Games in normal form

Agents $\xrightarrow[\text{Reasonable assumptions on preferences}]{}$ Utility function u

The agent wants to maximize u. Uncertainty can be accommodated in a sound way: the agent wants to maximize the expected value of u

Our previous example shows how when no other agent is present in the environment selecting the optimal action is quite easy actually!

Agent 1 $\xleftrightarrow{\text{GAME}}$ Agent 2

A FORMAL MODEL DESCRIBING THE INTERACTION BETWEEN UTILITY MAXIMIZING AGENTS

WHAT A GAME DEFINES?

- number and type of players
- for each player, her actions and her preferences
- a correspondence between actions and game outcomes
- sequential structure (if any), recall type, repetitions
- game information:
  - complete vs uncomplete
  - perfect vs imperfect (information set)

this is how to play the game

MECHANISM

(usually this is given)

WHAT A GAME DOES NOT DEFINE?

- description of a player's behavior

(usually this is the game's solution)

this is how players will play in the game

STRATEGIES

(usually this is not given)

WE START FROM THE SIMPLEST FORM OF GAME : STRATEGIC FORM

- aka "normal form" since many other representations can be transformed in this one : it's considered the canonical form

- degenerate sequential structure : players ACT simultaneously, no repetitions

- no uncertainty on players, complete information, deterministic correspondence from actions to outcomes

- finite number of actions to each player

DEF:

- $N = \{ - - \}$ set of players (we assume n players)
- $A_i = \{ \cdots \}$ set of action available to player i
- $A = \bigtimes_{i=1}^{n} A_i$ Joint action space $(a_1, a_2, \dots, a_n) \in A$ Action profile

$$\text{action profiles} \longleftrightarrow \text{outcomes} \quad \left(A = O \right)$$

- $u_i : A \longmapsto \mathbb{R}$ utility function for player i (notice that it depends on A and not on A_i)

LET'S SEE SOME EXAMPLES ...

MATCHING PENNIES

- 2 players, each one has a penny and can choose to display:
  - Head (H)
  - Tails (T)

- one player wins with equal pennies, the other player wins with different ones — w.l.o.g. PLAYER ① / PLAYER ②

P = a positive value

① \ ②	H	T
H	P, −P	−P, P
T	−P, P	P, −P

① is the row player
② is the column player

$$(a_1, a_2) \longrightarrow (\mu_1(a_1,a_2), \mu_2(a_1,a_2))$$

ACTION PROFILE → GAME OUTCOME

example:

$$\mu_1(H, H) = P, \quad \mu_2(H, H) = -P$$

$$\mu_1(T, H) = -P, \quad \mu_2(T, H) = P$$

THIS 2-PLAYER GAME HAS A SPECIAL FEATURE: IF A PLAYER WINS P THEN THE OTHER LOSES THE SAME AMOUNT.
FOR EACH AGENT WINNING ALSO MEANS MAKING THE OTHER LOSING AND VICE VERSA
THE BETTER FOR ME THE WORSE FOR YOU AND VICE VERSA

this is captured by this special regularity we can observe

$$\forall (a_1, a_2) \in A \qquad \mu_1(a_1,a_2) + \mu_2(a_1,a_2) = 0$$

ZERO SUM GAME

ZERO SUM IS A PARTICULAR CASE OF A MORE GENERAL "PURE COMPETITION SCHEME" FOR TWO-PLAYER GAMES

IN INFORMAL TERMS: STRATEGIC SITUATION WITH OPPOSITE OBJECTIVES

If player i receives some gain when the outcome changes from O_1 to O_2 then player $-i$ will receive a loss from the same situation. This must hold $\forall i \in \{1, 2\}$

IN FORMAL TERMS:

take two outcomes l_1 and l_2, then $u_1(l_1) \geq u_1(l_2)$ iff $u_2(l_1) \leq u_2(l_2)$

(it implies that $u_1(l_1) = u_1(l_2)$ then $u_2(l_1) = u_2(l_2)$)

it can be easily shown that

ZERO-SUM GAMES $\subseteq$ STRICTLY COMPETITIVE GAMES

also

if we consider games where (more generally)

$$\forall (a_1, a_2) \in A \qquad u_1(a_1, a_2) + u_2(a_1, a_2) = c$$

c is a constant, we can similarly show that

CONSTANT-SUM GAMES $\subseteq$ STRICTLY COMPETITIVE GAMES

(obviously ZERO-SUM GAMES $\subseteq$ CONSTANT-SUM GAMES)

or if we recall property P, we can also immediately see that actually

ZERO-SUM GAMES = CONSTANT-SUM GAME

any two-player constant sum game can be transformed in an equivalent zero-sum game by normalizing payoffs (positive affine transformation, subtract $\frac{c}{2}$, $a = 1$, $b = -\frac{c}{2}$)

