COSC-6590/GSCS-6390

Games: Theory and Applications

Lecture 09 - Two-Player Non-Zero-Sum Games

Luis Rodolfo Garcia Carrillo

School of Engineering and Computing Sciences
Texas A&M University - Corpus Christi, USA

Table of contents

Security Policies and Nash Equilibria

Security Policies and Nash Equilibria

Consider a two-player game G.
P_1 and P_2 select policies within **action spaces** Γ_1 and Γ_1.

Important: game has a distinct **outcome** for each player.

$$\text{In particular, when} \quad \begin{cases} P_1 \text{ uses policy } \gamma \in \Gamma_1 \\ P_2 \text{ uses policy } \sigma \in \Gamma_2 \end{cases}$$

we denote by

- $J_1(\gamma, \sigma)$: outcome of the game for P_1
- $J_2(\gamma, \sigma)$: outcome of the game for P_2

Each player wants to **minimize** their own outcome, and does not care about the outcome of the other player.

Security Levels and Policies

Definition 9.1 (Security policy).
The **security level** for P_i, $i \in \{1, 2\}$ is defined by

$$\bar{V}_{\Gamma_1,\Gamma_2}(J_i) := \underbrace{\inf_{\gamma \in \Gamma_i}}_{\substack{\text{minimize cost assuming} \\ \text{worst choice by } P_j}} \quad \underbrace{\sup_{\sigma \in \Gamma_j, j \neq i}}_{\substack{\text{worst choice by } P_j \\ \text{from } P_i\text{'s perspective}}} \quad J_i(\gamma, \sigma)$$

a **security policy** for P_i is any policy γ^* for which the infimum above is achieved, i.e.,

$$\bar{V}_{\Gamma_1,\Gamma_2}(J_i) := \inf_{\gamma \in \Gamma_i} \sup_{\sigma \in \Gamma_j, j \neq i} J_i(\gamma, \sigma) = \underbrace{\sup_{\sigma \in \Gamma_j, j \neq i} J_i(\gamma^*, \sigma)}_{\gamma^* \text{ achieves the infimum}}$$

A pair of policies (γ^*, σ^*) is a **minimax pair** if γ^* and σ^* are security policies for P_1 and P_2, respectively.

Nash Equilibria for non-zero-sum games

Here, Nash Equilibria (NE) captures the notion of **no regret**

- after knowing the choice made by the other player, each player finds that their own policy provided the lowest possible cost against the choice of the other player.

Definition 9.2 (Nash equilibrium). A pair of policies $(\gamma^*, \sigma^*) \in \Gamma_1 \times \Gamma_2$ is called a Nash equilibrium (NE) if

$$J_1(\gamma^*, \sigma^*) \leq J_1(\gamma, \sigma^*), \quad \forall \gamma \in \Gamma_1$$
$$J_2(\gamma^*, \sigma^*) \leq J_2(\gamma^*, \sigma), \quad \forall \sigma \in \Gamma_2$$

and the pair $(J_1(\gamma^*, \sigma_*), J_2(\gamma^*, \sigma^*))$ is the **Nash outcome of the game**.

Nash Equilibria for non-zero-sum games

For a **zero-sum game** in which $J_2(\gamma, \sigma) = -J_1(\gamma, \sigma)$, the 2nd eq. of the Nash Equilibrium becomes

$$-J_1(\gamma^*, \sigma^*) = -J_1(\gamma^*, \sigma), \quad \forall \sigma \in \Gamma_2$$

and we can re-write the Nash Equilibrium as

$$J_1(\gamma^*, \sigma) \leq J_1(\gamma^*, \sigma^*) \leq J_1(\gamma, \sigma^*) \qquad \forall \gamma \in \Gamma_1, \sigma \in \Gamma_2$$

Then, saddle-point equilibria are NE for zero-sum games.

Attention! However, it does not make sense to talk about saddle-point equilibria for a non-zero-sum game.

