Game Theory

Lecture 2: Static games with complete information – the normal form

Frieder Neunhoeffer

Lisbon School of Economics & Management
Universidade de Lisboa

Static games with complete information

- **Static games:** The players move simultaneously, so that none can observe and react to the choices made by the other players.

- **Complete information**: Each player knows fully the rules of the game. In particular, each one knows the "payoff functions" of all players.

- **Payoff function** of a player gives the reward the player gets as a result of the combination of strategies selected by all players.

Normal form of a game with pure strategies

Pure strategies are basic alternatives of an action.

- The **normal** or **"strategic" form** of a game contains three elements:

  1. A set of players $I = \{1,2,\dots,n\}$, where n is a positive integer.
  2. For any player $i \in I,\ S_i$ is the set of pure strategies that are available to him.
     - S_i can be either **finite**: $S_i = \{1,2,\dots,m_i\}$ for any integer $m_i \geq 2$.
     - Or S_i can be a **continuum**, where it takes the form of an *interval* $[a,b]$.
  3. Let $s = (s_1, s_2, \dots, s_i)$ be the vector or profile of strategies that are selected by all the players. Then, for any s and each player $i \in I, \pi_i(s)$ is the payoff function, that gives the reward (amount of utility or money) for player i, stemming from the combination of strategies s.

Normal form of a game with pure strategies

- Several definitions are useful:
  - The space of pure strategies of the game is $S \equiv \times_i S_i$. It contains all possible values of s.
  - The combined payoff function of the game is $\pi(s) \equiv (\pi_1(s), \pi_2(s), \dots, \pi_n(s))$

- **Definition of normal form of a game.** If only pure strategies are taken into account, any game can be defined by a triplet:

  $G \equiv (I, S, \Pi)$ where

  I is the set of players
  S is the space of pure strategies
  Π is the combined payoff function

Special case: 2 player finite games

- If there are only two players, 1 and 2, endowed with m_1 and m_2 discrete strategies respectively, the **normal form** can be written as a matrix $m_1 \times m_2$.
  - Each row h of the matrix represents a pure strategy $h \in S_1$, available to player 1.
  - Each column k of the matrix stands for a pure strategy $k \in S_2$, available to player 2.
  - The sub-matrix $A = (a_{hk}) = \pi_1(h, k)$ contains the payoffs of player 1.
  - The sub-matrix $B = (b_{hk}) = \pi_2(h, k)$ contains the payoffs of player 2.

- **Remark:** Usually, sub-matrices A and B are consolidated in a single matrix (A, B), that contains payoffs a_{hk}, b_{hk} in each matrix cell.

Example: Prisoner’s dilemma

- The payoff (years in prison) matrices of this game are:

$$A = \begin{pmatrix} 3 & 0 \\ 5 & 1 \end{pmatrix}, B = \begin{pmatrix} 3 & 5 \\ 0 & 1 \end{pmatrix}, (A, B) = \begin{pmatrix} 3,3 & 0,5 \\ 5,0 & 1,1 \end{pmatrix}$$

- Each player has two pure strategies: "Confess" and “Deny", respectively. The first strategy gives each player a higher payoff (fewer years in prison) than the second one, for any choice made by the other player. → Game theory predicts both select “Confess".
- If both chose “Deny" they get a payoff (1 year) strictly higher than the payoff that they achieve by choosing “Confess“ (3 years) → **dilemma**

Special case: 3 player finite games

- With three players and discrete strategy sets, it is possible to write the normal form using two matrices. In this case:
    1. The choice of player 1 is represented by the choice of the row.
    2. The choice of player 2 is represented by the choice of the column.
    3. The choice of player 3 is represented by the choice of the matrix.
- Each cell of each matrix contains three numbers, representing the payoffs of players 1, 2 and 3, respectively.
- <u>Example:</u>

$$\begin{pmatrix} 1,1,1 & 0,0,0 \\ 2,2,2 & 3,2,1 \end{pmatrix}, \begin{pmatrix} 3,4,5 & 1,1,1 \\ 6,1,0 & 0,0,0 \end{pmatrix}$$

Introducing mixed strategies

- **Definition:** A *mixed strategy* for player i is a distribution of probability over his set of pure strategies S_i. Let us assume that player i has m_i pure strategies. Then, a mixed strategy for player i can be represented by a vector x_i, whose generic element is p_{is}, the probability that player i assigns to playing pure strategy s.
- **Meaning of a mixed strategy:** Instead of choosing a basic alternative of action (i.e., a pure strategy), the player builds a random device (a coin, a dice or a roulette) that selects each pure strategy with a given probability. Then, he runs this mechanism before acting.
- Each mixed strategy of player i, x_i, has a *support*, labelled as $C(x_i)$, that is the set of pure strategies available to player i, to which he assigns positive probabilities, i.e.,

$$C(x_i) = \{s \in S_i : p_{is} > 0\}$$

Introducing mixed strategies

- Given the properties of probabilities, the set of mixed strategies of player i is (in geometric terms) the unit simplex in space m_i, Δ_i, as defined by:

$$\Delta_i \left\{ \mathrm{x}_i \in R_+^{m_i} \colon \sum_{s=1}^{m_i} x_{is} = 1 \right\}$$

- The set of mixed strategies of a player with **two pure strategies** is a line segment that connects points $\binom{0}{1}$ and $\binom{1}{0}$. In general, the **set of mixed strategies** has a **dimension** that is equal to the number of pure strategies minus 1, $\boldsymbol{m_i - 1}$. *This is so because each probability p_{is} can be written as 1 minus all the other probabilities*.