Strictly competitive games are equivalent to zero-sum games

FORMAL DEFINITION:

$(\sigma_1, \sigma_2) \longrightarrow (\sigma_1', \sigma_2')$ then

- NO CHANGE IN PAYOFFS
- THE PAYOFF OF ONE PLAYER INCREASES AND THE ONE OF THE OTHER DECREASES

Let us denote the game as $(A, -B)$ $\quad \sigma_i$ = column vector with player i P.

A = payoff matrix of player 1 $\quad u_1(\sigma) = \sigma_1^T A \sigma_2$

$-B$ = payoff matrix of player 2

Mathematically a game is SC if

$\sigma_1^T A \sigma_2 - \sigma_1'^T A \sigma_2'$ and } player 1 gain in $\sigma \to \sigma'$

$\underbrace{\sigma_1^T B \sigma_2 - \sigma_1'^T A \sigma_2'}_{\text{player 2 gain } (\cdot -1) \text{ in } \sigma \to \sigma'}$ are both null or with the same sign

— weak statement

$O = \{o_1\ o_2 \dots o_m\}$ IF:

for any $p, q \in [0,1]^m$ that are joint prob. distr. on O

WE HAVE THAT

$$\sum_i p_i u_1(o_i) \geq \sum_i q_i u_1(o_i) \text{ iff } \sum_i p_i u_2(o_i) \leq \sum_i q_i u_2(o_i)$$

then it is easy to see that there exists an affine transformation such that $u_1'(o_i) = -u_2'(o_i) \quad \forall o_i \in O$

this statement is "weak" because is using p and q as any $\mathbb{P}$ distribution over the outcomes. We know this is a more general space than distributions that are products of mixed strategies

$\downarrow$

the above argument is not a proof

TO PROVE IT WE NEED THE FOLLOWING THEOREM:

if $\forall\, \sigma, \sigma'$ $\quad \sigma_1^T A \sigma_2 - \sigma_1'^T A \sigma_2$ and $\sigma_1' B \sigma_2' - \sigma_1'^T B \sigma_2'$ have same sign

then B is an affine variant of A $\qquad B = aA + b J_{m,n} \qquad a > 0$

PROOF

Hp: $(A, -B)$ is a strictly competitive game

$a_{ij} = a_{max} \iff b_{ij} = b_{max}$

$a_{ij} = a_{min} \iff b_{ij} = b_{min}$

PROOF: Suppose that $a_{ij} = max$ and $b_{kl} = b_{max} > b_{ij}$, consider these strategy profiles:

$(i,j) \longleftrightarrow (k,l)$ when going $\longrightarrow$ both player loose utility

when going $\longleftarrow$ both players gain utility (recall the we must multiply by -1 the payoff of player 2)

this would contradict our Hp of strictly competitiveness

(Corollary $a_{max} = a_{min} \iff b_{max} = b_{min}$)

$\hookrightarrow$ if we are in this case then clearly B and A are affine variants

Let's see the case where: $a_{max} > a_{min}$ and $b_{max} > b_{min}$

Let's consider the following affine variants of A° and B (we operate a normalization)

J is the $m \times n$ ones matrix ($J_{m,n}$)

$$A' = \frac{1}{a_{max} - a_{min}} \left(A - a_{min} J \right)$$

$$B' = \frac{1}{b_{max} - b_{min}} \left(B - b_{min} J \right)$$

this normalizes payoffs between 0 and 1

we always have a 0 and a 1 in each of A' and B'

