Multiagent decision-making and control

Randomized feedback games

Maryam Kamgarpour

Professor of Engineering (IGM, STI), EPFL

ME-429, Spring 2025, EPFL

Course topics

1. Static games
2. Zero-sum games
3. Potential games
4. Extensive form games
5. Dynamic games, dynamic programming principle
6. Dynamic games, dynamic programming for games
7. Dynamic games, linear quadratic games, Markov games
8. Convex games, Nash equilibria characterization
9. Convex games, Nash equilibria computation
10. Auctions
11. Bayesian games
12. Learning in games
13. Final project presentations

Randomized feedback games

A multi-stage game in extensive form is a **feedback game** if

1. no information set spans over multiple stages
2. each "Player 1" node is the root of a separated sub-game.

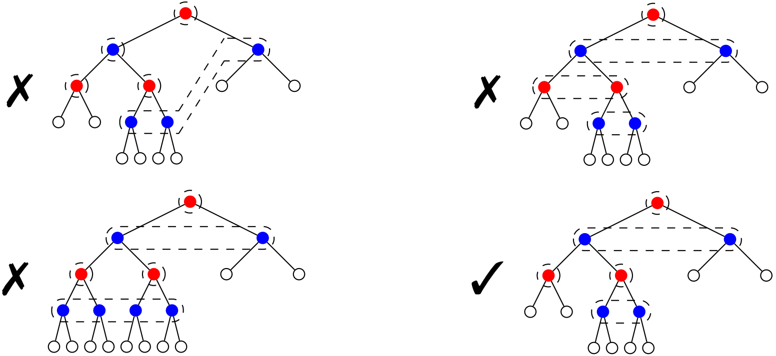

Feedback games

A multi-stage game in extensive form is a **feedback game** if

1. no information set spans over multiple stages
   - ***Each player knows the current stage of the game***
2. each “Player 1” node is the root of a separated sub-game.
   - ***Both player know what happened in the previous stages***

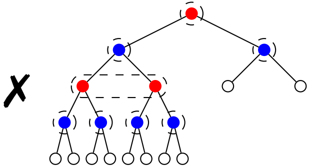

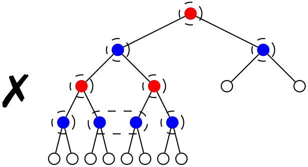

Perfect recall (information): Each player never forgets what he/she knows.

Towards randomized strategies

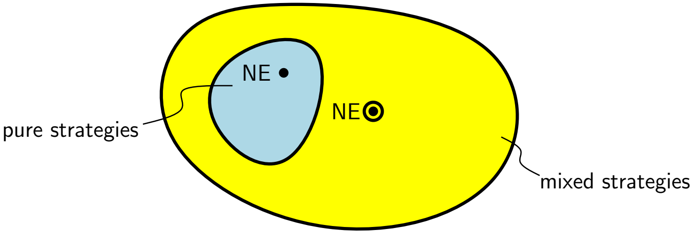

For games in matrix form, we **expanded** the set of strategies to include **mixed strategies**. Why?

Towards randomized strategies

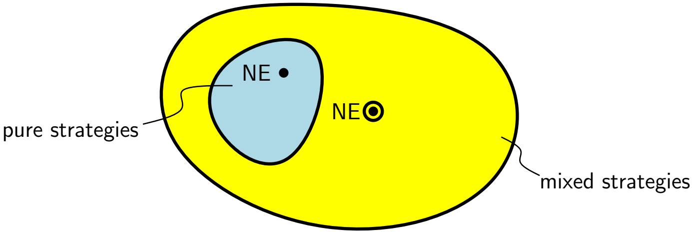

For games in matrix form, we **expanded** the set of strategies to include **mixed strategies**. Why?

