Multiagent Learning and Control

Static games

Maryam Kamgarpour

Professor of Engineering (IGM, STI), EPFL

ME-429, Spring 2025, EPFL

Course topics and dates

1. Static games
2. Zero-sum games
3. Potential games
4. Dynamic games, DP principle
5. Dynamic games, DP examples in games
6. Dynamic games, LQ games/brief Markov game
7. Convex games, Nash equilibria characterization
8. Convex games, Nash equilibria computation
9. Auctions
10. Bayesian games
11. Learning in finite action games
12. Extensive form games
13. Feedback games in extensive form
14. Final project presentations

Motivation - What are multiagent systems?

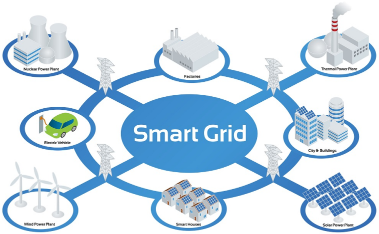

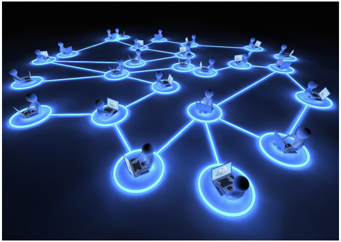

Elements of a Game

Definition

Game theory: mathematical models of decision-making of multiple players

Several things are needed to **characterize a game**:

- The **players** are the agents that make decisions
- The **actions** available to each player at each decision point
- The **information structure** specifies what each player knows before making each decision, in particular with respect to other players' decisions
- The **cost function** for each player, which depends on all players' decisions

Let's play: the Prisoner's dilemma

Two suspects interrogated by police about a crime, in separate rooms.

- If both suspects confess, they each serve 5 years in jail
- If none confesses, police puts each of them in jail for 1 year (evidence for a less serious crime)
- If one confesses, she gets free and the other serves 10 years in jail

Activity - get in groups of three and do the following:

1) Identify the components of the game:
   - players, actions and information of each player, cost for each action
2) play the game as follows:
   - student 1: police, students 2,3: prisoners
   - each prisoner informs the police of their decision privately
   - police informs them about their prison sentence
3) keep playing the game for 5 rounds

Matrix representation of a game

For games with

- two players P_1 and P_2
- a finite number of **actions** (**strategies**)
  $\Gamma = \{\gamma_1, \ldots, \gamma_m\}$, $\Sigma = \{\sigma_1, \ldots, \sigma_n\}$
- simultaneous play. What does each player know?
- costs: $J_1(\gamma_i, \sigma_j)$ and $J_2(\gamma_i, \sigma_j)$ players are minimizing their respective costs

we often adopt the following compact representation:

$$\begin{array}{cc} & \begin{array}{cc} \text{confess} & \text{silent} \end{array} \\ \begin{array}{c} \text{confess} \\ \text{silent} \end{array} & \begin{bmatrix} (5, 5) & (0, 10) \\ (10, 0) & (1, 1) \end{bmatrix} \end{array} \quad \text{or} \quad A = \begin{bmatrix} 5 & 0 \\ 10 & 1 \end{bmatrix} \quad B = \begin{bmatrix} 5 & 10 \\ 0 & 1 \end{bmatrix}$$

- Player 1 is the *row player*, Player 2 is the *column player*
- Each row/column corresponds to a possible action
- Each element of the matrix corresponds to the resulting costs for the two players, as an ordered pair (a_{ij}, b_{ij})

$$a_{ij} = J_1(\gamma_i, \sigma_j), \qquad b_{ij} = J_2(\gamma_i, \sigma_j)$$

Matrix games

The above definition can be extended to N players, by defining the set of actions of each player and the cost function (utilities) of each player.

Matrix games or equivalently *normal form games*, *strategic form games*:

- players act simultaneously
- players don't have knowledge of each other's choice of actions

Strategy: a complete description of how to play the game

- equivalent to control policy/law in control theory or decision rule in optimization
- in matrix games strategy and actions are equivalent, and the terms are used interchangeably (not the case in extensive form games/dynamic games).

Prisoner's dilemma applications

- Prisoner's dilemma has become a model of how individuals cooperate/compete when there is no possibility of binding agreements
- Variants of the game have been used to analyze and understand
  - politics: e.g. international climate negotiations for countries cutting carbon emissions: both countries are better off collectively if they mitigate, but they are individually better off if they pollute.
  - economics: e.g. firms pricing their products
  - biology: e.g. organisms competing for resources
  - COVID-19: e.g. states adopting lock-down

What is a good solution concept for the game?

What is a good model for decision-making in multiagent scenarios where no binding agreement is possible between players?
depends... let's make some assumptions

- each player has common knowledge of all players' cost functions
- rationality: each player wants to minimize her cost (maximizer her profit)
  - the above notion of rationality has been widely accepted
  - in single agent decision-making it implies action γ_1 is preferable to action γ_2 if $J(\gamma_1) \leq J(\gamma_2)$
  - in multi-agent decision-making, not so easy to generalize the above: $J_1(\gamma, \sigma)$ and each player cannot choose the action of other players

Based on the above, let's start by thinking what would the players do/not do...

