COSC-6590/GSCS-6390

Games: Theory and Applications

Lecture 04 - Mixed Policies

Luis Rodolfo Garcia Carrillo

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

Table of contents

Mixed Policies

Mixed Policies: Rock-Paper-Scissor

Making the following associations

$$\text{actions} = \begin{cases} \text{rock} & \equiv 1 \\ \text{paper} & \equiv 2 \\ \text{scissor} & \equiv 3 \end{cases} \qquad \text{outcomes} = \begin{cases} P_1 \text{ wins} & \equiv -1 \quad \text{minimizer} \\ P_2 \text{ wins} & \equiv +1 \quad \text{maximizer} \\ \text{draw} & \equiv 0 \end{cases}$$

The rock-paper-scissor game can be viewed as a matrix

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

For this game

- P_1's security level: $\bar{V}(A) = +1$. Any row is a security policy.
- P_2's security level: $\underline{V}(A) = -1$. Any col is a security policy.

Conclusion: we have a strict inequality: $\underline{V}(A) < \bar{V}(A)$

- game has no pure saddle-point equilibria.

Mixed Policies: Rock-Paper-Scissor

So far we studied **pure policies**

- choices of particular actions
  - perhaps based on some observation

We now introduce **mixed policies**

- choosing a probability distribution to select actions
  - perhaps as a function of observations

Players select their actions **randomly** according to a previously selected **probability distribution**.

Mixed Policies: Rock-Paper-Scissor

Consider a game specified by an $m \times n$ matrix A

- m actions for P_1 and n actions for P_2.

A **mixed policy** for P_1 is a set of numbers

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

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

A **mixed policy** for P_2 is a set of numbers

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

z_j: probability that P_2 uses to select the action $j \in \{1, 2, \dots, n\}$.

Mixed Policies: Rock-Paper-Scissor

Assumption: random selections by both players are done statistically independently.

Due to **randomness**, the same pair of mixed policies will lead to different outcomes as one plays the game again and again.

With **mixed policies**, players try to optimize the **expected outcome of the game**:

$$J = \sum_{i=1}^{m} \sum_{j=1}^{n} a_{ij} \text{ Prob } (P_1 \text{ selects } i \text{ and } P_2 \text{ selects } j)$$

Players make selections independently, then this simplifies to

$$J = \sum_{i,j} a_{ij} \text{ Prob } (P_1 \text{ selects } i) \text{ Prob } (P_2 \text{ selects } j) = \sum_{i,j} a_{ij} y_i z_j = y' A z$$

Mixed Policies: Rock-Paper-Scissor

$$J = y' A z$$

where y and z are the following column vectors:

$$y := \begin{bmatrix} y_1 \\ y_2 \\ \vdots \\ y_m \end{bmatrix} \qquad z := \begin{bmatrix} z_1 \\ z_2 \\ \vdots \\ z_n \end{bmatrix}$$

Notation.- symbol “ $'$ ” denotes matrix/vector transpose.

Objective (mixed policies):
P_1 wants to minimize the expected outcome $J = y' A z$.
P_2 wants to maximize the expected outcome $J = y' A z$.

Mixed Policies: Common Interpretations

Interpretation 1: Repeated game paradigm

P_1 and P_2 face each other multiple times.

In each game they choose their actions randomly according to preselected mixed policies (independently from each other, and from game to game).

Goal: minimize/maximize the cost/reward averaged over all the games played.

Paradigm appropriate for games in

- economics, e.g., in the advertising campaign game or the tax-payers auditing game;
- political/social **engineering**, e.g., in the crime/war deterrence games or the worker's compensation game.

Mixed Policies: Common Interpretations

Interpretation 2: Large population paradigm

There is large population of players P_1, and an equally large population of players P_2.

Players play pure policies, but percentage of players that play each pure policy matches the probabilities of the mixed policies.

Two players are selected randomly from each population (independently), they play against each other.

Goal: select a **good mix** for the populations so as to minimize/maximize the expected cost/reward.

Paradigm appropriate for games in

- tax auditing, crime deterrence, workers compensation, and some robust design problems.

Mixed Action Spaces

Mixed Action Spaces

Consider a game specified by an $m \times n$ matrix A

- m actions for P_1 and n actions for P_2.

With **pure policies**, the (**pure**) **action spaces** for P_1 and P_2 consist of the finite sets

$$\{1, 2, \ldots, m\} \qquad \text{and} \qquad \{1, 2, \ldots, n\}$$

With **mixed policies**, P_1 and P_2 choose distributions so their (**mixed**) **action spaces** consist of (infinite) continuous sets

$$\mathcal{Y} := \left\{ y \in \mathbb{R}^m : \sum_i y_i = 1, \;\; y_i \geq 0, \;\; \forall i \right\}, \quad \mathcal{Z} := \left\{ z \in \mathbb{R}^n : \sum_j z_j = 1, \;\; z_j \geq 0, \;\; \forall j \right\}$$

Sets such as $\mathcal{Y}$ and $\mathcal{Z}$ are called (**probability**) **simplexes**.