Bimatrix Games

Bimatrix Games

Pure bimatrix games are played by two players, each selecting policies from finite **action spaces**:

- P_1 has available m actions: $\Gamma_1 := \{1, 2, \ldots, m\}$
- P_2 has available n actions: $\Gamma_2 := \{1, 2, \ldots, n\}$

Outcomes for the players are quantified by two $m \times n$ matrices $A = [a_{ij}]$ and $B = [b_{ij}]$, one for P_1 and the other for P_2.

$$\text{When} \quad \begin{cases} P_1 \text{ selects action } i \in \Gamma_1 := \{1, 2, \ldots, m\} \\ P_2 \text{ selects action } j \in \Gamma_2 := \{1, 2, \ldots, n\} \end{cases}$$

$$\text{we have that} \quad \begin{cases} J_1 := a_{ij} \text{ is the outcome for } P_1 \\ J_2 := b_{ij} \text{ is the outcome for } P_2 \end{cases}$$

Note. P_1 selects a row of A/B and P_2 selects a column of A/B.

Bimatrix Games

P_1 and P_2 want to **minimize** their outcomes J_1 and J_2.

Zero-sum games: a special case of bimatrix games, for which $B = -A$. For example, when $B = -A$:

- as P_2 attempts to minimize $J_2 := b_{ij}$
- P_2 is also maximizing $-J_1 = a_{ij}$.

For the action spaces Γ_1 and Γ_2, the resulting security levels, policies, and Nash equilibria are called **pure**.

Note: since these action spaces are finite, security policies and minimax pairs always exist

Bimatrix Games

Example 9.1 (Prisoners' dilemma).

P_1 and P_2 are two prisoners arrested for a minor crime, but suspected of having committed a serious crime.

There is little evidence that incriminates them of the serious crime, so the prosecution's hope is that one of them incriminates the other.

The prisoners thus have two options:

$$\begin{cases} \text{action 1: do not confess} \\ \text{action 2: cooperate with the prosecution by testifying against other} \end{cases}$$

Let's make the following associations.

Bimatrix Games

Consider the bimatrix game defined by

$$A = \underbrace{\begin{bmatrix} 2 & 30 \\ 0 & 8 \end{bmatrix}}_{P_2 \text{ choices}} \Big\} \; P_1 \text{ choices} \qquad B = \underbrace{\begin{bmatrix} 2 & 0 \\ 30 & 8 \end{bmatrix}}_{P_2 \text{ choices}} \Big\} \; P_1 \text{ choices}$$

Outcomes: the number of years spent in jail resulting from their actions:

- **neither confesses:** they are both convicted of the minor crime and spend 2 years in jail
- **both cooperate:** they both get some deal, but still get 8 years for the serious crime
- **only one cooperates:** the one that testifies is released but the other spends 30 years in prison.

Bimatrix Games

Security levels and policies for this game

$$\bar{V}(A) = \min_i \max_j a_{ij} = 8, \qquad i^* = \arg\min_i \max_j a_{ij} = 2$$

$$\bar{V}(B') = \min_j \max_i b_{ij} = 8, \qquad j^* = \arg\min_j \max_i b_{ij} = 2$$

Note. The two **security policies** correspond to **confessing**.

Observation: notation for security level of P_2 is $\bar{V}(B')$

1. consistent with notation $\bar{V}(\cdot)$ used in zero-sum games, where the min is taken over rows and the max over cols.

Game has a single Nash equilibrium: the minimax pair

$(2, 2)$ **is a Nash equilibrium with outcome** $(8, 8)$

Bimatrix Games

$$A = \underbrace{\begin{bmatrix} 2 & 30 \\ 0 & 8 \end{bmatrix}}_{P_2 \text{ choices}} \Big\} \; P_1 \text{ choices} \qquad B = \underbrace{\begin{bmatrix} 2 & 0 \\ 30 & 8 \end{bmatrix}}_{P_2 \text{ choices}} \Big\} \; P_1 \text{ choices}$$

Paradox: pair of policies $(1, 1)$ would lead to outcome $(2, 2)$

- a strict improvement for both players.

