$
- when the column player chooses Even,
  $E[f(\xi)] = \alpha \cdot (-1) + (1-\alpha) \cdot 1 = 1 - 2\alpha$

If the column player applies mixed strategy $\xi^{(c)} = (\beta, 1-\beta)$

- when the row player chooses Odd,
  $E[f(\xi)] = \beta \cdot 1 + (1-\beta) \cdot (-1) = 2\beta - 1$
- when the row player chooses Even,
  $E[f(\xi)] = \beta \cdot (-1) + (1-\beta) \cdot 1 = 1 - 2\beta$

Odds and evens

Once again, we can solve the game graphically

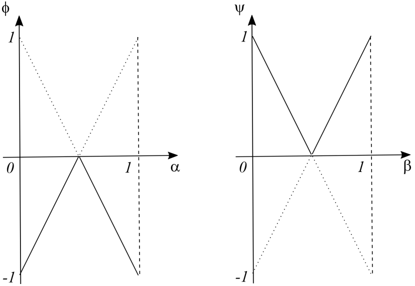

The optimal strategies are

- for the row player:

$$2\alpha^* - 1 = 1 - 2\alpha^* \;\Rightarrow\; \alpha^* = \frac{1}{2} \;\Rightarrow\; E[f(\xi^*)] = 0$$

- for the column player:

$$2\beta^* - 1 = 1 - 2\beta^* \;\Rightarrow\; \beta^* = \frac{1}{2} \;\Rightarrow\; E[f(\xi^*)] = 0$$

The general case

Both examples have two strategies for each player

What happens when the number of pure strategies increases?

- the linear functions depend on more than one probability: the lines become hyperplanes
- the number of linear functions increases: the worst case is still given by their lower or upper envelope

The first modification makes a graphical resolution less viable

Von Neumann and Morgenstern

- proved that the basic properties hold also in the general case
- found a way to compute the optimal strategies

The minimax theorem

Theorem (Von Neumann and Morgenstern)

For any two-player zero-sum game

- $u^{(r)} \leq v^{(r)} = v^{(c)} \leq u^{(c)}$
- at least one mixed strategy profile has this expected value
- this strategy profile is a saddle point with respect to mixed strategies

The guaranteed expected gain of the row player
coincides with the guaranteed expected loss of column

Adopting a different strategy implies the risk of a worse expected value

The minimax theorem

Sketch of the proof
(*exploiting the general theory of duality, that was unknown at the time*)

Given the lemma, the row player aims to optimise the following problem

$$\max \phi(\xi^{(r)}) = \min_{j \in X^{(c)}} \sum_{i \in X^{(r)}} \xi_i^{(r)} f_{ij}$$

$$\sum_{i \in X^{(r)}} \xi_i^{(r)} = 1$$

$$\xi_i^{(r)} \geq 0 \qquad i \in X^{(r)}$$

This can be turned into a Linear Programming problem:

$$\max v$$

$$v \leq \sum_{i \in X^{(r)}} \xi_i^{(r)} f_{ij} \qquad j \in X^{(c)}$$

$$\sum_{i \in X^{(r)}} \xi_i^{(r)} = 1$$

$$\xi_i^{(r)} \geq 0 \qquad i \in X^{(r)}$$

The minimax theorem

Similarly, the column player aims to optimise the following problem

$$
\begin{aligned}
\min \psi(\xi^{(c)}) = \max_{i \in X^{(r)}} \sum_{j \in X^{(c)}} \xi_j^{(c)} f_{ij} \\
\sum_{j \in X^{(c)}} \xi_j^{(c)} = 1 \\
\xi_j^{(c)} \geq 0 \qquad j \in X^{(c)}
\end{aligned}
$$

that can be linearised as

$$
\begin{aligned}
& \min w \\
& w \geq \sum_{j \in X^{(c)}} \xi_j^{(c)} f_{ij} \qquad i \in X^{(r)} \\
& \sum_{j \in X^{(c)}} \xi_j^{(c)} = 1 \\
& \xi_j^{(c)} \geq 0 \qquad j \in X^{(c)}
\end{aligned}
$$

The minimax theorem

$$\begin{array}{rcll} & \max v & & \\ \sum_{i \in X^{(r)}} \xi_i^{(r)} & = & 1 & \\ v - \sum_{i \in X^{(r)}} f_{ij} \xi_i^{(r)} & \leq & 0 & j \in X^{(c)} \\ \xi_i^{(r)} & \geq & 0 & i \in X^{(r)} \end{array} \qquad \begin{array}{rcll} & \min w & & \\ \sum_{j \in X^{(c)}} \xi_j^{(c)} & = & 1 & \\ w - \sum_{j \in X^{(c)}} f_{ij} \xi_j^{(c)} & \geq & 0 & i \in X^{(r)} \\ \xi_j^{(c)} & \geq & 0 & j \in X^{(c)} \end{array}$$

The two problems are reciprocally dual:

- max $\leftrightarrow$ min
- variables $\leftrightarrow$ constraints
  - one free variable $\leftrightarrow$ one equality constraint
  - all but one nonnegative variables $\leftrightarrow$ all but one "natural" constraints
- right-hand-side coefficients $\leftrightarrow$ cost coefficients
- transposed coefficient matrices

Duality implies that the two problems have the same optimal value