Mixed Action Spaces

Attention!

Pure policies still exist within the mixed action spaces.

For example, the vector

$$\begin{bmatrix} 0 & 1 & 0 & \cdots & 0 \end{bmatrix}' \in \mathcal{Y}$$

is the pure policy that consists of picking action 2, because

- the probability y_2 of picking this action is one
- the probabilities y_i, $i \neq 2$ of picking other actions are all equal to zero.

Mixed Security Policies and Saddle-Point Equilibrium

Mixed Security Policies and Saddle-Point Equilibrium

Definition 4.1 (Mixed security policy).

Consider a matrix game A. The **average security level** for P_1 (the minimizer) is

$$\bar{V}_m(A) := \underbrace{\min_{y \in \mathcal{Y}}}_{\substack{\text{minimize cost assuming} \\ \text{worst choice by } P_2}} \quad \underbrace{\max_{z \in \mathcal{Z}}}_{\substack{\text{worst choice by } P_2 \\ \text{from } P_1\text{'s perspective}}} \quad y' A z$$

Corresponding **mixed security policy** for P_1: any y^* that achieves the desired average security level, i.e.,

$$\underbrace{\max_{z \in \mathcal{Z}} y^{*\prime} A z = \bar{V}_m(A)}_{y^* \text{ achieves the minimum}} := \min_{y \in \mathcal{Y}} \max_{z \in \mathcal{Z}} \quad y' A z$$

$\in$ since there may be several y^* that achieve the minimum.

Mixed Security Policies and Saddle-Point Equilibrium

Definition 4.1 (Mixed security policy).

Conversely, the **average security level** for P_2 (maximizer) is

$$\underline{V}_m(A) := \underbrace{\max_{z\in\mathcal{Z}}}_{\substack{\text{maximize reward assuming}\\ \text{worst choice by } P_1}} \quad \underbrace{\min_{y\in\mathcal{Y}}}_{\substack{\text{worst choice by } P_1\\ \text{from } P_2\text{'s perspective}}} \quad y'Az$$

Corresponding **mixed security policy** for P_2: any z^* that achieves the desired average security level, i.e.,

$$\underbrace{\min_{y\in\mathcal{Y}} y^{*\prime}Az = \underline{V}_m(A)}_{z^* \text{ achieves the maximum}} := \max_{z\in\mathcal{Z}}\min_{y\in\mathcal{Y}} \quad y'Az$$

$\in$ since there may be several z^* that achieve the maximum.

Mixed Security Policies and Saddle-Point Equilibrium

Definition 4.2 (Mixed saddle-point equilibrium).

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

$$
\begin{aligned}
y^{*\prime} A z^* &\leq y' A z^*, \quad \forall y \in \mathcal{Y} \\
y^{*\prime} A z^* &\geq y{\ast}' A z, \quad \forall z \in \mathcal{Z}
\end{aligned}
$$

and $y^{*\prime} A z^*$ is the **mixed saddle-point value**.

These equations are often re-written as

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

Mixed Security Policies and Saddle-Point Equilibrium

Proposition 4.1 (Security levels/policies).
For every (finite) matrix A, the following properties hold:

P4.1 Average security levels are well defined and unique.

P4.2 Both players have mixed security policies (not necessarily unique).

P4.3 The average security levels always satisfy

$$\underbrace{\underline{V}(A)}_{\max_j \min_i a_{ij}} \leq \underbrace{\underline{V}_m(A)}_{\max_z \min_y y'Az} \leq \underbrace{\bar{V}_m(A)}_{\min_y \max_z y'Az} \leq \underbrace{\bar{V}(A)}_{\min_i \max_j a_{ij}}$$

Consequence: when there is a pure saddle-point equilibrium, $\underline{V}(A) = \bar{V}(A)$ are equal, and the average security levels are exactly the same as the (pure) security levels.

Mixed Security Policies and Saddle-Point Equilibrium

The inequality expresses a feature of **mixed policies**

$$\underbrace{\underline{V}(A)}_{\max_j \min_i a_{ij}} \quad \leq \quad \underbrace{\underline{V}_m(A)}_{\max_z \min_y y'Az} \quad \leq \quad \underbrace{\bar{V}_m(A)}_{\min_y \max_z y'Az} \quad \leq \quad \underbrace{\bar{V}(A)}_{\min_i \max_j a_{ij}}$$

They **lead to security levels that are better** than those of pure policies for both players.

- Left-most inequality means that the mixed security level for the maximizer is larger than the pure security level
- Right-most inequality means that the mixed security level for the minimizer is smaller than the pure security level.

Mixed Security Policies and Saddle-Point Equilibrium

Take for example

$$\underline{V}_m(A) = \max_{z \in \mathcal{Z}} \min_{y \in \mathcal{Y}} y' A z$$

Property **P4.1** states that the min and max in $\underline{V}_m(A)$ are achieved at specific points in $\mathcal{Y}$ and $\mathcal{Z}$, respectively.