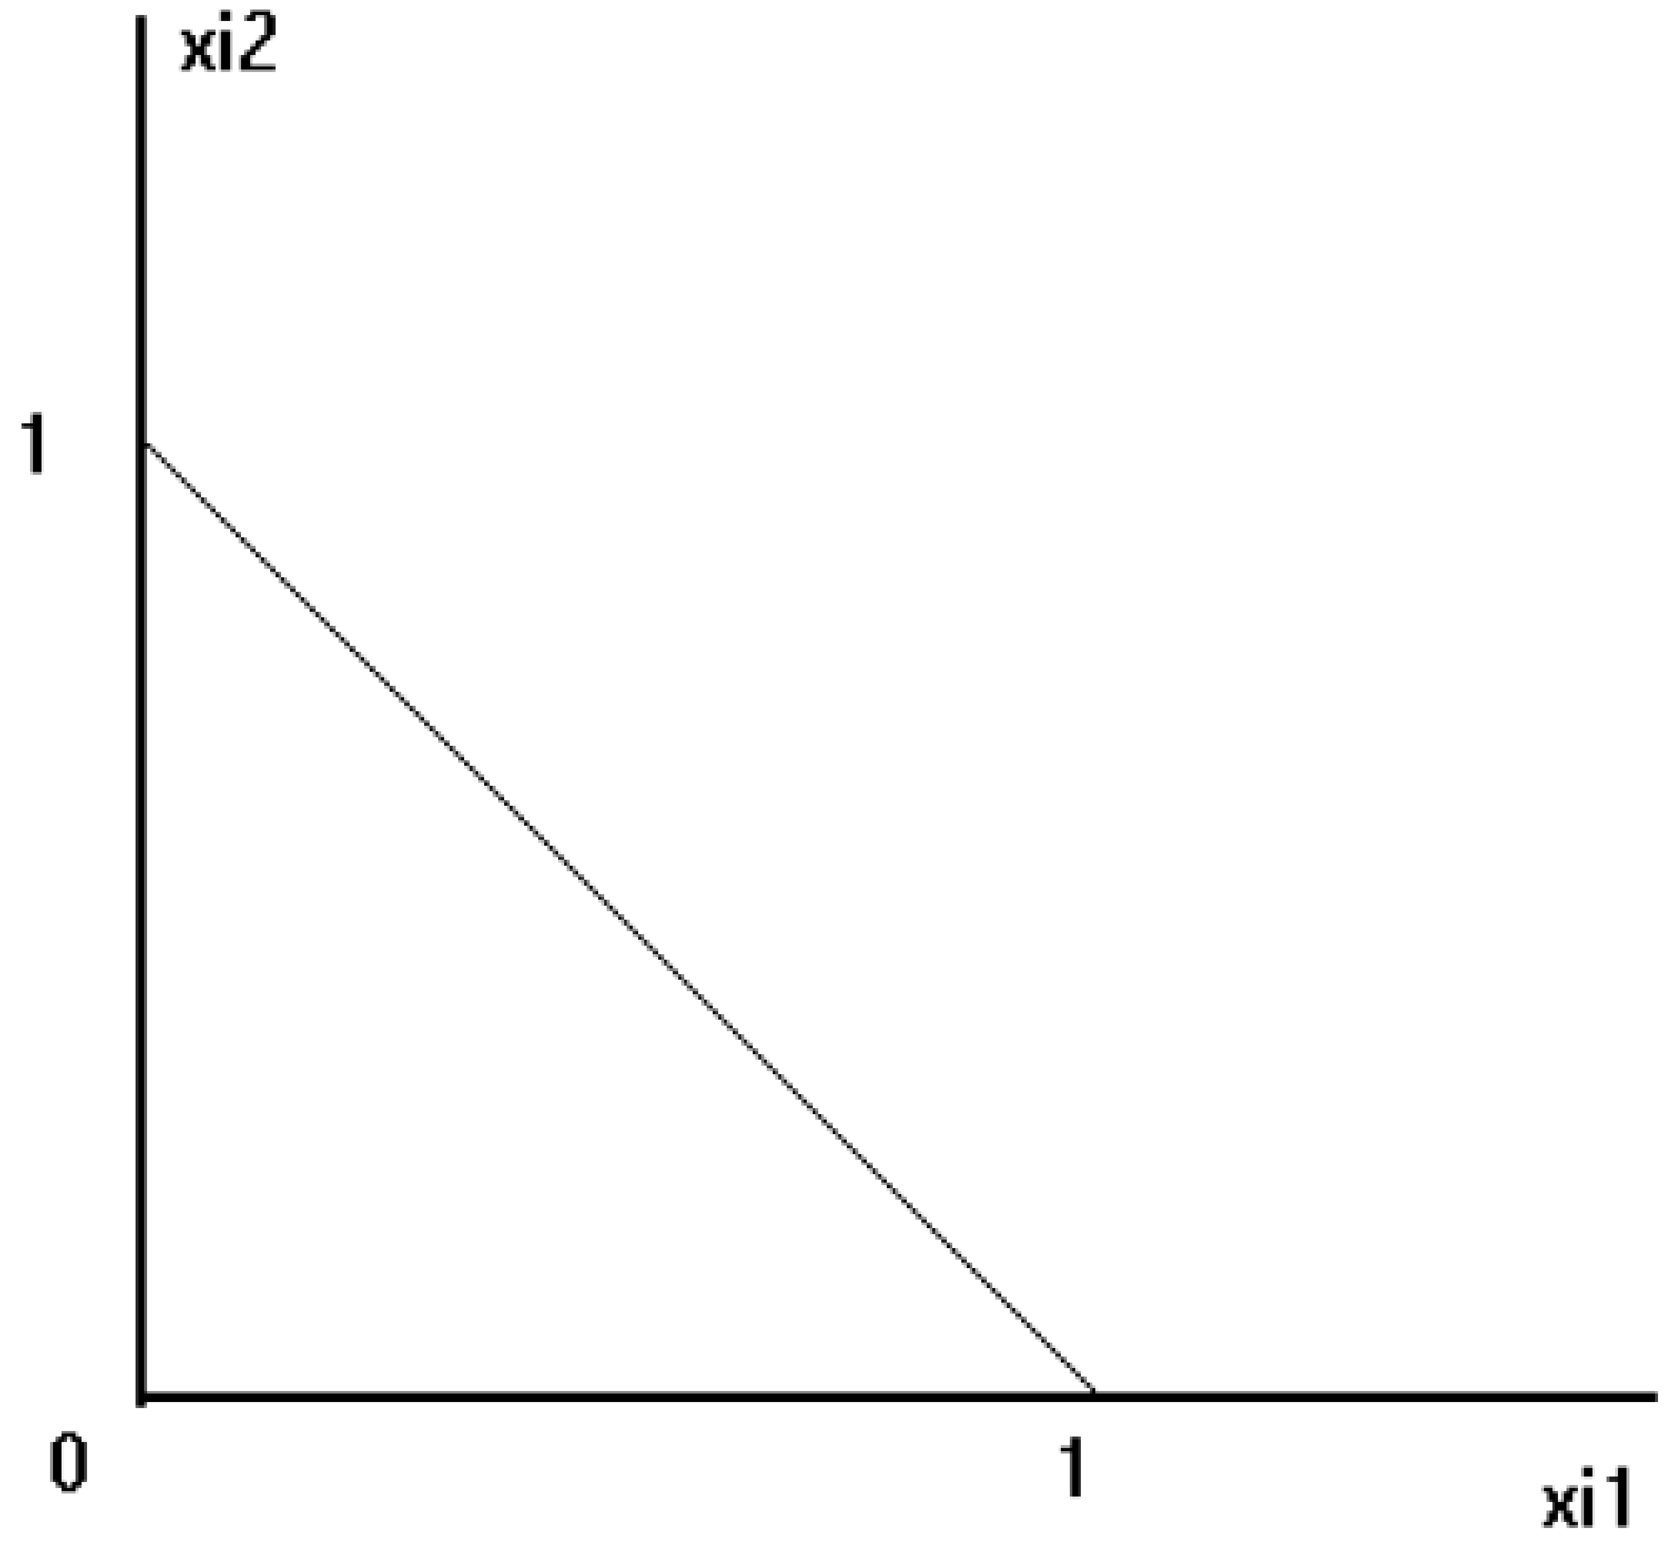

Definition of Projection

- Let us assume a space X of dimension n, whose elements are vectors $(x_1, x_2, \dots, x_n)$. A projection is a function that associates with each element $(x_1, x_2, \dots, x_n)$, a vector that is made up by a *part* of the initial coordinates (one dimension less). Hence, a *projection* $X \in R_n$ into R_{n-1} associates with an element $(x_1, x_2, \dots, x_n)$, an element of the form $(x_1, x_2, \dots, x_{n-1})$