Lets suppose to rearrange actions such as

$A'_{11} = B'_{11} = 1 \qquad A'_{22} = B'_{22} = 0$ other configurations can be addressed similarly

then $p \in [0,1]$

$\sigma_1^p = (p, 1-p, 0 \dots 0)^T \qquad D = B' - A'$

$\sigma_2^p = (p, 1-p, 0 \dots 0)^T \qquad |D_{rs}| = \max_{ij} |D_{ij}|$

- easy to see that

$$\boxed{\begin{array}{ll} \sigma_1^{p^T} A' \sigma_2^p = 0 & \text{if } p = 0 \\ \sigma_1^{p^T} A' \sigma_2^p = 1 & \text{if } p = 1 \end{array}}$$

for how we constructed the matrices and the strategy profile

then I can find $\bar{p}$ such that

$$\sigma_1^{\bar{p}^T} A' \sigma_2^{\bar{p}} = A'_{rs}$$

now consider these two strategy profiles

$(r, s) \longleftrightarrow \sigma^{\bar{p}}$

$(A'_{rs}, B'_{rs}) \longleftrightarrow (A'_{rs}, B'_{rs})$ otherwise the game would not be strictly competitive another way to put this:

$$\sigma_1^{\bar{p}^T} B' \sigma_2^{\bar{p}} = B'_{xs}$$

$$\sigma_1^{\bar{p}^T} (A' + D) \sigma_2^{\bar{p}} = B'_{xs}$$

$$\sigma_1^{\bar{p}^T} D \sigma_2^{\bar{p}} = B'_{xs} - A'_{xs} = D_{xs}$$

- $\sigma_1^{\bar{p}^T} D \sigma_2^{\bar{p}}$ is a weighted average over the elements of D
- The first element of D is D_{11} which by construction is $\emptyset$

$$D_{11} = 0$$

- The weight associated with D_{11} is $\bar{p}^2 > 0$ $(D_{xs} \neq 0)$

we cannot select with our weighted average the element of maximum absolute value if we are spending weight for the "zero"

~~there~~ So how come $\sigma_1^{\bar{p}^T} D \sigma_2^{\bar{p}} = D_{xs}$? $\rightarrow D = 0$

$$\Downarrow$$

$$A' = B'$$

Agents A and B play a poker game: one wins the amount that the other loses → ZERO SUM
(sports, robbery, war)

POSITIVE SUM: win-win situations (joint ventures, external funding)

NEGATIVE SUM: lose-lose situations (budget cuts)

Preliminary facts on SC games (2-player as always)

- ANY SC GAME ADMITS A ZERO-SUM REPRESENTATION

In any SC game all the outcomes must lie on a downward slope line

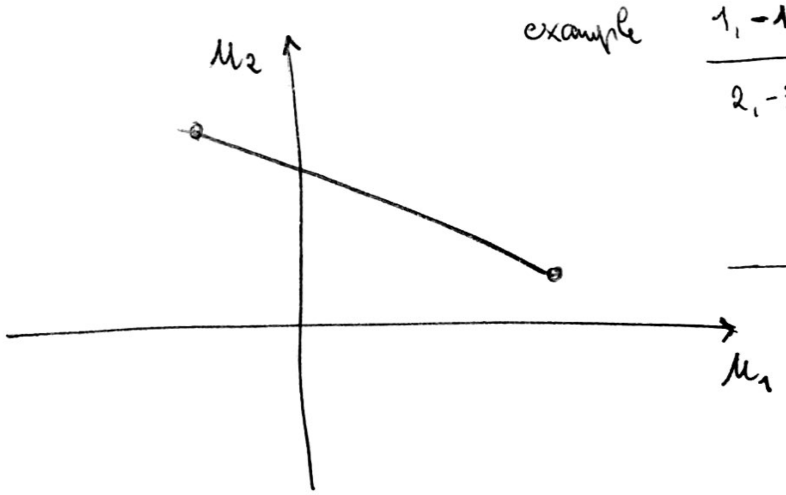

example

1, -1	3, -5
2, -3	1, -1

SC pure and mixed

We can always get this via positive affine transformations

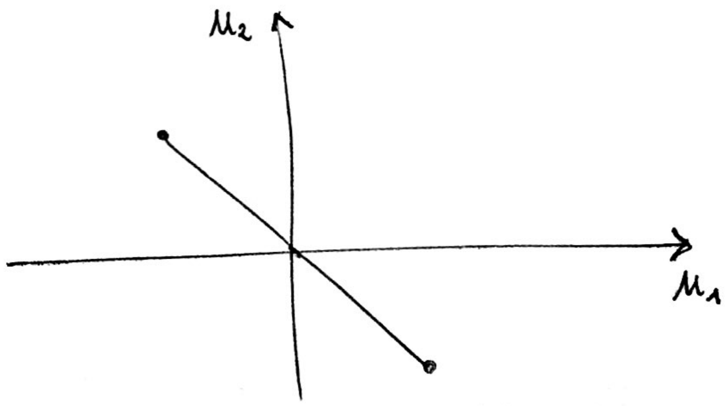

COORDINATION GAME

two mobile robots moving towards each other (they cannot communicate)