Strictly dominant and dominated actions

$$A = \begin{bmatrix} a_{11} & a_{12} & \dots & a_{1n} \\ a_{21} & \ddots & & \vdots \\ \vdots & & & \vdots \\ a_{m1} & \dots & \dots & a_{mn} \end{bmatrix}$$

Without loss of generality, lets' consider P_1, choosing rows:

- action i is **strictly dominant action** if

$$a_{ij} < a_{kj}, \quad \forall j \in \{1, \dots, n\}, \forall k \neq i$$

  - in words, row i is element-wise strictly less than the other rows
  - rational player will play strictly dominant action

- action k is **strictly dominated action** if $\exists$ i

$$a_{ij} < a_{kj}, \quad \forall j \in \{1, \dots, n\}$$

  - in words, there is some row i, that is strictly lower than row k element-wise
  - rational player will not play strictly dominated action

Action dominance to determine game outcome

Determine the outcome of the game by identifying strictly dominant actions:

	confess	silent
confess	(5, 5)	(0, 10)
silent	(10, 0)	(1, 1)

Determine the outcome of the game by removing strictly dominated strategies:

	confess	silent	suicide
confess	(5, 5)	(0, 10)	(5, 20)
silent	(10, 0)	(1, 1)	(0, 20)
suicide	(20, 5)	(20, 0)	(20, 20)

Observe:

- rationality leads to players playing strictly dominant strategies (in prisoner's dilemma individually rational behavior don't lead to jointly (socially) optimal decision)
- rationality and common knowledge of the game leads to elimination of strictly dominated strategies
- Next: strictly dominated strategies may not exist

Example: Autonomous cars

- Actions: each car can Remain (r) or Swerve (s)
- Cost function of each player $J_i(.,.)$, for $i = 1, 2$

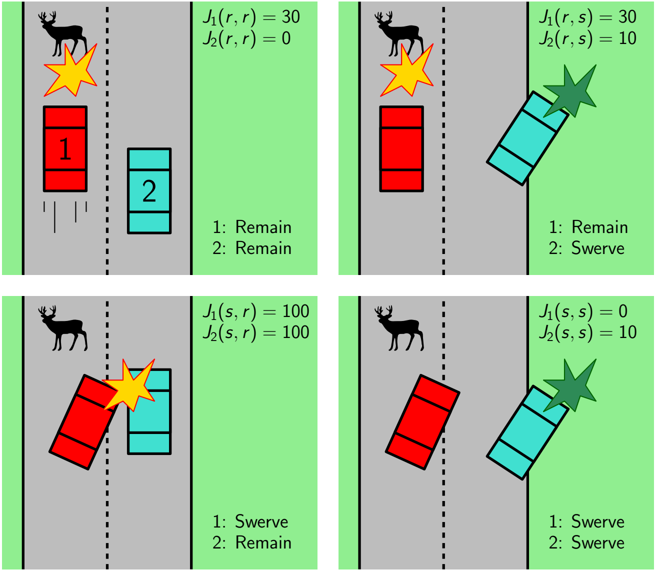

Think/Pair/Share: Action dominance

Are there any dominant or dominated actions for either of the players?

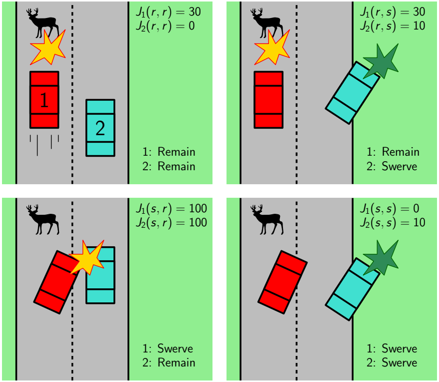

Write the game matrix. Discuss whether dominant/dominated actions exist

	remain	swerve
remain	(30,0)	(30,10)
swerve	(100,100)	(0,10)

there are no strictly dominant/dominated actions

Weakly dominant and dominated actions

$$A = \begin{bmatrix} a_{11} & a_{12} & \dots & a_{1n} \\ a_{21} & \ddots & & \vdots \\ \vdots & & & \vdots \\ a_{m1} & \dots & \dots & a_{mn} \end{bmatrix}$$

Without loss of generality, lets' consider P_1, choosing rows:

- action i is **weakly dominant action** if

$$\begin{aligned} a_{ij} \leq\ & a_{kj}, \quad \forall j \in \{1, \dots, n\}, \forall k \neq i \text{ and} \\ a_{il} <\ & a_{kl}, \quad \text{for some } l \in \{1, \dots, n\}\ \forall k \neq i \end{aligned}$$

  - row i is at least as good as other rows and has at least an element strictly lower than other rows

- action k is **weakly dominated action** if $\exists\ i$

$$\begin{aligned} a_{ij} \leq\ & a_{kj}, \quad \forall j \in \{1, \dots, n\} \text{ and} \\ a_{il} <\ & a_{kl}, \quad \text{for some } l \in \{1, \dots, n\} \end{aligned}$$

  - there is some row i, that is at least as good as row k element-wise, and row i has at least one element that is strictly better (lower) than row k

Iterated elimination of weakly dominated actions

Players are maximizers

$$\begin{array}{c} \\ T \\ M \\ B \end{array} \begin{array}{c} \begin{array}{cc} L & R \end{array} \\ \left[\begin{array}{cc} (1,1) & (0,0) \\ (1,1) & (2,1) \\ (0,0) & (2,1) \end{array} \right] \end{array}$$

for P1 : T is weakly dominated by M
remove

for P2 : L is weakly dominated by R

outcome : (2, 1)

for P1 : B is weakly dominated by M

for P2 : R is weakly dominated by L

outcome : (1, 1)