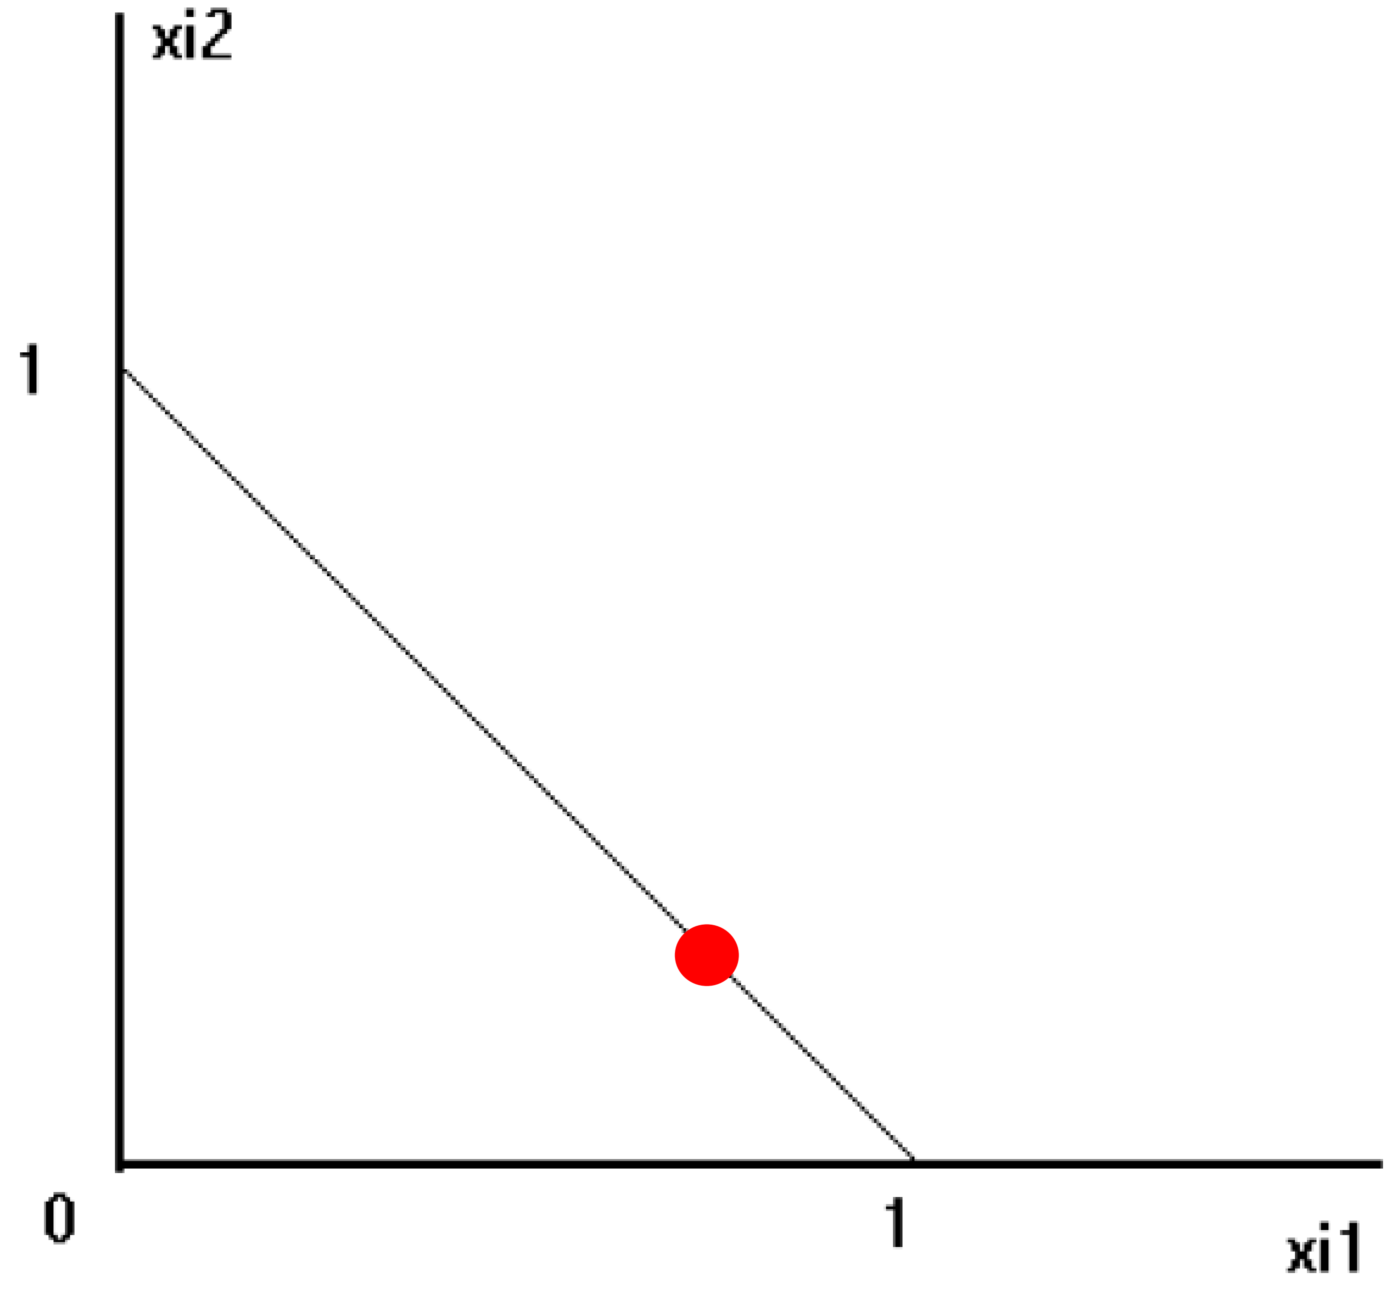

Definition of vertices (corners) of Δ_i

- Vertices (also called *corners*) of the set of mixed strategies of a player endowed with m_i pure strategies are the unit vectors in a space with dimension m_i:

$$e_i^1 = \begin{pmatrix} 1 \\ 0 \\ \dots \\ 0 \end{pmatrix}, e_i^2 = \begin{pmatrix} 0 \\ 1 \\ \dots \\ 0 \end{pmatrix}, \dots e_i^{m_i} = \begin{pmatrix} 0 \\ 0 \\ \dots \\ 1 \end{pmatrix}$$

- In the previous figures, the corners are points $\begin{pmatrix} 0 \\ 1 \end{pmatrix}$ and $\begin{pmatrix} 1 \\ 0 \end{pmatrix}$.
- **Meaning:** Each vertex represents the pure strategy s of player i. Consequently, the pure strategy s is nothing but a "special“ mixed strategy x_i that assigns the whole mass (i.e., 1) of probability to a single pure strategy s.
- Thus, each mixed strategy $x_i \in \Delta_i$ is a **convex** combination of its pure strategies e_i^s:

$$x_i = \sum_{s=1}^{m_i} x_{is} e_i^s;\ x_{is} \geq 0, \sum_{s=1}^{m_i} x_{is} = 1$$

Some more basic definitions

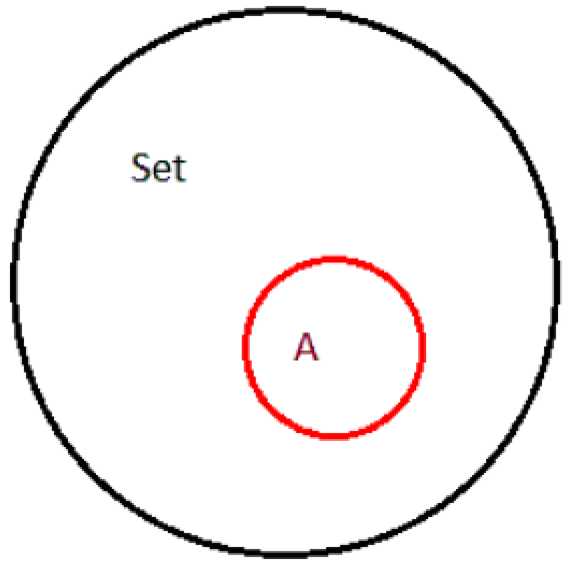

- An **interior point** has a neighborhood which is fully contained in the set. The interior of a set is the set of its interior points.
- A set is **open** if it is made up only of interior points, i.e., if it is coincident with its interior. Example: (1,2,3)