- **Guarantees of existence of NE for any finite action games** (the set is large enough!)
  Minmax Theorem in zero-sum games
- **Computational tools**
  Linear Programming in zero-sum games or for completely mixed strategies, non-convex quadratic programming in general games

Towards randomized strategies

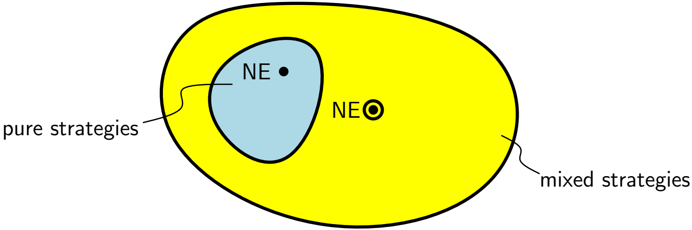

We could do the same for games in extensive form.
(Remember: they can all be converted in matrix form)

the set of **mixed strategies** in extensive games

- is computationally intractable
- is often unnecessarily large

Towards randomized strategies

(in feedback games)

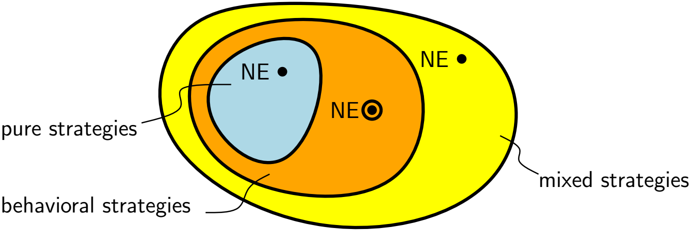

- We restrict our attention to the special class we considered before: **feedback games**
- We define **behavioral strategies**, a special class of randomized strategies
- For **feedback games**, the subset of **behavioral strategies** is
  - computationally tractable
  - large enough to contain a NE – no need for mixed strategies

Mixed strategies

Remember: A **pure strategy** γ **(or** σ is **a map** that associates an action to each information set.

$$\gamma : \{\mathcal{I}_1, \dots, \mathcal{I}_r\} \to \bigcup_i \mathcal{U}_i \qquad \mathcal{I}_i \mapsto \gamma(\mathcal{I}_i) \in \mathcal{U}_i$$

$$\sigma : \{\mathcal{J}_1, \dots, \mathcal{J}_s\} \to \bigcup_i \mathcal{V}_i \qquad \mathcal{J}_i \mapsto \sigma(\mathcal{J}_i) \in \mathcal{V}_i$$

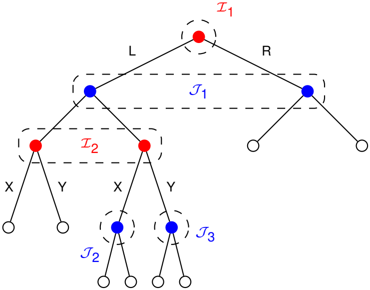

Example: Player 1

A strategy γ maps each **information set** into an **action**

$$\gamma(\mathcal{I}_1) \in \{L, R\}$$

$$\gamma(\mathcal{I}_2) \in \{X, Y\}$$

Mixed strategies

Consider the **set of pure strategies** for a player 1

$$\Gamma = \{\gamma_1, \gamma_2, \ldots, \gamma_n\}$$

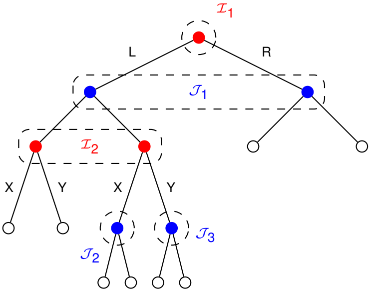

$$\Gamma = \{\gamma_1, \gamma_2, \gamma_3, \gamma_4\}$$
$$= \{\text{LX}, \text{LY}, \text{RX}, \text{RY}\}$$

where “LX” means

$$\gamma_1(\mathcal{I}_1) = L, \quad \gamma_1(\mathcal{I}_2) = X$$

Mixed strategies

Consider the **set of pure strategies** for a player i:

$$\Gamma = \{\gamma_1, \gamma_2, \ldots, \gamma_n\}$$

Mixed strategy

A mixed strategy $y \in \mathbb{R}^n$ for Player 1 (and equivalently $z \in \mathbb{R}^m$ for Player 2) corresponds to randomly selecting a pure strategy from the set of pure strategies Γ, according to the probabilities

$$y_1, \ldots, y_n$$

where

$$y_i \geq 0 \;\; \forall i, \quad \text{and} \quad \sum_i y_i = 1$$

Exactly the same definition as in games in matrix form
(remember the **probability simplices** $\mathcal{Y}$ and $\mathcal{Z}$)

Example

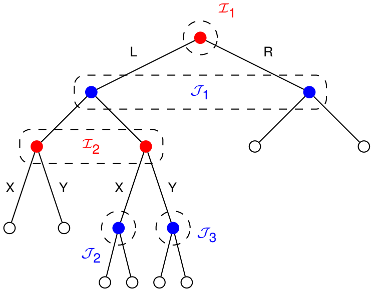

$$\Gamma = \{LX, LY, RX, RY\}$$

For example:

$$y = \begin{bmatrix} 0 \\ 0.5 \\ 0.5 \\ 0 \end{bmatrix}$$

Example

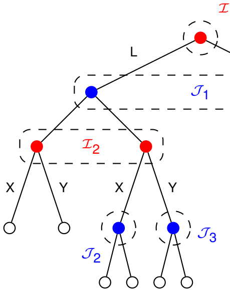

$$\Gamma = \{\text{LX}, \text{LY}, \text{RX}, \text{RY}\}$$

For example:

$$y = \begin{bmatrix} 0 \\ 0.5 \\ 0.5 \\ 0 \end{bmatrix}$$

50%

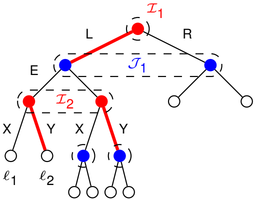

50%

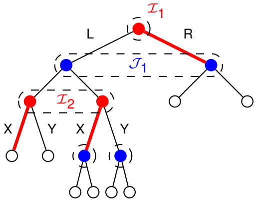

Mixed strategies

Given this definition, a number of results follow naturally.

Expected outcome of a game

$$J(y,z) = \sum_{\gamma \in \Gamma} \sum_{\sigma \in \Sigma} J(\gamma, \sigma)\text{Prob}(\text{P1 selects } \gamma\text{, P2 selects } \sigma)$$
$$= \sum_{i=1}^{n} \sum_{j=1}^{m} J(\gamma_i, \sigma_j) y_i z_j$$

Abuse of notation: $J(y,z)$ vs. $J(\gamma, \sigma)$!

If A_{ext} is the equivalent matrix form, then

$$J(y,z) = y^\top A_{\text{ext}} z$$

Mixed strategies

Mixed Nash equilibrium

$(y^*, z^*) \in \mathcal{Y} \times \mathcal{Z}$ is a **mixed saddle point** (Nash equilibrium) if

$$J(y^*, z) \leq J(y^*, z^*) \leq J(y, z^*)$$

for any $y \in \mathcal{Y}$, and any $z \in \mathcal{Z}$.

$J(y^*, z^*)$ is called the **saddle point value**.

Mixed strategies

Mixed Nash equilibrium

$(y^*, z^*) \in \mathcal{Y} \times \mathcal{Z}$ is a **mixed saddle point** (Nash equilibrium) if

$$J(y^*, z) \leq J(y^*, z^*) \leq J(y, z^*)$$

for any $y \in \mathcal{Y}$, and any $z \in \mathcal{Z}$.

$J(y^*, z^*)$ is called the **saddle point value**.

In extensive games the mixed saddle point is often non-unique. Why?