For noncooperative games, there is no paradox

- this solution is not robust
- either player can profit from deviating from it

Implementing (1,1) requires cooperation and mutual trust.

We are interested in **noncooperative** solutions: reached by the players without negotiation or faith/trust between them.

Admissible Nash Equilibria

Admissible Nash Equilibria

Example 9.2. Consider the bimatrix game defined by

$$A = \underbrace{\begin{bmatrix} 1 & 0 \\ 2 & -1 \end{bmatrix}}_{P_2 \text{ choices}} \Big\} \; P_1 \text{ choices} \qquad B = \underbrace{\begin{bmatrix} 2 & 3 \\ 1 & 0 \end{bmatrix}}_{P_2 \text{ choices}} \Big\} \; P_1 \text{ choices}$$

The security levels and policies for this game are

$$\bar{V}(A) = \min_i \max_j a_{ij} = 1, \qquad i^* = \arg\min_i \max_j a_{ij} = 1$$

$$\bar{V}(B') = \min_j \max_i b_{ij} = 2, \qquad j^* = \arg\min_i \max_j b_{ij} = 1$$

Unique security policies: Game has two Nash equilibria

$(1, 1)$ **is a Nash eq. with outcome** $(1, 2)$

$(2, 2)$ **is a Nash eq. with outcome** $(-1, 0)$

Admissible Nash Equilibria

Differences between bimatrix games and zero-sum matrix games:

Bimatrix games may have several Nash equilibria

- like zero-sum games that may have several saddle-point equilibria.

Nash equilibria are not always security policies

- unlike zero-sum games for which saddle-point equilibria are always security policies.

Different equilibria to bimatrix games may have different outcomes

- unlike zero-sum games for which saddle-point equilibria always have the same value.

Admissible Nash Equilibria

Definition 9.3 (Admissible Nash equilibria).

A Nash equilibrium $(\gamma^*, \sigma^*) \in \Gamma_1 \times \Gamma_2$ is **admissible** if there is no **better** Nash equilibrium in the sense that there is no other Nash equilibrium $(\bar{\gamma}^*, \bar{\sigma}^*) \in \Gamma_1 \times \Gamma_2$ such that

$$J_1(\bar{\gamma}^*, \bar{\sigma}^*) \leq J_1(\gamma^*, \sigma^*) \qquad J_2(\bar{\gamma}^*, \bar{\sigma}^*) \leq J_2(\gamma^*, \sigma^*)$$

with at least one of these inequalities strict.

i.e.: both players do no worse with $(\bar{\gamma}^*, \bar{\sigma}^*)$ and at least one of them does strictly better.

Admissible Nash Equilibria - Battle of the sexes (BoS)

BoS is a two-player coordination game.

A couple agreed to meet this evening, but cannot recall if they will be attending a **baby shower** or a **football game**

- the fact that they forgot is **common knowledge**.

The **husband** would prefer to go to the **baby shower**.
The **wife** would rather go to the **football game**.

They prefer to go to the **same place** rather than different ones

- additional harm might come from not only going to different locations, but going to the wrong one as well
- e.g. he goes to the football game while she goes to the baby shower, satisfying neither.

If they cannot communicate, where should they go?

Admissible Nash Equilibria - Battle of the sexes (BoS)

Example 9.3 Let's make the following associations:

$$\begin{cases} P_1\text{: the husband/boyfriend} \\ P_2\text{: the wife/girlfriend} \end{cases} \qquad \begin{cases} \text{action 1: } \textbf{going to a baby shower} \\ \text{action 2: } \textbf{going to a football game} \end{cases}$$

$$A = \underbrace{\begin{bmatrix} -2 & 1 \\ 0 & -1 \end{bmatrix}}_{P_2 \text{ choices}} \Big\} P_1 \text{ choices} \qquad B = \underbrace{\begin{bmatrix} -1 & 3 \\ 2 & -2 \end{bmatrix}}_{P_2 \text{ choices}} \Big\} P_1 \text{ choices}$$

Under the outcomes provided from the matrices A and B

- they have the most fun if they go together, but
- the husband prefers the baby shower,
- the wife prefers the football game.