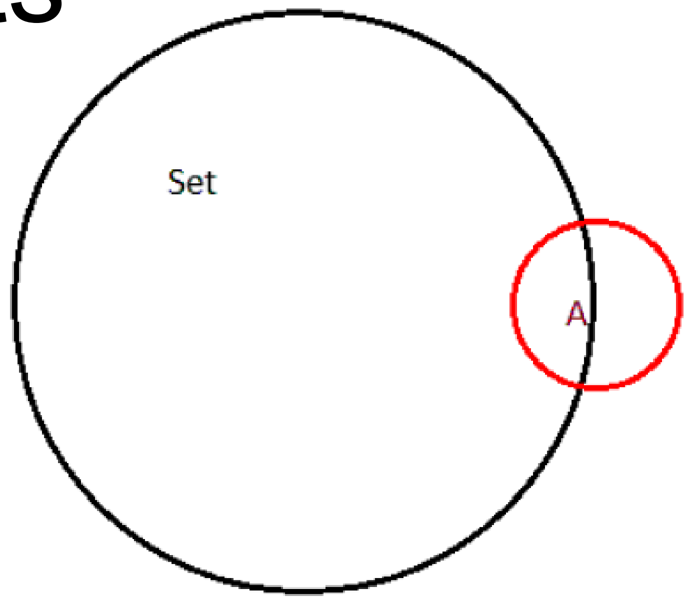

- An **boundary point** is a point for which any neighborhood contains both, points that belong and points that do not belong to the set. The **boundary of a set** is the set of its boundary points.
- A **closed set** is a set that owns its boundary. For instance, the set of mixed strategies of a player i, Δ_i, is a closed set. Example: [1,2,3]

Applied to mixed strategies

If we are concerned with the set of mixed strategies of a player i, Δ_i, then the above definitions have a more precise meaning:

- An interior point of Δ_i is a mixed strategy that is completely mixed, i.e., that assigns a positive probability to each pure strategy s available to player i.
- In our example the interior is coincident with the set of mixed strategies, leaving aside the boundary points or vertices $\binom{0}{1}$ and $\binom{1}{0}$.
- The boundary of the set of mixed strategies of player i is the set of strategies that assign positive probability only to some pure strategies, so that the support of the mixed strategy $x_i, C(x_i)$ is a proper subset of the set of pure strategies available to player i, S_i.