		B	
		Left	Right
A	Left	1, 1	0, 0
	Right	0, 0	1, 1

the agents are selfish

the mechanism make selfish behavior also good for everybody

coordination emerges as the way each player will play for their own good

→ NO CONFLICT OF INTERESTS

These type of games is called TEAM GAME

$$\forall a \in A_1 \times A_2 \times \ldots \times A_m \, , \ \forall i,j \in N, \ u_i(a) = u_j(a)$$

GAMES COMBINING BOTH ASPECTS (COORDINATION AND COOPERATION)

Battle of the sexes: Alice and Bob are going out for an evening and they like each other

Bob would like the pub

Alice would like the opera

They also would like to meet

They have to decide independently where to go

		B	
		Opera	Pub
A	Opera	2, 1	0, 0
	Pub	0, 0	1, 2

In this game we have both a team dynamic and a conflict of interest.

What should they decide to play?

To deal with this problem from a formal/computational point of view we have first to formalize their decisions ⟶ STRATEGIES

STRATEGIES

σ_i strategy for player i

$\sigma = (\sigma_1, \sigma_2, \dots, \sigma_n)$ strategy profile

IF σ_i PRESCRIBES TO PLAY A SINGLE ACTION : PURE STRATEGY

IF σ_i PRESCRIBES TO RANDOMIZE OVER TWO OR MORE ACTIONS ; MIXED STRATEGY

formally $0 \leq \sigma_i(a_i) \leq 1 \quad \forall a_i \in A_i$ s.t. $\sum_{a_i \in A_i} \sigma_i(a_i) = 1$

SOME MORE NOTATION $\sigma_{-i} = (\sigma_1, \sigma_2, \dots, \sigma_{i-1}, \sigma_{i+1}, \dots, \sigma_n)$

SUPPORT OF σ_i $\quad S(\sigma_i) = \{ a_i \in A_i \mid \sigma_i(a_i) > 0 \}$

$|S(\sigma_i)| = 1$ Pure strategy

$|S(\sigma_i)| = |A_i|$ Fully mixed strategy

If player i plays strategy σ_i what is her utility? This question in an ill-posed one in game theory (generally speaking)

The well-formed question would be: what's the utility of player i under strategy profile $\sigma = (\sigma_i, \sigma_{-i})$?

In general, σ will make the outcome of the game uncertain, basically inducing a lottery over the outcomes. But our utility-theoretic assumption allow us to correctly define the EXPECTED UTILITY ~~FOR~~ FOR PLAYER i OF σ

$$u_i(\sigma) = \sum_{a \in A} u_i(a) \prod_{j=1}^{n} \sigma_j(a_j)$$

(a_j is the action played by j in action profile a)

In a single-agent decision problem we would try to find σ_i^*

σ_i^* = OPTIMAL STRATEGY FOR AGENT i
(maximizes its expected utility)

however, in game theory to evaluate a candidate strategy σ_i, I need to know what the others intend to do!

I should consider beliefs about the others.

In general beliefs about the others must include their beliefs about me, which in turn must include my beliefs about them and so on.

SOLVING A GAME IS TRYING TO TELL HOW AGENTS WILL PLAY GIVEN THAT THEY ARE MUTUALLY AWARE OF THEIR INTERESTS AND RATIONALITY

↓

IT IS DONE BY DEFINING SUBSET OF OUTCOMES BY MEANS OF SOLUTION CONCEPTS (ways to select groups of outcomes)

A solution concept implicitly relates to some mutual belief scheme which can be said "reasonable" under some terms

Will the game truly end in some of such outcomes?

IT DEPENDS ON HOW MUCH YOU BELIEVE IN HOW THE SOLUTION CONCEPT FILLS THE MEANING OF "reasonable"

↓

That's why we do not have just one possible answer

That's why we do not have general consensus on what solution concept is better than another, but only "popular trends".

We said that a solution concept is a way to select outcomes.
The most trivial way to start with is to use a notion of optimality,

- We need to take the stance of an outsider/external entity
- We must consider the multi-objective nature of the problem

E.G., consider these two outcomes:

$$(3, 10^3) \text{ vs } (4, 1)$$ which one is optimal?

What an external unbiased observer would prefer the game to end?