Observation:

- the order of removal of weakly dominated actions will lead to different outcome

Exercise 1: Iterated elimination of dominated actions

Consider a game where $B = -A$, both players are minimizers, and A is equal to

$$\begin{array}{c} \\ a \\ b \\ c \end{array} \begin{array}{c} \begin{array}{cccc} x & y & z & w \end{array} \\ \begin{bmatrix} -1 & 1 & 3 & 3 \\ 1 & 0 & 2 & 2 \\ 4 & -5 & 3 & 2 \end{bmatrix} \end{array}$$

1. Are there any strictly/weakly dominant actions for either of the players?
2. Are there any strictly/weakly dominated actions for either of the players?
   - If so, remove the dominated action and repeat the above for the reduced game

- Repeat the above exercises for the case in which $B = A$.

Can action dominance lead to prediction of game outcome?

Summary of action dominance

- Rational players will play dominant actions if they exist
  - dominant actions do not usually exist
- Rational players will not play strictly dominated actions
  - can iteratively remove strictly dominated actions to reduce the game
  - the reduced game might still be large
  - strictly dominated actions do not always exist

How else can we define a solution concept for our game?

Security levels and strategies

The **security level** of P_1 is defined by

$$\bar{J}_1 := \min_{i \in \{1,\dots,m\}} \max_{j \in \{1,\dots,n\}} a_{ij}$$

$\bar{J}_1$ minimizes the worst-case cost of P_1

The **security strategy** of P_1 is defined by

$$\bar{i} \in \arg\min_{i \in \{1,\dots,m\}} \max_{j \in \{1,\dots,n\}} a_{ij}$$

Remark

- Security levels and strategies depends only on the players' own payoff matrix.
- Assumes the other player is completely adversarial, which is often not true.
- It is the framework used in robust optimization/robust control.

Security levels and security strategies

Compute the security strategies for

- prisoner's dilemma

$$
\begin{array}{c c}
 & \begin{array}{cc} \text{confess} & \text{silent} \end{array} \\
\begin{array}{c} \text{confess} \\ \text{silent} \end{array} &
\left[\begin{array}{cc} (5, 5) & (0, 10) \\ (10, 0) & (1, 1) \end{array}\right]
\end{array}
$$

$\bar{J}_1 = 5$

$\bar{i}_1 = \{\text{confess}\}$

$\bar{J}_2 = 5$

$\bar{i}_2 = \{\text{confess}\}$

- autonomous cars

$$
\begin{array}{c c}
 & \begin{array}{cc} \text{remain} & \text{swerve} \end{array} \\
\begin{array}{c} \text{remain} \\ \text{swerve} \end{array} &
\left[\begin{array}{cc} (30, 0) & (30, 10) \\ (100, 100) & (0, 10) \end{array}\right]
\end{array}
$$

$\bar{J}_1 = 30$

$\bar{i}_1$: remain

$\bar{J}_2 = 10$

$\bar{i}_2$: swerve

Summary of security strategies

- Always exist. why?
- Easy to compute.. A simple Matlab command
  `barJ1= min(max(A)), barJ2 = min(max(B))`
- Are they consistent with rationality or no regret assumption?

	remain	swerve
remain	(30, 0)	(30, 10)
swerve	(100, 100)	(0, 10)

Nash Equilibrium (informal definition)

- A pair of of actions (γ^*, σ^*) is a Nash Equilibrium if no player can do better by **unilaterally** changing her/his decision.
- Both players show **no regret** after observing the outcome.

	Confess	Stay silent
Confess	$(5, 5)$	$(0, 10)$
Stay silent	$(10, 0)$	$(1, 1)$

What is the Nash equilibrium strategy in prisoner's dilemma?

verify (confess, confess) Nash equilibrium

History

- Nash equilibrium concept proposed in the work of mathematician and philosopher, Antoine Augustin Cournot for analyzing economic firms' competition (we will see Cournot games in Convex game lecture)
- John Nash, mathematician, generalized the concept to any multi-player interaction

Nash equilibrium (formal definition)

Given a static two-player game described by the two payoff matrices A and B (for P_1 and P_2, respectively), we say that the pair of actions $i^* \in \{1, \ldots, m\}$ and $j^* \in \{1, \ldots, n\}$ are a **Nash Equilibrium** if

$$a_{i^*j^*} \leq a_{kj^*} \quad \forall k = 1, \ldots, m$$

and

$$b_{i^*j^*} \leq b_{i^*k} \quad \forall k = 1, \ldots, n.$$

In general terms, $\gamma^* \in \Gamma$ and $\sigma^* \in \Sigma$ are a Nash Equilibrium if

$$J_1(\gamma^*, \sigma^*) \leq J_1(\gamma, \sigma^*) \quad \forall \gamma \in \Gamma$$

and

$$J_2(\gamma^*, \sigma^*) \leq J_2(\gamma^*, \sigma) \quad \forall \sigma \in \Sigma$$

Best-response map

Given a pair of strategies (γ, σ) in $\Gamma \times \Sigma$, we define the **best-response maps**:

$$(\gamma, \sigma) \mapsto R(\gamma, \sigma) = (R_1(\sigma), R_2(\gamma)) \quad \text{where}$$

$$R_1(\sigma) = \{\gamma \in \Gamma \mid J_1(\gamma, \sigma) \leq J_1(\gamma', \sigma),\ \forall \gamma' \in \Gamma\}$$
$$R_2(\gamma) = \{\sigma \in \Sigma \mid J_2(\gamma, \sigma) \leq J_2(\gamma, \sigma'),\ \forall \sigma' \in \Sigma\}$$

- $R(\gamma, \sigma)$ is set-valued: $R : \Gamma \times \Sigma \to 2^{\Gamma \times \Sigma}$
- Example: players are maximizers . .