- we are minimizing/maximizing a continuous function over a compact set (i.e., bounded and closed)

Weierstrass' Theorem guarantees that such a min/max always exists at some point in the set.

The max is achieved at some point $z^* \in \mathcal{Z}$, then z^* can be used in a security policy for P_2 (maximizer), which justifies **P4.2**.

The same reasoning applies for $\bar{V}_m(A)$ and the corresponding mixed security policy.

Mixed Security Policies and Saddle-Point Equilibrium

Inequalities in property **P4.3** are straightforward to prove.

Start with the one in the left-hand side:

$$\underline{V}_m(A) := \max_{z\in\mathcal{Z}} \min_{y\in\mathcal{Y}} y'AZ \geq \max_{z\in\{e_1,e_2,\ldots,e_n\}} \min_{y\in\mathcal{Y}} y'Az = \max_j \min_{y\in\mathcal{Y}} y' \underbrace{Ae_j}_{j\text{th column of } A}$$

where $\{e_1, e_2, \ldots, e_n\} \subset \mathcal{Z}$ are the canonical basis of $\mathbb{R}^n$.

But then

$$\underline{V}_m(A) \geq \max_j \min_{y\in\mathcal{Y}} y' \begin{bmatrix} a_{1j} \\ a_{2j} \\ \vdots \\ a_{mj} \end{bmatrix} = \max_j \min_{y\in\mathcal{Y}} \sum_{i=1}^{m} y_i a_{ij}$$

Mixed Security Policies and Saddle-Point Equilibrium

An equality useful to optimize a linear function over a simplex:

$$\min_{y \in \mathcal{Y}} \sum_{i=1}^{m} y_i a_{ij} = \min_i a_{ij}$$

minimum is achieved by placing all the probability weight at the value of i for which the constant a_{ij} is the smallest.

We therefore conclude that

$$\underline{V}_m(A) \geq \max_j \min_i a_{ij} =: \underline{V}(A)$$

which is the left-most inequality in **P4.3**.

The right-most inequality can be proved in an analogous fashion, but starting with $\bar{V}_m(A)$, instead of $\underline{V}_m(A)$.

Mixed Security Policies and Saddle-Point Equilibrium

The middle inequality can be proved in exactly the same way that we used to prove that $\underline{V}(A) \geq \bar{V}(A)$

$$\underline{V}_m(A) = \underbrace{\min_y y' A z^*}_{\text{where } z^* \text{ is a mixed security policy}} \leq \min_y \max_z y' A z =: \bar{V}_m(A)$$

Mixed Security Policies and Saddle-Point Equilibrium

Note 5. To prove **P4.1**, apply **Weierstrass' Theorem** twice

1. Note $y'Az$ is a continuous function of y (for each fixed z), which is being minimized over the compact set $\mathcal{Y}$.
   By Weierstrass Theorem, there exists a $y \in \mathcal{Y}$ at which the minimum is achieved, i.e.,

   $$y^{*\prime}Az = \min_{y\in\mathcal{Y}} y'Az =: f(z)$$

   Value of the minimum $f(z)$ is a continuous function of z.
2. Use the fact that $f(z)$ is continuous and is being maximized over the compact set $\mathcal{Z}$.
   By Weierstrass Theorem, there exists a $z \in \mathcal{Z}$ at which the maximum is achieved, i.e.,

   $$f(z^*) = \max_{z\in\mathcal{Z}} f(z) = \max_{z\in\mathcal{Z}} \min_{y\in\mathcal{Y}} y'Az$$

Mixed Security Policies and Saddle-Point Equilibrium

Note 6.

Lemma 4.1 (Optimization of linear functions over simplexes).

Consider a probability simplex

$$\mathcal{X} := \left\{ x \in \mathbb{R}^m : \sum_i x_i = 1, \quad x_i \geq 0, \quad \forall i \right\}$$

and a linear function f of the form

$$f(x) = \sum_{i=1}^{m} a_i x_i$$

Then

$$\min_{x \in \mathcal{X}} f(x) = \min_{i \in \{1,2,\ldots,m\}} a_i, \qquad \max_{x \in \mathcal{X}} f(x) = \max_{i \in \{1,2,\ldots,m\}} a_i$$

Mixed Security Policies and Saddle-Point Equilibrium

The existence of a mixed saddle-point equilibrium is closely related to the average security levels for the two players.

Theorem 4.1 (Mixed saddle-point equilibrium vs. security levels).

A matrix game defined by A has a mixed saddle-point equilibrium **if and only if**

$$\underline{V}_m(A) := \max_{z \in \mathcal{Z}} \min_{y \in \mathcal{Y}} y' A z = \min_{y \in \mathcal{Y}} \max_{z \in \mathcal{Z}} y' A z =: \bar{V}_m(A)$$

Notation. In short, the min and max commute.

Mixed Security Policies and Saddle-Point Equilibrium

In particular

1. if (y^*, z^*) is a mixed saddle-point equilibrium then y^* and z^* are mixed security policies for P_1 and P_2, respectively and the equation is equal to the mixed saddle-point value.
2. if the equation holds and y^* and z^* are mixed security policies for P_1 and P_2, then (y^*, z^*) is a mixed saddle-point equilibrium and its value is equal to the equation.