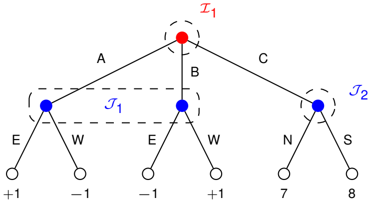

Many mixed Nash equilibria

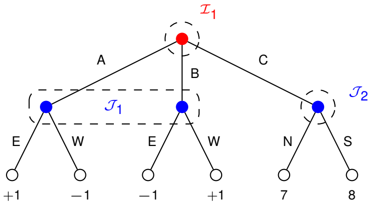

	$\sigma(\mathcal{J}_1) = E$ $\sigma(\mathcal{J}_2) = N$	$\sigma(\mathcal{J}_1) = E$ $\sigma(\mathcal{J}_2) = S$	$\sigma(\mathcal{J}_1) = W$ $\sigma(\mathcal{J}_2) = N$	$\sigma(\mathcal{J}_1) = W$ $\sigma(\mathcal{J}_2) = S$
$\gamma(\mathcal{I}_1) = A$	$+1$	$+1$	-1	-1
$\gamma(\mathcal{I}_1) = B$	-1	-1	$+1$	$+1$
$\gamma(\mathcal{I}_1) = C$	$+7$	$+8$	$+7$	$+8$

Write down NE strategy for Player 2.

Mixed strategies

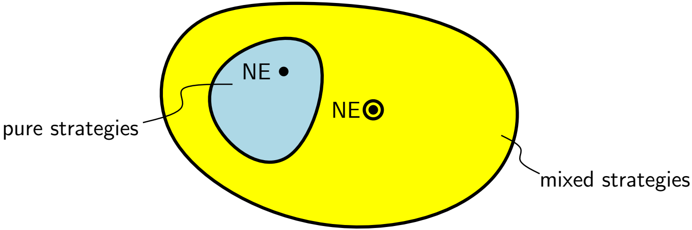

The Minmax Theorem

In any **game** (not only matrix game) the average security levels of the players in mixed strategies coincide, that is

$$\underline{V}_m = \max_{z \in \mathcal{Z}} \min_{y \in \mathcal{Y}} J(y,z) = \min_{y \in \mathcal{Y}} \max_{z \in \mathcal{Z}} J(y,z) = \overline{V}_m$$

and therefore **a mixed NE always exists**.

What is wrong with mixed strategies

- The set of mixed strategies is **very large** to explore.
  - Directly connected to the size of A_{ext}.

What is wrong with mixed strategies

- The set of mixed strategies is **very large** to explore.
  - Directly connected to the size of A_{ext}.
- Not all NE in mixed strategies are **subgame perfect**.

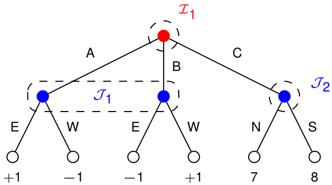

$\sigma(\mathcal{J}_2)$ is **irrelevant** for the value of the game.

Behavioral strategies

Behavioral strategies vs mixed strategies

Behavioral strategies are randomized strategies in which the randomization is done **over actions as the game is played** and not **over pure policies before the game starts**.

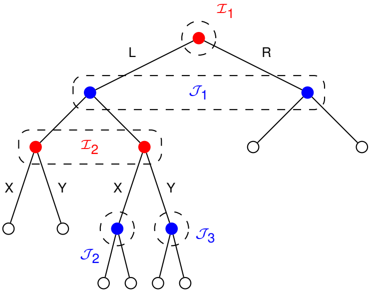

Pure strategy

A pure strategy γ maps each **information set** into an **action**

$$\gamma\left(\mathcal{I}_1\right) = L, \quad \gamma\left(\mathcal{I}_2\right) = X$$

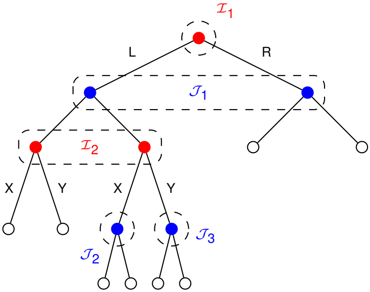

Mixed strategy

A mixture of pure strategies from

$$\Gamma = \{LX, LY, RX, RY\}$$

For example 50% LX, 50% LY

$$y = \begin{bmatrix} 0.5 & 0.5 & 0 & 0 \end{bmatrix}^\top$$

Behavioral strategy

A random action at each IS, chosen independetly.