Nash equilibrium → (4, 3)

$(T, L) \in R(T, L)$

	L	*C*	*R*
T	(4, 3)	(5, 1)	(6, 2)
M	(2, 1)	(9, 4)	(3, 6)
B	(3, 0)	(9, 6)	(2, 8)

P1

$R_1(L) = \{T\}$
$R_1(C) = \{M, B\}$
$R_1(R) = \{T\}$

P2

$R_2(T) = \{L\}$
$R_2(M) = \{R\}$
$R_2(B) = \{R\}$

Nash equilibria are fixed points of best-response maps

- best-response map, $R = (R_1, R_2) : \Gamma \times \Sigma \to 2^{\Gamma \times \Sigma}$

$$R_1(\sigma) = \{\gamma \in \Gamma \mid J_1(\gamma, \sigma) \leq J_1(\gamma', \sigma),\ \forall \gamma' \in \Gamma\}$$
$$R_2(\gamma) = \{\sigma \in \Sigma \mid J_2(\gamma, \sigma) \leq J_2(\gamma, \sigma'),\ \forall \sigma' \in \Sigma\}$$

- $(\gamma, \sigma) \in \Gamma \times \Sigma$ is a fixed point of the best-response map:

$$(\gamma, \sigma) \in R(\gamma, \sigma)$$

- Nash equilibrium $\gamma^* \in \Gamma$ and $\sigma^* \in \Sigma$:

$$J_1(\gamma^*, \sigma^*) \leq J_1(\gamma, \sigma^*) \quad \forall \gamma \in \Gamma$$
$$J_2(\gamma^*, \sigma^*) \leq J_2(\gamma^*, \sigma) \quad \forall \sigma \in \Sigma$$

(γ^*, σ^*) is a Nash equilibrium if & only if it is a fixed point of the best response map

Observations - Nash equilibrium (NE) solution concept

Nash equilibrium rationality properties:

- Each player optimizes her cost given actions of other players
  - strictly dominated strategies cannot be Nash equilibrium (verify this)
  - NE is a subset of set of actions remaining after iterated elimination of strictly dominated strategies (this requires a formal proof..)
- Dominant strategies if they exist are Nash equilibrium
  - NE: optimal action given what others are doing; dominant strategy: optimal action regardless of what others are doing

Can Nash equilibria predict the outcome of a game?

Consider the autonomous car example:

$$A = \begin{bmatrix} 30 & 30 \\ 100 & 0 \end{bmatrix} \qquad B = \begin{bmatrix} 0 & 10 \\ 100 & 10 \end{bmatrix}$$

- Nash equilibria (no regret strategy)
  - (Remain, Remain), outcome (30, 0)
  - (Swerve, Swerve), outcome (0, 10)
- The two Nash equilibria are not **interchangeable**, that is, (Remain, Swerve) and (Swerve, Remain) are not Nash equilibria.

Which strategy will the two cars play?

Exercise 2: the Stag Hunt[1]

Two hunters have to choose their prey between stags and hares.

- If they both go after the same stag (*D* days of food), they will succeed.
- One hunter alone will not be able to hunt the stag
- One hunter alone can hunt a hare ($R < D/2$ days of food).

For the parameterized game, answer the following questions:

1) Are there dominated actions? 2) What are the security strategies for each player? 3) What are the Nash equilibria?

	Stag	Hare
Stag		
Hare		

[1] *old story that appeared in Jean-Jacques Rousseau's Discourse on Inequality, 1775, read more here*

Multiple Nash Equilibria

Multiple NE

Assume a game has multiple Nash equilibria, with different outcomes.
Can we predict what NE strategy will each player play?

Are some Nash equilibria "preferable"?

- **Example 1:** $(1, 2)$ and $(-1, 0)$ players' costs under 2 NE
- **Example 2:** $(1, 0)$ and $(-1, 2)$

Partial order

$$(-1, 0) < (1, 2) \quad \text{but} \quad \begin{array}{l}(1,0) \not< (-1,2)\\(1,0) \not> (-1,2)\end{array}$$

Admissible Nash Equilibria

A Nash equilibrium (γ^*, σ^*) is **admissible** if there is no other Nash equilibrium $(\tilde{\gamma}, \tilde{\sigma})$ such that

$$J_1(\tilde{\gamma}, \tilde{\sigma}) \leq J_1(\gamma^*, \sigma^*) \quad \textbf{and} \quad J_2(\tilde{\gamma}, \tilde{\sigma}) \leq J_2(\gamma^*, \sigma^*)$$

with at least one of the two inequalities strict.