Consequence: all mixed saddle-point equilibria must have the same mixed saddle-point values, which is called the **value of the game**, and denoted by $V_m(A)$.

Key difference between pure and mixed policies

- the equation in Theorem 4.1 holds for every matrix A.

This fact is known as the **Minimax Theorem**.

Mixed Security Policies and Saddle-Point Equilibrium

Example 4.1 (Network routing game).

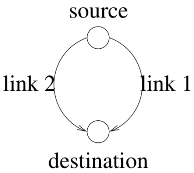

A matrix game between the router P_1 and the attacker P_2 if we make the following associations

$$P_1\text{'s actions:} \begin{cases} \text{send packet through link } 1 \equiv 1 \\ \text{send packet through link } 2 \equiv 2 \end{cases} \qquad P_2\text{'s actions:} \begin{cases} \text{attack link } 1 \equiv 1 \\ \text{attack link } 2 \equiv 2 \end{cases}$$

$$\text{outcomes:} \begin{cases} \text{packet arrives} & \equiv -1 & P_1 \text{ wins} \\ \text{packet is intercepted} & \equiv +1 & P_2 \text{ wins} \end{cases}$$

Mixed Security Policies and Saddle-Point Equilibrium

The associatons lead to the following matrix

$$A = \underbrace{\left.\begin{bmatrix} +1 & -1 \\ -1 & +1 \end{bmatrix}\right\}}_{P_2 \text{ choices}} P_1 \text{ choices}$$

with pure security levels

$$\underline{V}(A) = -1, \qquad \bar{V}(A) = +1,$$

showing there are no saddle-point equilibria in pure policies.

For mixed policies, we have (using **Lemma 4.1**)

$$\begin{aligned} \underline{V}_m(A) = \max_{z\in\mathcal{Z}} \min_{y\in\mathcal{Y}} y'Az &= \max_{z\in\mathcal{Z}} \min_{y\in\mathcal{Y}} y_1(z_1 - z_2) + y_2(z_2 - z_1) \\ &= \max_{z\in\mathcal{Z}} \min\{z_1 - z_2, z_2 - z_1\} \end{aligned}$$

Mixed Security Policies and Saddle-Point Equilibrium

To compute the maximum over z, note that

$$\min\{z_1 - z_2, z_2 - z_1\} \begin{cases} = 0 & z_1 = z_2 \\ < 0 & z_1 \neq z_2 \end{cases}$$

The maximum over z is obtained for $z_1 = z_2 = \frac{1}{2}$, leading to

$$\underline{V}_m(A) = 0$$

with a mixed security policy for P_2 given by $z^* := \begin{bmatrix} \frac{1}{2} & \frac{1}{2} \end{bmatrix}$.

To compute the minimum over y, we have

$$\begin{aligned} \bar{V}_m(A) = \min_{y \in \mathcal{Y}} \max_{z \in \mathcal{Z}} y' A z &= \min_{y \in \mathcal{Y}} \max_{z \in \mathcal{Z}} z_1(y_1 - y_2) + z_2(y_2 - y_1) \\ &= \min_{y \in \mathcal{Y}} \max\{y_1 - y_2, y_2 - y_1\} = 0 \end{aligned}$$

with a mixed security policy for P_1 given by $y^* := \begin{bmatrix} \frac{1}{2} & \frac{1}{2} \end{bmatrix}$.

Mixed Security Policies and Saddle-Point Equilibrium

The result shows that

$$\underline{V}_m(A) = \bar{V}_m(A) = 0$$

Conlusion: (From **Theorem 4.1**) this game has a mixed saddle-point equilibrium (y^*, z^*), which consists of each player selecting each of the links with equal probabilities.

Mixed Security Policies and Saddle-Point Equilibrium

Example 4.2 (Rock-paper-scissors). Consider a matrix representation of the rock-paper-scissors game

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

For this game $\underline{V}(A) = -1$ and $\bar{V}(A) = 1$.

For mixed policies, using **Lemma 4.1**, we conclude that

$$\begin{aligned} \underline{V}_m(A) &= \max_{z \in \mathcal{Z}} \min_{y \in \mathcal{Y}} y' A z \\ &= \max_{z \in \mathcal{Z}} \min_{y \in \mathcal{Y}} y_1(z_2 - z_3) + y_2(z_3 - z_1) + y_3(z_1 - z_2) \\ &= \max_{z \in \mathcal{Z}} \min \{z_2 - z_3, z_3 - z_1, z_1 - z_2\} \end{aligned}$$