For example, **in** $\mathcal{I}_1$, 50% L, 50% R, while **in** $\mathcal{I}_2$, 50% X, 50% Y.

$$\gamma^b(\mathcal{I}_1) = \begin{bmatrix} 0.5 \\ 0.5 \end{bmatrix} \qquad \gamma^b(\mathcal{I}_2) = \begin{bmatrix} 0.5 \\ 0.5 \end{bmatrix}$$

Behavioral strategies

- **Pure strategies**
  A map that assignes an action to each information set
  $$\mathcal{I}_i \mapsto u_i = \gamma(\mathcal{I}_i) \in \mathcal{U}_i$$
- **Mixed strategies**
  A **probability distribution** over the pure strategies $\gamma_i \in \Gamma$
  $$y \in \mathcal{Y} \subset \mathbb{R}^n, \quad n = |\Gamma|, \qquad y_i = \mathbb{P}(\gamma_i)$$
- **Behavioral strategies**
  A map that assignes a **probability distribution** over the available actions to each information set
  $$\mathcal{I}_i \mapsto \gamma^b(\mathcal{I}_i) \in \mathcal{Y}_i \subset \mathbb{R}^{|\mathcal{U}_i|}$$

Mixed vs. behavioral

Are all behavioral strategies also mixed strategies?

In a feedback game, a behavioral strategy γ^b corresponds to some mixed strategy y (similarly, for player 2).

Proof: We can construct y element by element.
Consider the pure strategy γ_i, for $i = 1, 2, \ldots, n$. For each information set $\mathcal{I}_j$, let $\gamma_i(\mathcal{I}_j) = u_j^* \in \mathcal{U}_i$.
Then, the elements of the equivalent mixed strategy y are defined as

$$\begin{aligned} y_i &= \mathbb{P}(\gamma_i) = \\ &= \mathbb{P}(u_1 = u_1^*, u_2 = u_2^*, \ldots, u_r = u_r^*) = \\ &= \mathbb{P}(u_1 = u_1^*)\,\mathbb{P}(u_2 = u_2^*) \cdots \mathbb{P}(u_r = u_r^*) \end{aligned}$$

Note: we used the independence of the randomization at different information sets (in the behavioral strategy), and the fact that you don't visit the same IS twice.

Mixed vs. behavioral

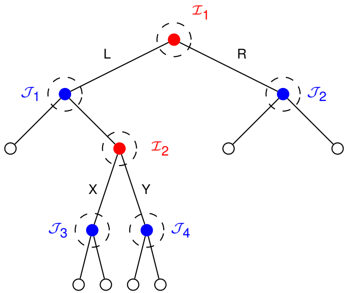

Behavioral strategy

In $\mathcal{I}_1$, 20% L, 80% R,
in $\mathcal{I}_2$, 50% X, 50% Y.

$$\gamma^b(\mathcal{I}_1) = \begin{bmatrix} 0.2 \\ 0.8 \end{bmatrix} \qquad \gamma^b(\mathcal{I}_2) = \begin{bmatrix} 0.5 \\ 0.5 \end{bmatrix}$$

Corresponding **mixed strategy**?

Pure strategies: $\Gamma = \{LX, LY, RX, RY\}$

What is $\mathbb{P}(LX)$?

$$\mathbb{P}(LX) = \mathbb{P}(u_1 = L, u_2 = X) = \mathbb{P}(u_1 = L)\,\mathbb{P}(u_2 = X) = 0.2 \cdot 0.5 = 0.1$$

$$y = \left[\mathbb{P}(LX), \mathbb{P}(LY), \mathbb{P}(RX), \mathbb{P}(RY)\right]^T = \left[0.1, 0.1, 0.4, 0.4\right]^T$$

Mixed vs. behavioral

Are all mixed strategies also behavioral strategies?