Note: A strategy set (γ^p, σ^p) is Pareto optimal if there is no other *strategy* satisfying the above two inequalities, with one being strict

On existence of admissible Nash equilibria

What are the admissible Nash equilibria in Stag hunt?
Is there an admissible Nash equilibrium in the autonomous cars?

Example: autonomous cars

hard to predict the game's outcome

- **Multiple admissible Nash equilibria** which
  - have **different values**
  - are **not interchangeable**

Few options available:

- Both Players play their **security strategy**
  - **(remain, swerve)**, with outcome $(30, 10)$ (worse than both NE...)
- **Mechanism design**, i.e. change the Players' costs to induce a unique admissible Nash equilibrium
  - example: a **fine** if you cross lanes and hit another car
  - example: sharing the cost of hitting the deer
  - Could either of the above strategies help resolve the non-uniqueness and lead to a "good" Nash equilibrium?

Game: Rock, Paper, Scissors

Pair up with your colleague and play Rock, Paper, Scissors for 10 rounds!

Now, consider only one round of the game.

$A =$

	Rock	Paper	Scissors
Rock	(0,0)	(-1,1)	(1,-1)
Paper	(1,-1)	(0,0)	(-1,1)
Scissors	(-1,1)	(1,-1)	(0,0)

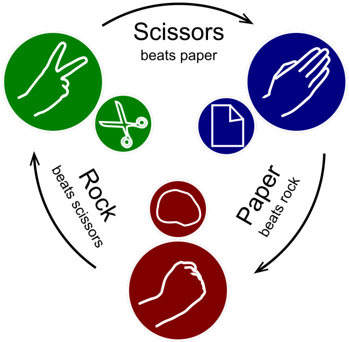

Think/Share/Pair

players are maximizers

- Write the payoff matrix for a single round of the game
  1: P_2 wins, -1: P_1 wins, 0: draw , & vice versa
- Are there dominant / dominated actions? No
- What are the security strategies? any row/column gives security level -1
- Is there a Nash equilibrium? No

Recap

- Nash equilibrium definition and properties
- Nash equilibrium may be
  - unique (prisoner's dilemma)
  - multiple admissible NE (autonomous cars)
  - unique admissible NE (stag-hung) t
  - not exist (rock-paper-scissors)

What is a rational solution concept that can exist in all games considered?

Mixed strategies

- A **mixed strategy** is a probability distribution on the actions

$$\Gamma = \{\gamma_1, \ldots, \gamma_m\} \quad \rightarrow \quad y = \begin{bmatrix} y_1 \\ \vdots \\ y_m \end{bmatrix} \qquad \Sigma = \{\sigma_1, \ldots, \sigma_n\} \quad \rightarrow \quad z = \begin{bmatrix} z_1 \\ \vdots \\ z_n \end{bmatrix}$$

$$y \in \mathcal{Y} = \left\{ y \mid \sum_i y_i = 1, y_i \geq 0 \right\} \qquad z \in \mathcal{Z} = \left\{ z \mid \sum_i z_i = 1, z_i \geq 0 \right\}$$

Pure strategies still exist within the mixed strategy space.

- Example

$$\text{pure strategy } \gamma_i \qquad \equiv \qquad y = \begin{bmatrix} 0 \\ \vdots \\ 0 \\ y_i = 1 \\ 0 \\ \vdots \\ 0 \end{bmatrix} \in \mathcal{Y}.$$

Expected outcome

Expected outcome for player 1

$$J(y,z) = \sum_{i,j} a_{ij}\mathbb{P}[\gamma_i, \sigma_j] = \sum_{i,j} a_{ij}\mathbb{P}[\gamma_i]\mathbb{P}[\sigma_j] = \sum_{i,j} a_{ij} y_i z_j = y^\top A z$$

Similarly, defined for player 2. There is a slight abuse of notation. Do you see it?

- **expected outcome** of the game

$$J_1(y,z) = \sum_{i=1}^{m}\sum_{j=1}^{n} y_i z_j a_{ij} = y^\top A z$$

$$J_2(y,z) = \sum_{i=1}^{m}\sum_{j=1}^{n} y_i z_j b_{ij} = y^\top B z$$

Multiple interpretations:

- repeated game: think of rock-paper-scissors
- mixture of pure actions (if the game allows that)

Exercise 3: Properties of mixed strategies and payoffs

- Let $\mathcal{X} \subset \mathbb{R}^n$. Write the definition of X being convex.

$\forall x_1, x_2 \in X, \quad \lambda \in [0,1] \qquad \lambda x_1 + (1-\lambda) x_2 \in X$

- Show that the set of mixed strategies $\mathcal{Y}$ is a convex sets (clearly, same goes for $\mathcal{Z}$).
- Let $f : \mathcal{X} \to \mathbb{R}$. Write the definition of f being convex.

$\forall x_1, x_2 \in X, \quad \lambda \in [0,1] \quad f(\lambda x_1 + (1-\lambda) x_2) \leq \lambda f(x_1) + (1-\lambda) f(x_2)$

- Show that $J_1(y, z) = y^\top Az$ is linear in y for each fixed z and hence, it is convex in y for each fixed z.