Mixed strategy set Δ_i when player i has three pure strategies

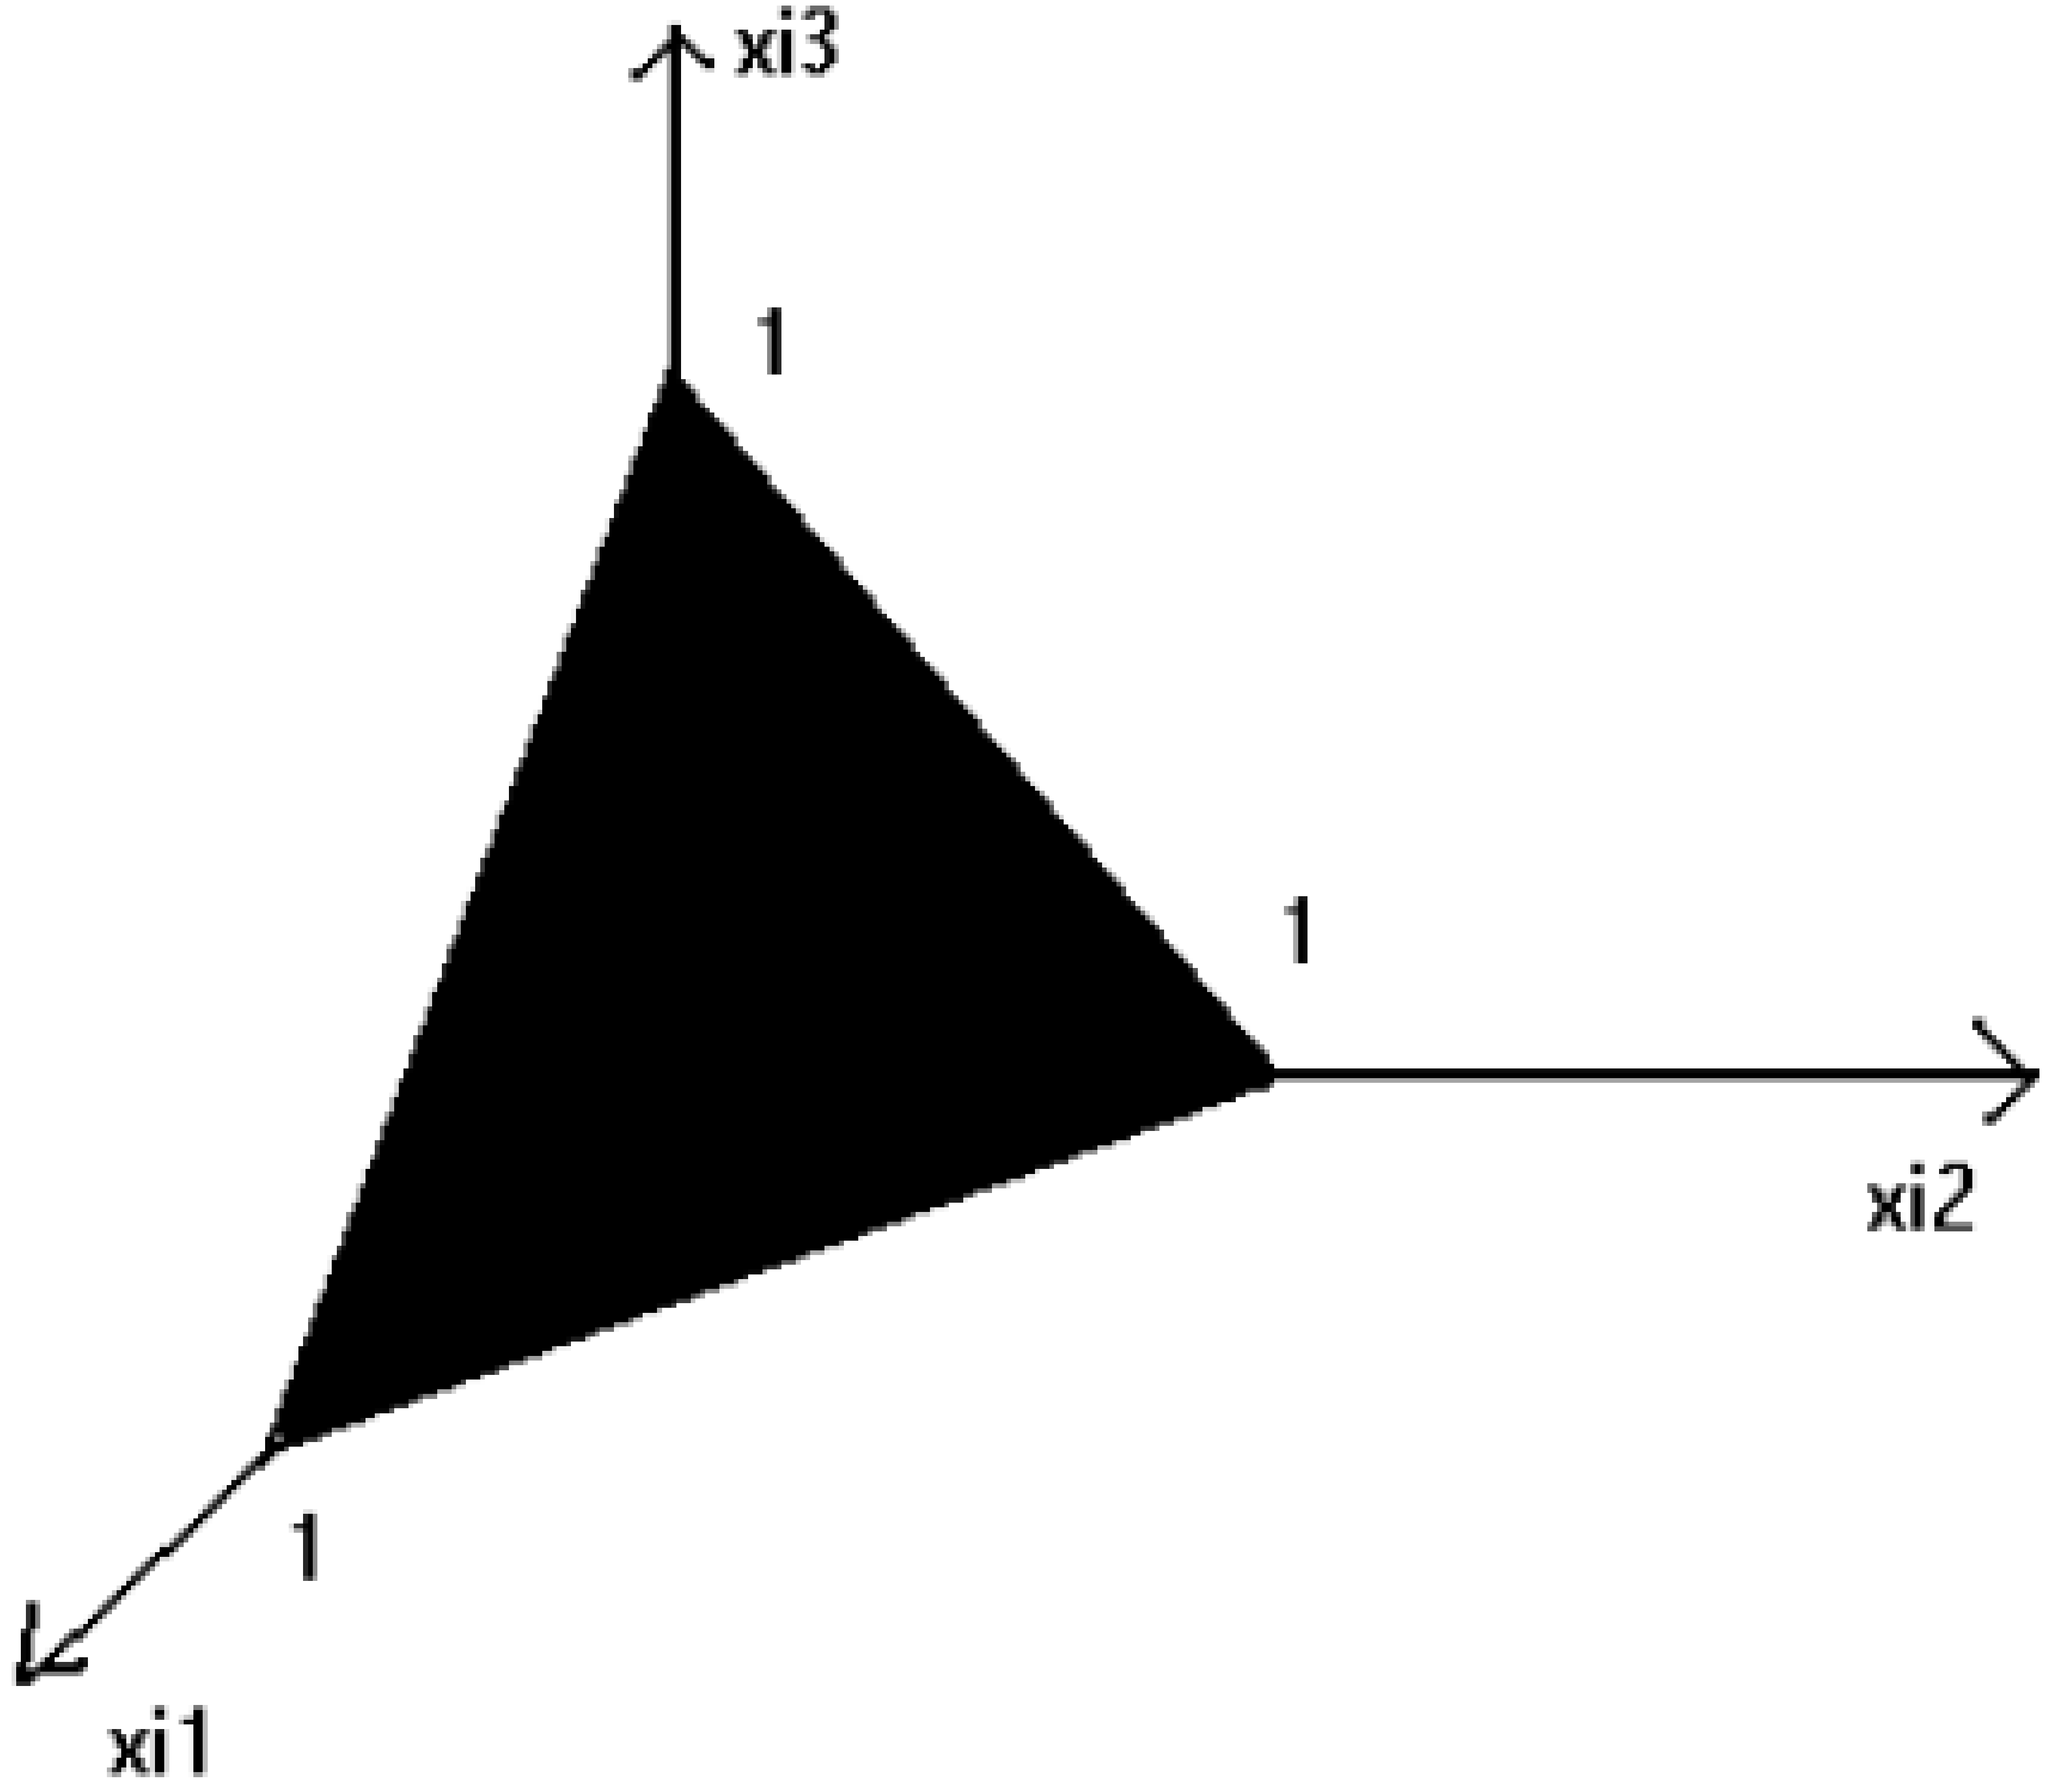

Interacting mixed strategies of all players

Up to now we have considered the strategies chosen by a unique player (i). Henceforth, we will consider the mixed strategies selected by all the participants in the game. The key concepts here are:

- **Profile of mixed strategies:** It is the vector $(x_1, x_2, \dots, x_n)$, where $x_i \in \Delta_i$ is the mixed strategy selected by player i.
- **Space of mixed strategies** Θ is the Cartesian product $\times_i \Delta_i$, whose generic element is $(x_1, x_2, \dots, x_n)$.

Interacting mixed strategies of all players

- Let us assume that there are two players and that each player has available two pure strategies. Then, each player has a strategy set Δ_i, and Θ is four-dimensional.
- But we can take linear projections of each strategy set and reduce Θ to two dimensions.
- Bi-dimensional projection of the space of mixed strategies →