Mixed Security Policies and Saddle-Point Equilibrium

$$\underline{V}_m(A) = \max_{z \in \mathcal{Z}} \min \{z_2 - z_3, z_3 - z_1, z_1 - z_2\}$$

which is maximized with $z_1 = z_2 = z_3 = \frac{1}{3}$. This means that

$$\underline{V}_m(A) = 0$$

corresponding to a mixed security policy for P_2 given by $z^* := \begin{bmatrix} \frac{1}{3} & \frac{1}{3} & \frac{1}{3} \end{bmatrix}$. Similarly,

$$\bar{V}_m(A) = \min_{y \in \mathcal{Y}} \max_{z \in \mathcal{Z}} y' A z = \cdots = 0$$

with a mixed security policy for P_1 given by $y^* := \begin{bmatrix} \frac{1}{3} & \frac{1}{3} & \frac{1}{3} \end{bmatrix}$

Mixed Security Policies and Saddle-Point Equilibrium

Note. Why? To get the maximum over z we must have $z_2 \geq z_3$ since otherwise $z_2 - z_3 < 0$ and we would get the minimum smaller than zero. For similar reasons, we must also have $z_3 \geq z_1$ and $z_1 \geq z_2$. The only way to satisfy these three inequalities simultaneously is to have all z_j equal to each other.

The result shows that

$$\underline{V}_m(A) = \bar{V}_m(A) = 0$$

This game has a mixed saddle-point equilibrium (y^*, z^*)

- each player selecting rock, paper, or schissors with equal probabilities.

Note. A lot of work to conclude what every 7 year old learns in the school yard.

General Zero-Sum Games

General Zero-Sum Games

Consider a two-player zero-sum game G.
P_1 and P_2 select policies within **action spaces** Γ_1 and Γ_2.

For a pair of policies $\gamma \in \Gamma_1$, $\sigma \in \Gamma_2$, $J(\gamma, \sigma)$ is the **outcome of the game** when P_1 uses policy γ and P_2 uses policy σ.

P_1 wants to minimize outcome $J(\gamma, \sigma)$, P_2 wants to maximize it.

Note: we now allow the outcome J to depend on the policies in an arbitrary fashion.

- we need to adapt previous definitions of security policies and levels.

General Zero-Sum Games

Definition 4.3 (Security policy). The **security level** for P_1 (the minimizer) is defined by

$$\bar{V}_{\Gamma_1,\Gamma_2} := \underbrace{\inf_{\gamma\in\Gamma_1}}_{\substack{\text{minimize cost assuming}\\ \text{worst choice by } P_2}} \quad \underbrace{\sup_{\sigma\in\Gamma_2}}_{\substack{\text{worst choice by } P_2\\ \text{from } P_1\text{'s perspective}}} \quad J(\gamma,\sigma)$$

Security policy for P_1

- any policy $\gamma^* \in \Gamma_1$ for which the infimum is achieved, i.e.,

$$\underbrace{\sup_{\sigma\in\Gamma_2} J(\gamma^*,\sigma) = \bar{V}_{\Gamma_1,\Gamma_2}(G)}_{\gamma^* \text{ achieves the infimum}} := \inf_{\gamma\in\Gamma_1} \sup_{\sigma\in\Gamma_2} J(\gamma,\sigma)$$

General Zero-Sum Games

Notation.
- The **infimum** of a set is its largest lower bound
- The **supremum** of a set is its smallest upper bound.

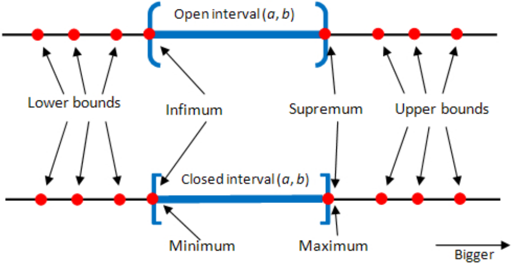

General Zero-Sum Games

The **security level** for P_2 (the maximizer) is defined by

$$\underline{V}_{\Gamma_1,\Gamma_2} := \underbrace{\sup_{\sigma\in\Gamma_2}}_{\substack{\text{maximize rewards assuming}\\ \text{worst choice by } P_1}} \underbrace{\inf_{\gamma\in\Gamma_1}}_{\substack{\text{worst choice by } P_1\\ \text{from } P_2\text{'s perspective}}} J(\gamma,\sigma)$$

Security policy for P_2

- any policy $\sigma^* \in \Gamma_2$ for which the supremum is achieved, i.e.,

$$\underbrace{\sup_{\gamma\in\Gamma_1} J(\gamma,\sigma^*) = \underline{V}_{\Gamma_1,\Gamma_2}(G)}_{\sigma^* \text{ achieves the maximum}} := \sup_{\sigma\in\Gamma_2} \inf_{\gamma\in\Gamma_1} J(\gamma,\sigma)$$

General Zero-Sum Games

Definition 4.4 (Saddle-point equilibrium).

A pair of policies $(\gamma^*, \sigma^*) \in \Gamma_1 \times \Gamma_2$ is called a **saddle-point equilibrium** if

$$\begin{aligned} J(\gamma^*, \sigma^*) &\leq J(\gamma, \sigma^*), &\qquad \forall \gamma \in \Gamma_1 \\ J(\gamma^*, \sigma^*) &\geq J(\gamma^*, \sigma), &\qquad \forall \sigma \in \Gamma_2 \end{aligned}$$

and $J(\gamma^*, \sigma^*)$ is the **saddle-point value**.