Mixed Nash equilibrium

Given a static two-player game described by the two payoff matrices A and B (for P_1 and P_2, respectively), a pair of mixed strategies $y^* \in \mathcal{Y}$ and $z^* \in \mathcal{Z}$ is a **mixed Nash Equilibrium** if

$$(y^*)^\top A z^* \leq y^\top A z^* \quad \forall y \in \mathcal{Y}$$

and

$$(y^*)^\top B z^* \leq (y^*)^\top B z \quad \forall z \in \mathcal{Z}$$

In general terms, $y^* \in \mathcal{Y}$ and $z^* \in \mathcal{Z}$ are a Nash Equilibrium if

$$J_1(y^*, z^*) \leq J_1(y, z^*) \quad \forall y \in \mathcal{Y}$$

and

$$J_2(y^*, z^*) \leq J_2(y^*, z) \quad \forall z \in \mathcal{Z}$$

Remark

Not easy to compute - even in 2-player case!

Certifying Mixed Nash Equilibria

To **certify** that a given pair of mixed strategies is a mixed Nash Equilibrium, we can perform a finite number of checks.

Proposition

a a ⟺ b

A pair of strategies (y^*, z^*) is a Nash equilibrium if and only if

$$J_1(y^*, z^*) \le J_1(y, z^*) \quad \forall \text{ pure strategy } y$$
$$J_2(y^*, z^*) \le J_2(y^*, z) \quad \forall \text{ pure strategy } z$$

b

Proof.

$\Rightarrow$ follows directly from the definition of Nash equilibrium. a ⟹ b

$\Leftarrow$ b. ⟹ a

from b: $p^* = J_1(y^*, z^*) \le e_i^T A z^*$, $e_i = \begin{bmatrix} 0 \\ \vdots \\ 1 \\ 0 \\ 0 \end{bmatrix}$ → i-th entry $\forall i = 1, \ldots, n$

if we consider any $y \in \mathcal{Y}$

$y_i p^* \le y_i e_i^T A z^*$

$p^* = \sum_{i=1}^{n} y_i p^* \le \sum y_i e_i^T A z^* = y^T A z^*$

□

How to compute mixed Nash equilibria

intractable: non-convex problem even for 2 player games
See Lecture 10 in Hepsanha's book. However, in certain cases, we can solve a system of linear equations for computing Nash equilibria:

Lemma - on characterizing completely mixed Nash equilibria

1. If y^* and z^* are **completely mixed Nash Equilibria** (i.e., none of their elements are zero), then they **need** to satisfy

   all elements strictly positive

$$\begin{cases} Az^* = p^* \mathbf{1} & \in \mathbb{R}^m \\ (y^*)^\top B = q^* \mathbf{1}^\top & \in \mathbb{R}^n \\ \mathbf{1}^\top y^* = 1, & \\ \mathbf{1}^\top z^* = 1. & \end{cases}$$

   probability distribution

2. Suppose that y^*, z^* satisfy the above linear system of equations. If $y_i^* \geq 0, i = 1, \ldots, n$, $z_j^* \geq 0, j = 1, \ldots, m$, then (y^*, z^*) is a mixed Nash equilibrium.

Completely mixed Nash equilibria proof

We will prove part 1 of the Lemma above. Part 2's proof can be found in Theorem 10.1 in Hespnaha's book.

Proof of Lemma, part 1: Let (y^*, z^*) be a completely mixed Nash equilibrium. Then,

$$(y^*)^\top A z^* = \min_{y \in \mathcal{Y}} y^\top A z^* = \min_{y} \sum_i y_i \underbrace{(Az^*)_i}_{i\text{-th entry of } Az^* \in \mathbb{R}^n} .$$

Suppose one entry of the vector Az^*, call it i is strictly larger than any other one. Then, clearly, we would have to choose $y_i = 0$ in the minimum above. This implies that the Nash equilibrium will not be completely mixed.

Hence, to have a completely mixed Nash equilibrium, all entries must have the same cost, or equivalently, there must exist a $p^* \in \mathbb{R}$ such that $Az^* = p^* 1_n$, where $1_n \in \mathbb{R}^n$ is vector of ones. Same reasoning applies to $(y^*)^\top B$.

Exercise 4: Computing mixed Nash equilibria in matching pennies

Consider the matching pennies game[2]

	Left	Right
Left	$(1, -1)$	$(-1, 1)$
Right	$(-1, 1)$	$(1, -1)$

Compute the mixed Nash equilibria.

$$A = \begin{bmatrix} 1 & -1 \\ -1 & 1 \end{bmatrix} \qquad B = \begin{bmatrix} -1 & 1 \\ 1 & -1 \end{bmatrix}$$

$$z^* = \begin{bmatrix} z_1^* \\ z_2^* \end{bmatrix}, \quad z_1^* + z_2^* = 1$$

$$y^* = \begin{bmatrix} y_1^* \\ y_2^* \end{bmatrix}, \quad y_1^* + y_2^* = 1$$

$$A z^* = p^* \mathbf{1}$$

$$\begin{bmatrix} 1 & -1 \\ -1 & 1 \end{bmatrix} \begin{bmatrix} z_1^* \\ z_2^* \end{bmatrix} = \begin{bmatrix} p^* \\ p^* \end{bmatrix}$$

$$\begin{bmatrix} y_1^* & y_2^* \end{bmatrix} \begin{bmatrix} -1 & 1 \\ 1 & -1 \end{bmatrix} = \begin{bmatrix} q^* & q^* \end{bmatrix}$$

$$z_1^* = z_2^* = 1/2, \quad y_1^* = y_2^* = 1/2, \quad p^* = q^* = 0$$

[2] This is also known as penalty kick. The penalty kick would have slightly different numbers for each action based on actual penalty kick data! see Preface of Karlin & Peres book.