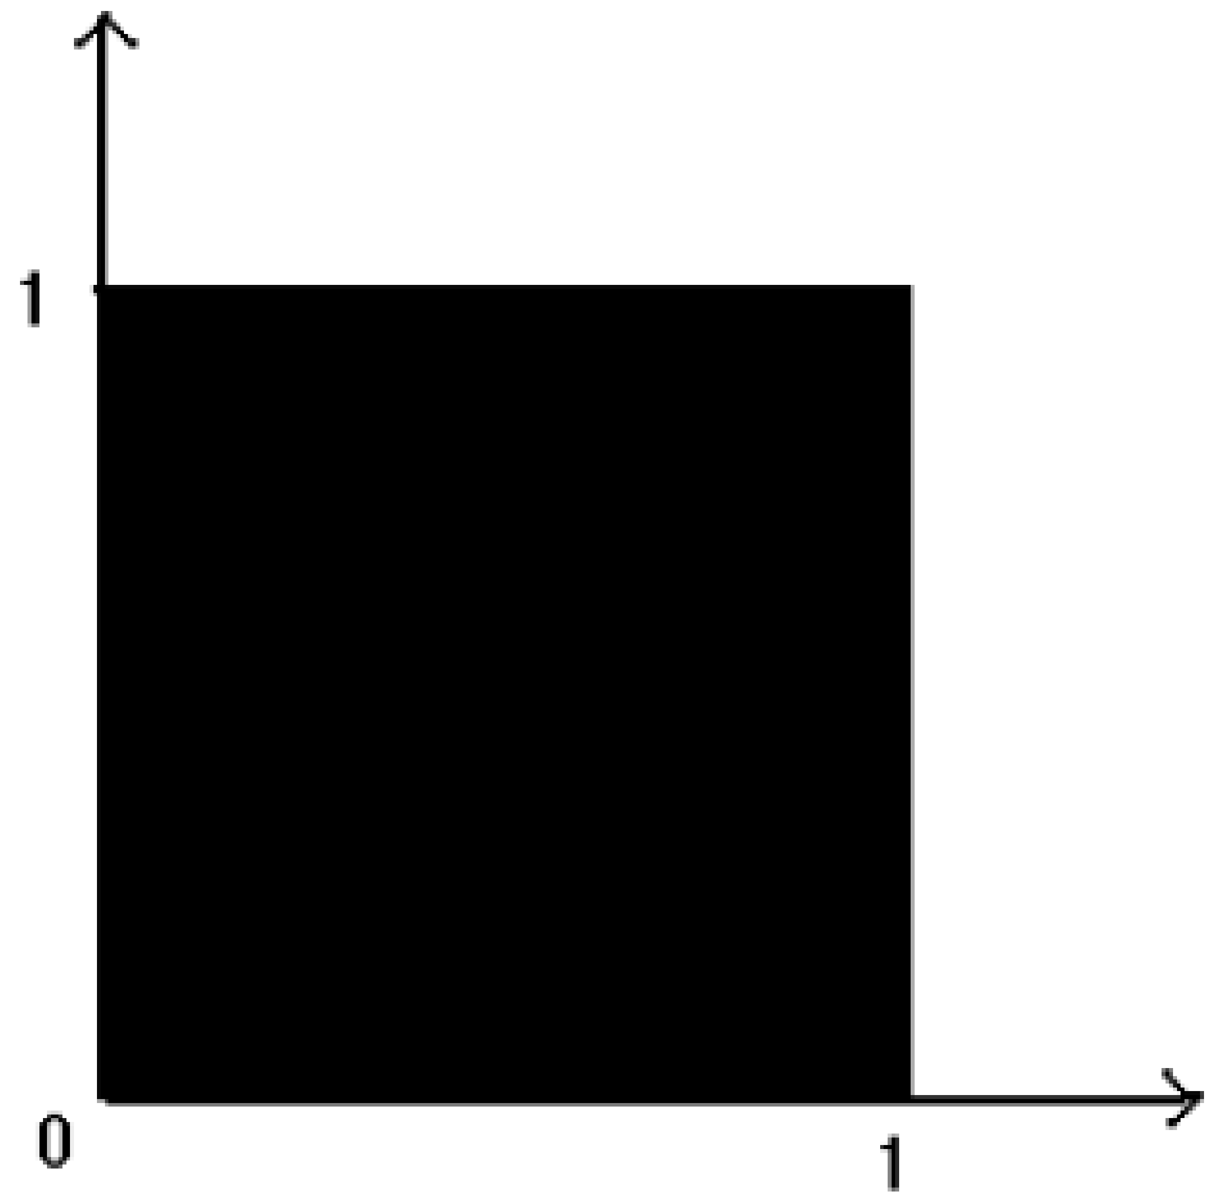

- The dimension of the space of mixed strategies of the game, Θ, is

$$\sum_{i=1}^{n}(m_i - 1) = m - n$$

$$and\ m = m_1 + m_2 + \cdots + m_n$$

$$and\ m_i \equiv number\ of\ pure\ strategies\ available\ to\ player\ i$$

- It is clear that $m > n$, as each player has at least two pure strategies.

Practical remark

Instead of writing a profile of mixed strategies as $x = (x_1, x_2, \dots, x_n)$, $\in \Theta$, we use often the notation (x_i, y_{-i}). If we are concerned mainly with the behavior of player i, we label his strategy as x_i, while we assume that the other players select their strategies according to the profile $y \in \Theta$.

Payoff functions in mixed strategies

- We assume that the random devices of all players are *statistically independent*. Hence, the probability that a given profile of pure strategies $s = (s_1, s_2, \ldots, s_n), \in S$ is used when the profile of mixed strategies $x = (x_1, x_2, \ldots, x_n)$ is adopted by all the players, is simply the associated product defined by,

$$x(s) = \prod_{i=1}^{n} x_{is_i}$$

- where x_{is_i} is the probability that the mixed strategy $x_i \in \Delta_i$, selected by player i, assigns to the pure strategy $s_i \in S_i$.
- The expected payoff v_i of player i, associated with the profile of mixed strategies $x = (x_1, x_2, \ldots, x_n) \in \Theta$, is

$$v_i(x) = \sum_{s \in S} x(s)\, \pi_i(s).$$

- Then, the payoff combined function for all the players is $v(x) = (v_1(x), v_2(x), \ldots, v_n(x))$.

Payoff functions in mixed strategies

- The normal form of a game where only pure strategies are allowed is, as we have seen before, the triple:

$$G = (I, S, \Pi)$$

- whereas the game rules, if mixed strategies are allowed, become

$$G' = (I, \Theta, v)$$