No (not even for feedback games).

Mixed vs. behavioral

Are all mixed strategies also behavioral strategies?

No (not even for feedback games).

Counterexample:

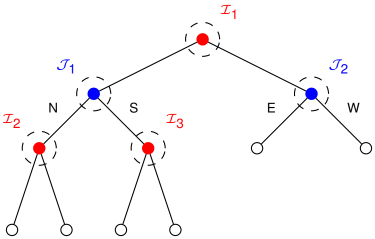

Mixed strategy

$$\Sigma = \{NE, NW, SE, SW\}$$

For example 50% NE, 50% SW

$$z = \begin{bmatrix} 0.5 & 0 & 0 & 0.5 \end{bmatrix}^\top$$

Can we describe the same mixed strategies as a behavioral strategy?

$\mathbb{P}(NW) = 0$ only if $\mathbb{P}(N) = 0$ or $\mathbb{P}(W) = 0$, which implies that either $\mathbb{P}(NE) = 0$ or $\mathbb{P}(SW) = 0$.

Mixed vs. behavioral

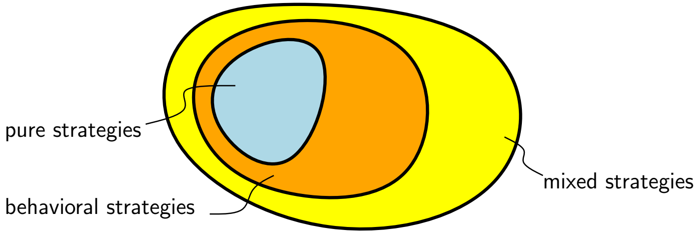

Not surpising: **degrees of freedom** in mixed / behavioral strategies

- **Mixed strategies**

$$y \in \mathcal{Y} \subset \mathbb{R}^n, \quad n = |\Gamma|,$$

  therefore $|\mathcal{U}_1| \times |\mathcal{U}_2| \times \cdots \times |\mathcal{U}_r| - 1$ degrees of freedom.

- **Behavioral strategies**

$$(y_1, y_2, \ldots, y_r), \quad y_i \in \mathcal{Y}_i \subset \mathbb{R}^{|\mathcal{U}_i|}$$

  therefore $(|\mathcal{U}_1| - 1) + (|\mathcal{U}_2| - 1) + \cdots + (|\mathcal{U}_r| - 1)$ degrees of freedom.

Mixed vs. behavioral

We have defined the smaller set of **behavioral strategies**.

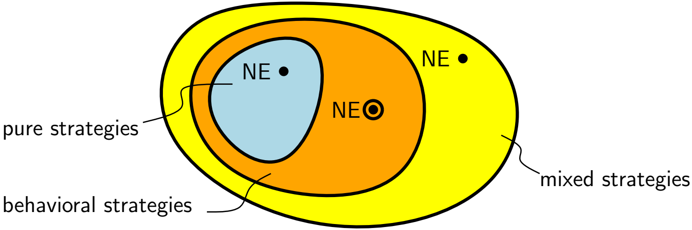

Mixed vs. behavioral

We have defined the smaller set of **behavioral strategies**.

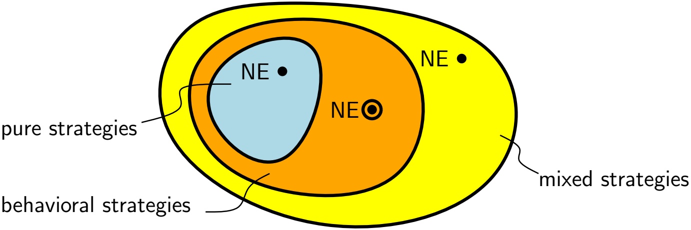

In the following, we will see that, for **feedback games**,

1. It is **not restrictive** to consider only behavioral strategies.

Mixed vs. behavioral

We have defined the smaller set of **behavioral strategies**.