Exercise 5: Computing Nash equilibria in a coordination game

A couple agreed to meet after work but hasn't decided where. The woman (Player 1) likes to see a soccer match and the man likes to see a ballet. They both prefer to attend the same event to not. The rewards are given below:

	soccer	ballet
soccer	$(3, 2)$	$(1, 1)$
ballet	$(0, 0)$	$(2, 3)$

- Identify the pure Nash equilibria. Are they interchangeable? Is there an admissible Nash equilibrium?

- Find the mixed Nash equilibrium. What is the probability of the couple attending the same event?

Existence of a Nash Equilibrium

John Nash's theorem, 1949

A mixed strategy Nash equilibrium exists in every finite action game.

The proof uses a fixed point theorem: let $\mathcal{X} \subset \mathbb{R}^n$ non-empty, convex compact.

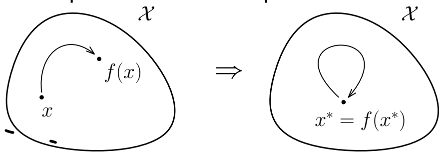

Brouwer's Fixed Point theorem (1910)
Let $f : \mathcal{X} \to \mathcal{X}$ be continuous. Then, f has a fixed point.

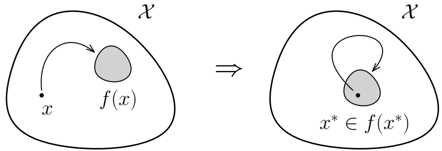

Kakutani's Fixed Point theorem (1941)
Let $R : \mathcal{X} \to 2^{\mathcal{X}}$ be a set-valued map with a closed graph such that $\forall x \in \mathcal{X}$, $R(x)$ is convex and non-empty. Then, R has a fixed point.

$((y, z), R(y, z))$

Closed sets, bounded sets, compact sets

A subset $\mathcal{X} \subset \mathbb{R}^n$ is compact if and only if it is closed and bounded.

- Write the definition of a closed set.

a set is closed if it contains all its limit points

- Write the definition of a bounded set. $X \subseteq \mathbb{R}^n$ is bounded

if $\exists M > 0$ s.t. $\|x\| \leq M \quad \forall x \in X$

- Show that $\mathcal{Y}$, $\mathcal{Z}$ are compact.

↳ any norm in $\mathbb{R}^n$

↳ show closed & bounded

Properties of the best-response map for mixed strategies

See details in Fudenberg & Tirole, Theorem 1.1 in Section 1.3.1

Let $R : \mathcal{Y} \times \mathcal{Z} \to 2^{\mathcal{Y} \times \mathcal{Z}}$ be the response map corresponding to mixed strategies.

- Why is $R(y, z)$ non-empty?

  P1 $R_1(z) = \operatorname{argmin}_{y \in \mathcal{Y}} y^T A z$

  min of a continuous function over a compact set ⇒ minimum is achieved

- Show $R(y, z)$ is convex for every $(y, z) \in \mathcal{Y} \times \mathcal{Z}$.

  $(y_1, z_1), (y_2, z_2), \lambda \in [0,1]$ show

  $R(\lambda(y_1, z_1) + (1-\lambda)(y_2, z_2)) \leq \lambda R(y_1, z_1) + (1-\lambda) R(y_2, z_2)$

- Verify that $R(y, z)$ has a closed graph. let $X = \mathcal{Y} \times \mathcal{Z}$,

Graph of $R(\cdot)$; $G = \{ (x, y) \in X \times X \text{ s.t. } y \in R(X) \}$

show if $\exists$ sequence $\{(x_i, y_i)\}_i \in G$ that converges to a point $(\bar{x}, \bar{y})$, then $(\bar{x}, \bar{y}) \in G$

Fixed point of the best-response map

We can apply **Kakutani's fixed point theorem** and conclude that the best-response map

$$(y, z) \mapsto R(y, z) = (R_1(z), R_2(y))$$

where R_1, R_2 are player 1 and 2's best-response maps in mixed strategies, has a fixed point.

By definition of $R(y, z)$, its fixed points correspond to mixed Nash Equilibria. □

N players

The proof can be repeated almost identically for N player games, by extending the corresponding definitions of mixed Nash Equilibrium.

John Nash's Ph.D. thesis

In 1949, John Nash provided a one-page proof that games with any number of players have a mixed "Nash equilibrium" in his 27-page Ph.D. thesis. He used Brouwer's fixed point theorem in his thesis and Kakutani's fixed point theorem in his published paper in 1950. The approach using Kakutani's theorem is more general as we will see in the lectures on convex games.

In 1994, John Nash was awarded the Nobel Prize for this pioneering work.

Table of Contents

Existence of Equilibrium Points

I have previously published [Proc. N. A. S. 36 (1950) 48-49] a proof of the result below based on Kakutani's generalized fixed point theorem. The proof given here uses the Brouwer theorem.