Obviously we need to talk about PARETO OPTIMALITY / EFFICIENCY:

PARETO DOMINATION: σ^1 Pareto dominates σ^2 if

- $\forall i \in N \quad u_i(\sigma^1) \geq u_i(\sigma^2)$
- $\exists j \in N \quad u_j(\sigma^1) > u_j(\sigma^2)$

Our external observer can, by choosing σ_1 over σ_2, make at least one player better off without having the others complain.

σ is Pareto efficient if $\neg\exists\, \sigma'$ that Pareto dominates it

our solution concept $PE = \{\sigma \mid \sigma \text{ is Pareto efficient}\}$

we can easily verify that

$|PE| \geq 1$, PE contains at least a pure strategy profile

What happens in zero-sum games? (all strategy profiles are in PE) the set of points on PE is also called Pareto Curve

- usually it can be visualized as a (continuous or piecewise) north-east border of the payoff polygon (two-player games)

Let's try to take a look

Coordination game

1, 1	0, 0
0, 0	1, 1

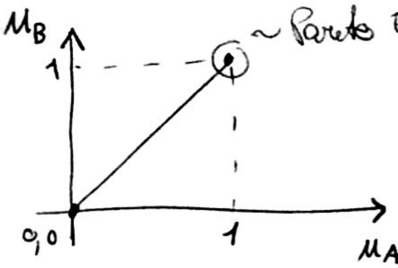

~ Pareto Efficient

in common payoff games all the PE outcomes are with the same payoffs.

zero-sum

-1, 1	1, -1
1, -1	-1, 1

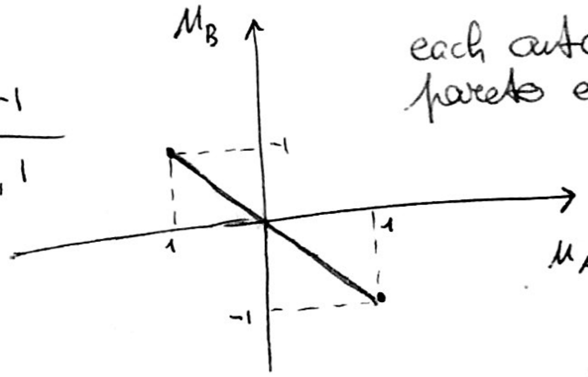

each outcome is pareto efficient!

battle of the sexes

2, 1	0, 0
0, 0	1, 2

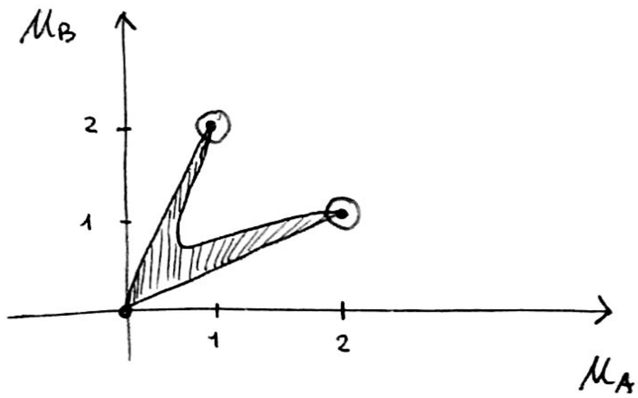

Another example

2, 4	3, 1
2, 2	4, 1

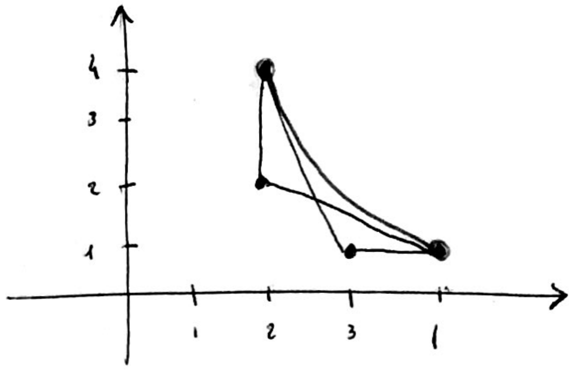

So...

- there is always a pure outcome at least
- computing PE is typically HARD, lots of work in approximating PE
- IN GAMES WITH TWO PLAYERS: EASY

In fact...

A	B
C	D

2 x 2 bimatrix game

PE can be computed exactly in P-time

THERE ARE ONLY ~19 CASES IN WHICH A, B, C, D CAN BE DISPOSED ON THE PLANE