Admissible Nash Equilibria

Battle of the sexes is a bimatrix game

$$A = \underbrace{\left.\begin{bmatrix} -2 & 1 \\ 0 & -1 \end{bmatrix}\right\}}_{P_2 \text{ choices}} P_1 \text{ choices} \qquad B = \underbrace{\left.\begin{bmatrix} -1 & 3 \\ 2 & -2 \end{bmatrix}\right\}}_{P_2 \text{ choices}} P_1 \text{ choices}$$

The security levels and policies for this game are

$$\bar{V}(A) = \min_i \max_j a_{ij} = 0, \qquad i^* = \arg\min_i \max_j a_{ij} = 2$$

$$\bar{V}(B') = \min_j \max_i b_{ij} = 2, \qquad j^* = \arg\min_j \max_i b_{ij} = 1$$

Game has two Nash equilibria

$(1,1)$ **is a Nash eq. with outcome** $(-2,-1)$

$(2,2)$ **is a Nash eq. with outcome** $(-1,-2)$

both are admissible: none is **better** than the other.

Admissible Nash Equilibria

Attention! The minimax pair $(2,1)$ is not a Nash equilibrium.

The minimax $(1,2)$ nor $(2,1)$ are Nash equilibria:

- in both cases, both players regret their choices
- after knowing the other player's decision, they wish they had done things differently.

Two additional important properties of bimatrix games

1.- Bimatrix games may have several admissible NE.

2.- Nash equilibria are not interchangeable:

- (γ_1^*, σ_1^*) and (γ_2^*, σ_2^*) may both be Nash equilibria
- but (γ_1^*, σ_2^*) and (γ_2^*, σ_1^*) may not be.

Mixed Policies

Mixed Policies

Players select their actions randomly, according to previously selected probability distributions

Consider a game specified by two $m \times n$ matrices A and B that determine the outcomes for P_1 and P_2.

Mixed Policy for P_1: a set of numbers

$$\{y_1, y_2, \ldots, y_m\}, \quad \sum_{i=1}^{m} y_i = 1, \quad y \geq 0, \ \ \forall i \in \{1, 2, \ldots, m\}$$

y_i: probability that P_1 uses to select action $i \in \{1, 2, \ldots, m\}$.

Mixed Policy for P_2: a set of numbers

$$\{z_1, z_2, \ldots, z_n\}, \quad \sum_{i=1}^{n} z_j = 1, \quad z \geq 0, \ \ \forall j \in \{1, 2, \ldots, n\}$$

z_i: probability that P_2 uses to select action $j \in \{1, 2, \ldots, n\}$.

Mixed Policies

Random selections by players are statistically independently.

P_1 and P_2 try to minimize their expected outcomes

$$J_1 = \sum_{i,j} a_{ij} y_i z_j = y' A z \qquad J_2 = \sum_{i,j} b_{ij} y_i z_j = y' B z$$

where $y := [y_1 \;\; y_2 \;\; \cdots \;\; y_m]$ and $z := [z_1 \;\; z_2 \;\; \cdots \;\; z_n]$.

Use the concepts of security levels, security policies, and NE with the understanding that:

1.- Action spaces are the sets $\mathcal{Y}$ and $\mathcal{Z}$ of all mixed policies for players P_1 and P_2, respectively

2.- For a pair of mixed policies $y \in \mathcal{Y}$ for P_1 and $z \in \mathcal{Z}$ for P_2

- $J_1(y, z) := y' A z$ is the outcome for P_1
- $J_2(y, z) := y' B z$ is the outcome for P_2

Mixed Policies

Definition 9.4 (Mixed Nash equilibrium).