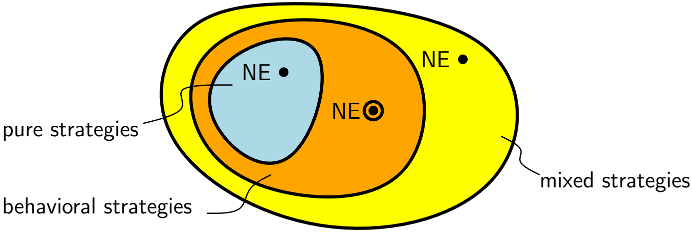

In the following, we will see that, for **feedback games**,

1. It is **not restrictive** to consider only behavioral strategies.
2. The set of behavioral strategies is **computationally tractable**
   - There is an algorithm to find saddle point behavioral strategies
   - The low search space dimension makes the algorithm efficient

Mixed vs. behavioral

We have defined the smaller set of **behavioral strategies**.

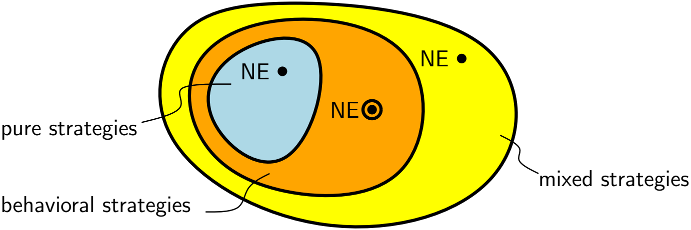

In the following, we will see that, for **feedback games**,

1. It is **not restrictive** to consider only behavioral strategies.
2. The set of behavioral strategies is **computationally tractable**
   - There is an algorithm to find saddle point behavioral strategies
   - The low search space dimension makes the algorithm efficient
3. The proposed algorithm returns a **subgame perfect** strategy.

NE in behavioral strategies

For feedback games we can look for Nash equilibria in the sets of **behavioral strategies** Γ^b and Σ^b, knowing that

- there is one
- it has the same expected outcome of any other mixed NE.

NE in behavioral strategies

For feedback games we can look for Nash equilibria in the sets of **behavioral strategies** Γ^b and Σ^b, knowing that

- there is one
- it has the same expected outcome of any other mixed NE.

Behavioral Nash equilibrium

$(\gamma^{b*}, \sigma^{b*}) \in \Gamma^b \times \Sigma^b$ is a **saddle point** (Nash equilibrium) in behavioral strategies if

$$J(\gamma^{b*}, \sigma^b) \leq J(\gamma^{b*}, \sigma^{b*}) \leq J(\gamma^b, \sigma^{b*})$$

for any $\gamma^b \in \Gamma^b$, and any $\sigma^b \in \Sigma^b$.

NE in behavioral strategies

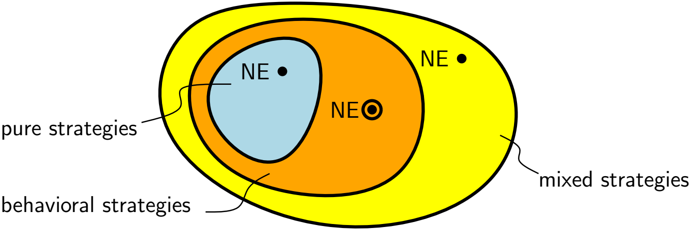

Summary

For feedback games, a Nash equilibrium in behavioral strategies exists and has the same value as any mixed strategy Nash equilibrium.

Single stage game

How to search for a Nash equilibrium behavioral strategy?