- In the case of finite games with two players, the normal form of the game can be expressed by the payoff matrices (A, B), $A = (a_{hk})$ being the payoff matrix of player 1, whose pure strategies are represented by the rows, and $B = (b_{hk})$, the payoff matrix of player 2, whose pure strategies are expressed by the columns. Then, the expected payoffs of the players related with a profile of strategies $x = (x_1, x_2)$ is,

$$v_1(x) = \sum_{h=1}^{m_1} \sum_{k=1}^{m_2} x_{1h} a_{hk} x_{2k} = x_1 A x_2$$

$$v_2(x) = \sum_{h=1}^{m_1} \sum_{k=1}^{m_2} x_{1h} b_{hk} x_{2k} = x_1 B x_2$$

Example

- Let us consider the game *Prisoner's Dilemma*. The payoff (years in prison) matrices of the game are, as we saw before,

$$A = \begin{pmatrix} 3 & 0 \\ 5 & 1 \end{pmatrix}, B = \begin{pmatrix} 3 & 5 \\ 0 & 1 \end{pmatrix}$$

- We have remarked that, in this game, the first strategy is more advantageous for each player, regardless of the move of the opponent.
- This property is conserved if mixed strategies are allowed. In this case, the expected payoff functions of player 1 is,

$$E\big(\pi_1(x)\big) = x_1 A x_2 = x_{11}(3x_{21} + 0x_{22}) + x_{12}(5x_{21} + 1x_{22})$$

- Moreover, since

$$x_{11} + x_{12} = 1 \leftrightarrow x_{11} = 1 - x_{12}$$

- we obtain, through substitution of x_{11} into $E\big(\pi_1(x)\big)$, that

$$E\big(\pi_1(x)\big) = (1 - x_{12})3x_{21} + x_{12}(5x_{21} + 1x_{22})$$
$$= 3x_{21} + x_{12}(2x_{21} + x_{22})$$

- Inspection of the expected payoff function shows that it is an increasing function of x_{12}. Until player 1 selects its first pure strategy (i.e., $x_{12} = 1$), it pays always to increase the probability that the mixed strategy assigns to the first pure strategy.
- Given the game symmetry, the corresponding reasoning holds for player 2.

References

- GIBBONS, R. (1992), A Primer in Game Theory, New York, Harvester Wheatsheaf, Chapter 1, Sections 1.1.A., 1.3.A, pages 2-4, 29-33.
- MAS-COLLEL, A. et AL. (1995), Microeconomic Theory, chapter 7, sections 7.A, 7.B, 7.D and 7.E, pages: 219-221 and 228-233.
- WEIBULL, Jörgen W. (1997), Evolutionary Game Theory, Cambridge Mass./London, England, chapter 1: 1-9.