A pair of policies $(y^*, z^*) \in \mathcal{Y} \times \mathcal{Z}$ is a mixed Nash equilibrium if

$$y^{*\prime} A z^* \leq y' A z^*, \quad \forall y \in \mathcal{Y} \qquad\qquad y^{*\prime} B z^* \leq y^{*\prime} A z, \quad \forall z \in \mathcal{Z}$$

and $(y^{*\prime} A z^*, y^{*\prime} B z^*)$ is the mixed Nash outcome of the game.

The introduction of mixed policies enlarges the action spaces for both players to the point that Nash equilibria now always exist.

Theorem 9.1 (Nash).
Every bimatrix game has at least one mixed Nash equilibrium.

Best-Response Equivalent Games and Order Interchangeability

Best-Response Eq. Games and Order Interchangeability

Consider two general two-player games G and H with the same **action spaces** Γ_1 and Γ_2 but different outcomes.

For the same pair of policies $\gamma \in \Gamma_1$ and $\sigma \in \Gamma_2$

- G has outcomes $G_1(\gamma, \sigma)$ for P_1 and $G_2(\gamma, \sigma)$ for P_2
- H has outcomes $H_1(\gamma, \sigma)$ for P_1 and $H_2(\gamma, \sigma)$ for P_2

Definition 9.5 (Best-response equivalent)
The games G and H are **best-response equivalent** (BRE) if they have the same set of Nash equilibria (NE), i.e., a pair of policies (γ, σ) is a NE for G if and only if it is a NE for H.

BRE allows us to characterize a class of games for which we have order interchangeability for NE.

Best-Response Eq. Games and Order Interchangeability

Proposition 9.1 (Order interchangeability).
The NE of game G are interchangeable if G is best-response equivalent to zero-sum game H that is zero-sum.

Proof Proposition 9.1. is consequence of the facts that

1. zero-sum games enjoy the order interchangeability property
2. if two games are BRE they have the same NE.

If (γ_1^*, σ_1^*) and (γ_2^*, σ_2^*) are both NE for G, then because of BRE (γ_1^*, σ_2^*) and (γ_2^*, σ_1^*) are also NE for H.
Since H is zero-sum, then (γ_1^*, σ_2^*) and (γ_2^*, σ_1^*) are also NE for H. Because of BRE, (γ_1^*, σ_2^*) and (γ_2^*, σ_1^*) must be NE for G.

It is possible to show two games are BRE by examining the functions defining their outcomes, without computing their NE.

Best-Response Eq. Games and Order Interchangeability

Lemma 9.1. Suppose that there exist two monotone strictly increasing scalar functions $\alpha : \mathbb{R} \to \mathbb{R}$ and $\beta : \mathbb{R} \to \mathbb{R}$ such that

$$H_1(\gamma, \sigma) = \alpha(G_1(\gamma, \sigma)), \quad H_2(\gamma, \sigma) = \beta(G_2(\gamma, \sigma)), \qquad \forall \gamma \in \Gamma_1, \sigma \in \Gamma_2$$

then G and H are best-response equivalent.

Proof of Lemma 9.1. Show that if (γ^*, σ^*) is a NE of G then it is also a NE of H.
First, assume we are given a NE (γ^*, σ^*) of G, for which

$$G_1(\gamma^*, \sigma^*) \leq G_1(\gamma, \sigma^*), \quad \forall \gamma, \qquad G_2(\gamma^*, \sigma^*) \leq G_2(\gamma^*, \sigma), \quad \forall \sigma$$

Applying the monotone functions α and β to both sides of the left and right-hand side inequalities, respectively, we obtain

$$\alpha(G_1(\gamma^*, \sigma^*)) \leq \alpha(G_1(\gamma, \sigma^*)), \quad \forall \gamma, \qquad \beta(G_2(\gamma^*, \sigma^*)) \leq \beta(G_2(\gamma^*, \sigma)), \quad \forall \sigma$$

Best-Response Eq. Games and Order Interchangeability

From this we conclude that

$$H_1(\gamma^*, \sigma^*) \leq H_1(\gamma, \sigma^*), \quad \forall\gamma, \qquad H_2(\gamma^*, \sigma^*) \leq H_2(\gamma^*, \sigma), \quad \forall\sigma$$

which confirms that (γ^*, σ^*) is indeed a NE of H.