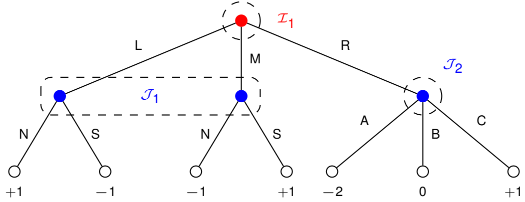

Step 1: For each $\mathcal{J}_i$, $i = 1, \ldots, s$, construct the corresponding **matrix game** where the edges entering in $\mathcal{J}_i$ are the actions for Player 1, and the edges leaving $\mathcal{J}_i$ are the actions for player 2.

	N	S
L	$+1$	-1
M	-1	$+1$

	A	B	C
R	-2	0	$+1$

Single stage game

How to search for a Nash equilibrium behavioral strategy?

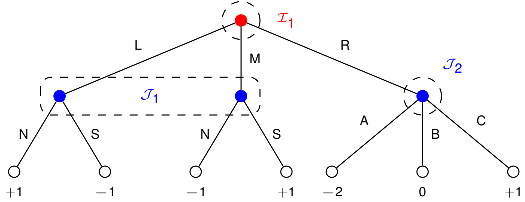

Step 2: Compute the **mixed NE** for each matrix game. The resulting NE mixed strategy for Player 2 is his **NE behavioral strategy**.

	N	S
L	$+1$	-1
M	-1	$+1$

	A	B	C
R	-2	0	$+1$

$$\sigma^*(\mathcal{J}_1) = \begin{bmatrix} 0.5 \\ 0.5 \end{bmatrix}, \quad \sigma^*(\mathcal{J}_2) = \begin{bmatrix} 0 \\ 0 \\ 1 \end{bmatrix}$$

Single stage game

How to search for a Nash equilibrium behavioral strategy?

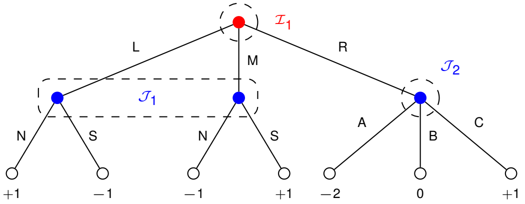

Step 3: Assign the **value** of the corresponding matrix game to each information set $\mathcal{J}_i$.

	N	S
L	$+1$	-1
M	-1	$+1$

	A	B	C
R	-2	0	$+1$

$$\sigma^*(\mathcal{J}_1) = \begin{bmatrix} 0.5 \\ 0.5 \end{bmatrix}, \quad \sigma^*(\mathcal{J}_2) = \begin{bmatrix} 0 \\ 0 \\ 1 \end{bmatrix}$$

$$V_{m,1} = 0, \quad V_{m,2} = +1$$

Single stage game

How to search for a Nash equilibrium behavioral strategy?

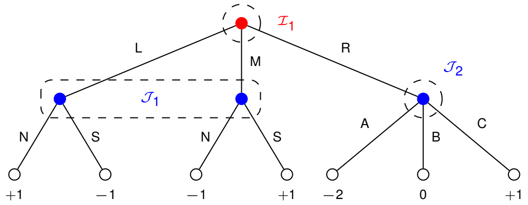

Step 4: The behavioral NE for Player 1 is given by the mixed strategy corresponding to the most favorable set $\mathcal{J}_i$.

	N	S
L	$+1$	-1
M	-1	$+1$

	A	B	C
R	-2	0	$+1$

$$\sigma^*(\mathcal{J}_1) = \begin{bmatrix} 0.5 \\ 0.5 \end{bmatrix}, \quad \sigma^*(\mathcal{J}_2) = \begin{bmatrix} 0 \\ 0 \\ 1 \end{bmatrix}$$

$$V_{m,1} = 0, \quad V_{m,2} = +1$$

$$\gamma^*(\mathcal{I}_1) = \begin{bmatrix} 0.5 \\ 0.5 \\ 0 \end{bmatrix}, \quad V_m = 0$$