These equations are often re-written as

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

General Zero-Sum Games

Notation 1 (Infimum and supremum). The **infimum** of a set $\mathcal{S} \subset \mathbb{R}$ is the largest number y^* such that

$$y^* \leq s, \qquad \forall s \in \mathcal{S}$$

When there is no number y^* that satisfies the equation, the infimum is defined to be $-\infty$. Under this definition every set $\mathcal{S} \subset \mathbb{R}$ has an infimum, but the infimum may not belong to the set. When the infimum does belong to the set it is said to be a **minimum**. For example,

$$\inf\left\{e^{-x} : x \geq 0\right\} = 0, \quad \inf\left\{e^{-x} : x \leq 0\right\} = \min\left\{e^{-x} : x \leq 0\right\} = 1$$

Note that the infimum of the set in the left hand side is not a minimum since 0 does not belong to the set, but the infimum of the set in the right hand side does belong to the set.

General Zero-Sum Games

The **supremum** of a set $\mathscr{L} \subset \mathbb{R}$ is the smallest number z^* such that

$$z^* \geq s, \qquad \forall s \in \mathcal{S}$$

When there is no number z^* that satisfies the equation, the supremum is defined to be $+\infty$. Under this definition every set $\mathcal{S} \subset \mathbb{R}$ has a supremum. When the supremum does belong to the set it is said to be a **maximum**.

General Zero-Sum Games

Proposition 4.2 (Security levels/policies/saddle-point equilibria). These properties hold for every zero-sum game:

P4.4 Security levels are well defined and unique.

Note. However, security levels may take infinite values (including $-\infty$ or $+\infty$) and security policies may not exist.

P4.5 The security levels always satisfy the following inequality:

$$\underline{V}_{\Gamma_1,\Gamma_2}(G) := \sup_{\sigma\in\Gamma_2} \inf_{\gamma\in\Gamma_1} J(\gamma,\sigma) \leq \inf_{\gamma\in\Gamma_1} \sup_{\sigma\in\Gamma_2} J(\gamma,\sigma) =: \bar{V}_{\Gamma_1,\Gamma_2}(G)$$

P4.6 When $\underline{V}_{\Gamma_1,\Gamma_2}(G) = \bar{V}_{\Gamma_1,\Gamma_2}(G)$ and there exist security policies γ^* and σ^* for P_1 and P_2, then(σ^*,γ^*) is a saddle-point equilibrium.
Its value $J(\sigma^*,\gamma^*)$ is equal to $\underline{V}_{\Gamma_1,\Gamma_2}(G) = \bar{V}_{\Gamma_1,\Gamma_2}(G)$.

General Zero-Sum Games

P4.7 If there exists a saddle-point equilibrium (σ^*, γ^*) then γ^* and σ^* are security policies for P_1 and P_2, respectively and $\underline{V}_{\Gamma_1,\Gamma_2}(G) = \bar{V}_{\Gamma_1,\Gamma_2}(G) = J(\sigma^*, \gamma^*)$.

P4.8 If (σ_1^*, γ_1^*) and (σ_2^*, γ_2^*) are both saddle-point equilibria then (σ_1^*, γ_2^*) and (σ_2^*, γ_1^*) are also saddle-point equilibria. All these equilibria have exactly the same value.

Notation. A consequence of this is that all saddle-point equilibria must have exactly the same saddle-point values, which is called the value of the game and is denoted by V_{Γ_1,Γ_2}.

General Zero-Sum Games

Example 4.3 (Resistive circuit design). Recall the robust design problem for a resistive circuit discussed in Section 1.3.

- P_1: **designer**. Picks the nominal resistance R_{nom} to minimize the current error

$$e(R_{\text{nom}}, \delta) = \left| \frac{1}{R} - 1 \right| = \left| \frac{1}{(1+\delta)R_{\text{nom}}} - 1 \right|$$

- P_2: **nature**. Picks value of $\delta \in [-0.1, 0.1]$ to maximize e.

We saw that the policy

$$\pi_1^* : P_1 \text{ selects } R_{\text{nom}} = \frac{100}{99}$$

is a security policy for P_1 and leads to a security level $\bar{V} = 0.1$

- for any value of $\delta \in [-0.1, 0.1]$, P_1 could get the error to be equal to 0 by an appropriate choice of R_{nom}.

General Zero-Sum Games

The security level for P_2 is $\underline{V} = 0$ and any policy is a security policy for P_2. In view of the fact that

$$\underline{V} = 0 < \bar{V} = 0.1$$

Conclusion: game does not have a saddle-point equilibrium.