Suppose we are given a NE (γ^*, σ^*) of H, for which

$$H_1(\gamma^*, \sigma^*) \leq H_1(\gamma, \sigma^*), \quad \forall\gamma, \qquad H_2(\gamma^*, \sigma^*) \leq H_2(\gamma^*, \sigma), \quad \forall\sigma$$

From which we obtain

$$\alpha(G_1(\gamma^*, \sigma^*)) \leq \alpha(G_1(\gamma, \sigma^*)), \quad \forall\gamma, \qquad \beta(G_2(\gamma^*, \sigma^*)) \leq \beta(G_2(\gamma^*, \sigma)), \quad \forall\sigma$$

Functions α and β are monotone strictly increasing, then

$$G_1(\gamma^*, \sigma^*) \leq G_1(\gamma, \sigma^*), \quad \forall\gamma, \qquad G_2(\gamma^*, \sigma^*) \leq G_2(\gamma^*, \sigma), \quad \forall\sigma$$

which confirms that (γ^*, σ^*) is indeed a Nash equilibrium of G.

Best-Response Eq. Games and Order Interchangeability

Attention!

Given a two-player game G one can try to find monotonically strictly increasing functions $\alpha, \beta : \mathbb{R} \to \mathbb{R}$ such that Lemma 9.1 holds with $H_1(\gamma, \sigma) = -H_2(\gamma, \sigma)$, for all γ, σ which would allow us to conclude that the NE of G are interchangeable and all Nash outcomes are equal to each other.

Specializing this to a bimatrix game G defined by a pair of $m \times n$ matrices A and B and pure policies, this amounts to finding monotonically strictly increasing functions $\alpha, \beta : \mathbb{R} \to \mathbb{R}$ such that

$$\alpha(a_{ij}) = -\beta(b_{ij}) \quad \forall i, j$$

Best-Response Eq. Games and Order Interchangeability

Restricting our search, e.g., to polynomial functions of the type

$$\alpha(s) := \sum_{k=1}^{\ell} a_k s^k \qquad \beta(s) := \sum_{k=1}^{\ell} b_k s^k$$

the previous equality

$$\alpha(a_{ij}) = -\beta(b_{ij}) \ \ \forall i, j$$

leads to linear equations on the polynomial coefficients, which are easy to solve.

One would still need to verify the monotonicity of the polynomials so obtained (over the range of possible game outcomes).

Best-Response Eq. Games and Order Interchangeability

Monotone function.

Function $f : \mathbb{R} \to \mathbb{R}$ is said to be **monotone non-decreasing** if

$$x \geq y \Rightarrow f(x) \geq f(y), \quad \forall x, y \in \mathbb{R}$$

and it is said to be **monotone strictly increasing** if

$$x > y \Rightarrow f(x) > f(y), \quad \forall x, y \in \mathbb{R}$$

which is also equivalent to say that

$$f(x) \leq f(y) \Rightarrow x \leq y, \quad \forall x, y \in \mathbb{R}$$

Best-Response Eq. Games and Order Interchangeability

Attention!

The lack of interchangeability is an **unpleasant** possibility in non-zero-sum games and leads to the following hierarchy of two-player games:

1. Games with single NE or with multiple but interchangeable NE with equal values are the most **predictable** for noncooperative rational players.

- this class of games includes all zero-sum games and the prisoners' dilemma.

Best-Response Eq. Games and Order Interchangeability

2. Games with a single admissible NE or with multiple but interchangeable admissible NE with equal values are still fairly predictable for noncooperative rational players.