The method is to set up a sequence of continuous mappings:

$$\mathfrak{s} \to \mathfrak{s}'(\mathfrak{s},1);\ \mathfrak{s} \to \mathfrak{s}'(\mathfrak{s},2);\ \cdots$$

whose fixed points have an equilibrium point as limit point. A limit mapping exists, but is discontinuous, and need not have any fixed points.

THEO. 1: Every finite game has an equilibrium point.

Proof: Using our standard notation, let $\mathfrak{s}$ be an n-tuple of mixed strategies, and $p_{i\alpha}(\mathfrak{s})$ the pay-off to player i if he uses his pure strategy $\pi_{i\alpha}$ and the others use their respective mixed strategies in $\mathfrak{s}$. For each integer λ we define the following continuous functions of $\mathfrak{s}$:

$$q_i(\mathfrak{s}) = \max_\alpha p_{i\alpha}(\mathfrak{s}),$$

$$\phi_{i\alpha}(\mathfrak{s},\lambda) = p_{i\alpha}(\mathfrak{s}) - q_i(\mathfrak{s}) + 1/\lambda, \text{ and}$$

$$\phi^+_{i\alpha}(\mathfrak{s},\lambda) = \max[0, \phi_{i\alpha}(\mathfrak{s},\lambda)].$$

Now $\sum_\alpha \phi^+_{i\alpha}(\mathfrak{s},\lambda) \geq \max_\alpha \phi^+_{i\alpha}(\mathfrak{s},\lambda) = 1/\lambda > 0$ so that

$$c'_{i\alpha}(\mathfrak{s},\lambda) \doteq \frac{\phi^+_{i\alpha}(\mathfrak{s},\lambda)}{\sum_\beta \phi^+_{i\beta}(\mathfrak{s},\lambda)}$$

is continuous.

Define $s'_i(\mathfrak{s},\lambda) = \sum_\alpha \pi_{i\alpha} c'_{i\alpha}(\mathfrak{s},\lambda)$ and $\mathfrak{s}'(\mathfrak{s},\lambda) = (s'_1, s'_2, \dots s'_n)$. Since all the operations have preserved continuity, the mapping $\mathfrak{s} \to \mathfrak{s}'(\mathfrak{s},\lambda)$ is con-

-6-

tinuous; and since the space of n-tuples, $\mathfrak{s}$, is a cell, there must be a fixed point for each λ. Hence there will be a subsequence $\mathfrak{s}_\mu$, converging to $\mathfrak{s}^*$, where $\mathfrak{s}_\mu$ is fixed under the mapping $\mathfrak{s} \to \mathfrak{s}'(\mathfrak{s},\lambda(\mu))$.

Now suppose $\mathfrak{s}^*$ were not an equilibrium point. Then if $\mathfrak{s}^* = (s_1^*, \dots s_n^*)$ some component s_i^* must be non-optimal against the others, which means s_i^* uses some pure strategy $\pi_{i\alpha}$ which is non-optimal. [see (4), pg. 4] This means that

$$p_{i\alpha}(\mathfrak{s}^*) < q_i(\mathfrak{s}^*)$$

which justifies writing

$$p_{i\alpha}(\mathfrak{s}^*) - q_i(\mathfrak{s}^*) < -\epsilon.$$

From continuity, if μ is large enough,

$$\left|[p_{i\alpha}(\mathfrak{s}_\mu) - q_i(\mathfrak{s}_\mu)] - [p_{i\alpha}(\mathfrak{s}^*) - q_i(\mathfrak{s}^*)]\right| < \epsilon/2 \text{ and } 1/\lambda(\mu) < \epsilon/2.$$

Adding, $p_{i\alpha}(\mathfrak{s}_\mu) - q_i(\mathfrak{s}_\mu) + 1/\lambda(\mu) < 0$ which is simply

$$\phi_{i\alpha}(\mathfrak{s}_\mu, \lambda(\mu)) < 0, \text{ whence } \phi^+_{i\alpha}(\mathfrak{s}_\mu,\lambda(\mu)) = 0, \text{ whence}$$

$$c'_{i\alpha}(\mathfrak{s}_\mu, \lambda(\mu)) = 0.$$

From this last equation we know that $\pi_{i\alpha}$ is not used in $\mathfrak{s}_\mu$ since

$$\mathfrak{s}_\mu = \sum_\alpha \pi_{i\alpha} c'_{i\alpha}(\mathfrak{s}_\mu, \lambda(\mu))$$

, because $\mathfrak{s}_\mu$ is a fixed point.

And since $\mathfrak{s}_\mu \to \mathfrak{s}^*$, $\pi_{i\alpha}$ is not used in $\mathfrak{s}^*$, which contradicts our assumption.

Hence $\mathfrak{s}^*$ is indeed an equilibrium point.

Summary

- Key elements of a game, and how to recognize them
- How to formalize static games in matrix form
- Strategy dominance
- (Pure) security levels and strategies
- (Pure) Nash equilibria
- Multiple Nash equilibria
- Admissible Nash equilibria
- Mixed strategies
- Mixed Nash equilibria
- Certifying and computing mixed Nash equilibria
- Nash's Theorem: existence of mixed Nash equilibria

Further readings

To deepen your understanding:

- Chapters, 1 and 10, *Hespanha, Noncooperative Game Theory*
- Chapters 4 and 5, *Karlin and Peres, Game Theory, Alive*
- Chapter 1, *Game theory, Fudenberg & Tirole*