Suppose resistors are being drawn from a box that contains

$$\begin{cases} 45\% \text{ of the resistors with } \delta = -0.1 \\ 55\% \text{ of the resistors with } \delta = +0.1 \end{cases}$$

When P_1 uses the policy π_1 and a resistor is drawn randomly from the box, the expected value of the error is

$$E[e] = 0.45 \times \left| \frac{1}{(1-0.1)\frac{100}{99}} - 1 \right| + 0.55 \times \left| \frac{1}{(1+0.1)\frac{100}{99}} - 1 \right| = 0.1$$

General Zero-Sum Games

We can view this box of resistors as P_2's mixed policy, by imagining that nature picked the distribution of resistors in the box in order to maximize the expected value of the error. This corresponds to the following mixed policy for P_2

$$\pi_2^* : P_2 \text{ selects } \begin{cases} -0.1 & \text{with probability } 0.45 \\ +0.1 & \text{with probability } 0.55 \end{cases}$$

We can state the following:

1. Restricting to pure policies, the security levels are

$$\underline{V} = 0 < \bar{V} = 0.1$$

The policy π_1^* is a security policy for P_1 and any policy for P_2 is a security policy for this player. The gap between the security levels indicates that there are no pure saddle points.

General Zero-Sum Games

2. Enlarging the universe of policies to consider mixed policies, the average security levels of the game become

$$\underline{V}_m = \bar{V}_m = 0.1$$

Observation: we allowed P_2 to raise her security level to 0.1.

The policy π_1^* is still a mixed security policy for P_1, but now π_2^* is a mixed security policy for P_2.

The equality between the security levels indicates that (π_1^*, π_2^*) is a mixed saddle-point equilibrium.

To regard π_1^* as a mixed policy, think of this policy as

- the probability distribution for R_{nom} that places all probability mass at $\frac{100}{99}$.

General Zero-Sum Games

The fact that $R_{\text{nom}} = \frac{100}{99}$ is part of a mixed saddle-point equilibrium, tells us that this design is not conservative in the sense that:

1. selection of $R_{\text{nom}} = \frac{100}{99}$ guarantees that the current error remains below 0.1 for any resistor R with error below 10%, and

2. there exists a population of resistors with error below 10%. Extracting R randomly from this population, we get $E[e] = 0.1$.

If the security policy $R_{\text{nom}} = \frac{100}{99}$ was not a mixed saddle-point equilibrium then there would be no underlying distribution of resistors that could lead to $E[e] = 0.1$.

- indicating that by choosing $R_{\text{nom}} = \frac{100}{99}$ we were **protecting** our design against a **phantom** worst-case distribution of resistances that actually did not exist.

Practice Exercises

Practice Exercises

4.1 (Resistive circuit design). Consider the robust design problem. Show that the average security levels of this game are

$$\underline{V}_m = \bar{V}_m = 0.1$$

and that the policies (π_1^*, π_2^*) defined by

$$\pi_1^* : P_1 \text{ selects } R_{\text{nom}} = \frac{100}{99}$$

$$\pi_2^* : P_2 \text{ selects } \begin{cases} -0.1 & \text{with probability } 0.45 \\ +0.1 & \text{with probability } 0.55 \end{cases}$$

form a mixed saddle-point equilibrium.

Hint: Show that π_1^* is a pure security policy for P_1, leading to $\bar{V} = 0.1$. Next, verify that π_1^* and π_2^* satisfy the general conditions for a mixed saddle-point equilibrium.

Practice Exercises

Solution to Exercise 4.1.

To show that $R_{\text{nom}} = \frac{100}{99}$ is a pure security policy for P_1 with security level equal to 0.1, we start by computing

$$\bar{V} = \min_{R_{\text{nom}} \geq 0} \max_{\delta \in [-0.1, 0.1]} \left| \frac{1}{(1+\delta) R_{\text{nom}}} - 1 \right|$$

$$= \min_{R_{\text{nom}} \geq 0} \max_{\delta \in [-0.1, 0.1]} \begin{cases} \frac{1}{(1+\delta) R_{\text{nom}}} - 1 & \delta \leq \frac{1}{R_{\text{nom}} - 1} \\ 1 - \frac{1}{(1+\delta) R_{\text{nom}}} & \delta > \frac{1}{R_{\text{nom}} - 1} \end{cases}$$

Top branch is monotone decreasing with respect to δ.
Bottom branch is monotone increasing with respect to δ.

Inner maximization is achieved at the extreme points for δ.

Practice Exercises

We conclude that

$$\bar{V} = \min_{R_{\text{nom}} \geq 0} \begin{cases} \max\left\{\frac{1}{(1-0.1)R_{\text{nom}}} - 1, 1 - \frac{1}{(1+0.1)R_{\text{nom}}}\right\} & -0.1 \leq \frac{1}{R_{\text{nom}}} - 1, 0.1 > \frac{1}{R_{\text{nom}}} - 1 \\ \frac{1}{(1-0.1)R_{\text{nom}}} - 1 & -0.1 \leq \frac{1}{R_{\text{nom}}} - 1, 0.1 \leq \frac{1}{R_{\text{nom}}} - 1 \\ 1 - \frac{1}{(1+0.1)R_{\text{nom}}} & -0.1 > \frac{1}{R_{\text{nom}}} - 1, 0.1 > \frac{1}{R_{\text{nom}}} - 1 \end{cases}$$

$$= \min_{R_{\text{nom}} \geq 0} \begin{cases} \max\left\{\frac{1}{(1-0.1)R_{\text{nom}}} - 1, 1 - \frac{1}{(1+0.1)R_{\text{nom}}}\right\} & R_{\text{nom}} \in \left(\frac{10}{11}, \frac{10}{9}\right] \\ \frac{1}{(1-0.1)R_{\text{nom}}} - 1 & R_{\text{nom}} \leq \frac{10}{11} \\ 1 - \frac{1}{(1+0.1)R_{\text{nom}}} & R_{\text{nom}} > \frac{10}{9} \end{cases}$$