- e.g., bimatrix game in **Example 9.2** or the one defined by

$$A = B = \begin{bmatrix} 0 & 2 \\ 2 & 1 \end{bmatrix}$$

  with a single admissible NE $(1, 1)$.
  Note that $(2, 2)$ is also NE, but it is not admissible.

3. In games with multiple admissible NE that are interchangeable but have different values, noncooperative rational players will likely end up in a NE, but it will generally be difficult to predict which.

Best-Response Eq. Games and Order Interchangeability

4. Games with multiple admissible NE that are not interchangeable are problematic. It is unclear whether or not the players will find a common equilibrium.

- e.g., battle of the sexes or the bimatrix game defined by

$$A = B = \begin{bmatrix} 0 & 1 \\ 1 & 0 \end{bmatrix}$$

with two admissible but non-interchangeable NE $(1, 1)$ and $(2, 2)$ with the same value $(0, 0)$.

When played repeatedly, these games can lead to persistent oscillations in the policies used by the players: they may try to constantly adjust to the most recent policy used by the other.

Best-Response Eq. Games and Order Interchangeability

What options do we have for the latter types of games in a noncooperative setting in which one should not rely on negotiation/trust between players?

1. The players may simply use security policies, leading to minimax solutions. Such solutions are often costly for both players and therefore not efficient.

2. When possible, the reward structure of the game should be changed to avoid inefficient solutions and policy oscillations in repeated games.

It is possible to **reshape** the reward structure of a game in economics (and engineering) through pricing, taxation, or other incentives/deterrents.

Practice Exercises

Practice Exercises

9.1 (Order interchangeability for Nash).

Consider two NE (γ_1^*, σ_1^*) and (γ_2^*, σ_2^*) for a two-player game.

Show that if these two equilibria are interchangeable in the sense that (γ_1^*, σ_2^*) and (γ_2^*, σ_1^*) are also NE, then

$$G_1(\gamma_1^*, \sigma_1^*) = G_1(\gamma_2^*, \sigma_1^*), \qquad G_1(\gamma_2^*, \sigma_2^*) = G_1(\gamma_1^*, \sigma_2^*)$$
$$G_2(\gamma_1^*, \sigma_1^*) = G_2(\gamma_1^*, \sigma_2^*), \qquad G_2(\gamma_2^*, \sigma_2^*) = G_2(\gamma_2^*, \sigma_1^*)$$

Solution to Exercise 9.1.
Since (γ_1^*, σ_1^*) is a Nash equilibrium, we must have

$$G_1(\gamma_1^*, \sigma_1^*) \leq G_1(\gamma_2^*, \sigma_1^*), \qquad G_2(\gamma_1^*, \sigma_1^*) \leq G_2(\gamma_1^*, \sigma_2^*),$$

Practice Exercises

but since (γ_2^*, σ_1^*) and (γ_1^*, σ_2^*) are also Nash equilibria we must also have that

$$G_1(\gamma_2^*, \sigma_1^*) \leq G_1(\gamma_1^*, \sigma_1^*), \qquad G_2(\gamma_1^*, \sigma_2^*) \leq G_2(\gamma_1^*, \sigma_1^*),$$

therefore we actually have

$$G_1(\gamma_1^*, \sigma_1^*) = G_1(\gamma_2^*, \sigma_1^*), \qquad G_2(\gamma_1^*, \sigma_1^*) = G_2(\gamma_1^*, \sigma_2^*),$$

Similarly, using the facts that (γ_2^*, σ_2^*), (γ_1^*, σ_2^*), and (γ_2^*, σ_1^*) are all Nash equilibria, we can also conclude that

$$G_1(\gamma_2^*, \sigma_2^*) = G_1(\gamma_1^*, \sigma_2^*), \qquad G_2(\gamma_2^*, \sigma_2^*) = G_2(\gamma_2^*, \sigma_1^*),$$

which concludes the proof.

End of Lecture

09 - Two-Player Non-Zero-Sum Games

Questions?