Considering the three branches separately, optimum is achieved at the top branch when

$$\frac{1}{(1-0.1)R_{\text{nom}}} - 1 = 1 - \frac{1}{(1+0.1)R_{\text{nom}}} \Leftrightarrow R_{\text{nom}} = \frac{100}{99}$$

This shows that π_1^* is a security policy for P_1, corresponding to a security level of $\bar{V} = 0.1$

Practice Exercises

To prove (π_1, π_2) is a mixed saddle-point equilibrium, show that π_1^* and π_2^* satisfy the required conditions.

Since π_1^* is a pure security policy for P_1 with security level equal to 0.1, when P_1 uses this policy the error in the current satisfies

$$e\left(\frac{100}{99}, \delta\right) = \left|\frac{1}{(1+\delta)\frac{100}{99}} - 1\right| \leq 0.1 \qquad \forall \delta \in [-0.1, 0.1]$$

This means that, for any mixed policy π_2^* for P_2, we must have

$$E_{\pi_1^*, \pi_2}\left[e\left(\frac{100}{99}, \delta\right)\right] \leq 0.1, \qquad \forall \pi_2$$

where π_1^*, π_2^* means we are fixing the (pure) policy π_1 and taking δ to be a random variable with distribution π_2.

Practice Exercises

Suppose we fix P_2's policy to be π_2^* and consider an arbitrary policy π_1 for P_1, corresponding to a distribution for R_{nom} with probability density function $f(r)$. Then

$$\begin{aligned}
E_{\pi_1,\pi_2^*} =& 0.45\int_0^{\infty} e(r,-0.1)f(r)dr + 0.55\int_0^{\infty} e(r,0.1)f(r)dr\\
=& \int_0^{\frac{10}{9}} 0.45\left(\frac{1}{(1-0.1)r}-1\right) fr(dr) + \int_{\frac{10}{9}}^{\infty} 0.45\left(1-\frac{1}{(1-0.1)r}\right) f(r)dr\\
&+ \int_0^{\frac{10}{11}} 0.55\left(\frac{1}{(1+0.1)r}-1\right) fr(dr) + \int_{\frac{10}{11}}^{\infty} 0.55\left(1-\frac{1}{(1+0.1)r}\right) f(r)dr\\
=& \int_0^{\frac{10}{11}} \left(0.45\left(\frac{1}{(1-0.1)r}-1\right) + 0.55\left(\frac{1}{(1+0.1)r}-1\right)\right) f(r)dr\\
&+ \int_{\frac{10}{11}}^{\frac{10}{9}} \left(0.45\left(\frac{1}{(1-0.1)r}-1\right) + 0.55\left(1-\frac{1}{(1+0.1)r}\right)\right) f(r)dr\\
&+ \int_{\frac{10}{9}}^{\infty} \left(0.45\left(1-\frac{1}{(1-0.1)r}\right) + 0.55\left(1-\frac{1}{(1+0.1)r}\right)\right) f(r)dr
\end{aligned}$$

Practice Exercises

$$E_{\pi_1,\pi_2^*} = \int_0^{\frac{10}{11}} \left(\frac{1}{r} - 1\right) f(r)dr + \int_{\frac{10}{11}}^{\frac{10}{9}} \frac{1}{10} f(r)dr + \int_{\frac{10}{9}}^{\infty} \left(1 - \frac{1}{r}\right) f(r)dr$$

$$\text{Since} \qquad r \leq \frac{10}{11} \quad \Rightarrow \quad \frac{1}{r} - 1 \geq \frac{11}{10} - 1 = \frac{1}{10}$$

$$\text{And} \qquad r \geq \frac{10}{9} \quad \Rightarrow \quad 1 - \frac{1}{r} \geq 1 - \frac{9}{10} = \frac{1}{10}$$

Conclusion: E_{π_1,π_2^*} is minimized by selecting the probability mass for the distribution π_1 to be in interval $[\frac{10}{11}, \frac{10}{9}]$, leading to

$$\max_{\pi} E_{\pi_1,\pi_2^*} = \frac{1}{10} = 0.1 = E_{\pi_1^*,\pi_2^*}$$

This and $E_{\pi_1^*,\pi_2}$ establish that π_1^* and π_2^* satisfy the conditions for a saddle-point equilibrium.

End of Lecture

04 - Mixed Policies

Questions?