
Algorithmic Game Theory and Applications

Lecture 2: Mixed Strategies, Expected Payoffs, and Nash Equilibrium

Kousha Etesami



Finite Strategic Form Games

Recall the “strategic game” definition, now “finite”:

Definition A finite strategic form game Γ , with n -players, consists of:

1. A set $N = \{1, \dots, n\}$ of Players.
2. For each $i \in N$, a finite set

$$S_i = \{1, \dots, m_i\}$$

of (*pure*) strategies. Let $S = S_1 \times S_2 \times \dots \times S_n$ be the set of possible combinations of (pure) strategies.

3. For each $i \in N$, a *payoff (utility) function* $u_i : S \mapsto \mathbb{R}$, describes the payoff $u_i(s_1, \dots, s_n)$ to player i under each combination of strategies.

(Each player prefers to maximize its own payoff.)

We are assuming for convenience that each strategy set is given as $S_i = \{1, \dots, m_i\}$, so we can enumerate each player’s strategies: “strategy 1”, “strategy 2”, ..., “strategy m_i ”.

Mixed (Randomized) Strategies

We define “mixed” strategies for general finite games.

Definition A **mixed** (i.e., **randomized**) **strategy** x_i for Player i , with $S_i = \{1, \dots, m_i\}$, is a probability distribution over S_i . In other words, it is a vector $x_i = (x_i(1), \dots, x_i(m_i))$, such that $x_i(j) \geq 0$ for $1 \leq j \leq m_i$, and

$$x_i(1) + x_i(2) + \dots + x_i(m_i) = 1$$

Intuition: Player i uses randomness to decide which strategy to play, based on the probabilities in x_i .

Let X_i be the set of mixed strategies for Player i .
For an n -player game, let

$$X = X_1 \times \dots \times X_n$$

denote the set of all possible combinations, or **“profiles”**, of mixed strategies.

Expected Payoffs

Let $x = (x_1, \dots, x_n) \in X$ be a profile of mixed strategies.

For $s = (s_1, \dots, s_n) \in S$ a combination of pure strategies, let

$$x(s) := \prod_{j=1}^n x_j(s_j)$$

be the probability of combination s under mixed profile x . (We are assuming the players make their random choices independently.)

Definition: The **expected payoff** of Player i under a mixed strategy profile $x = (x_1, \dots, x_n) \in X$, is:

$$U_i(x) := \sum_{s \in S} x(s) * u_i(s)$$

I.e., it is the “weighted average” of what Player i wins under each pure combination s , weighted by the probability of that combination.

Key Assumption: Every player’s goal is to maximize its own expected payoff.

Discussion: this assumption is sometimes dubious.

some notation

We call a mixed strategy $x_i \in X_i$ pure if $x_i(j) = 1$ for some $j \in S_i$, and $x_i(j') = 0$ for $j' \neq j$. We denote such a pure strategy by $\pi_{i,j}$.

I.e., the “mixed” strategy $\pi_{i,j}$ does not randomize at all: it picks (with probability 1) exactly one strategy, j , from the set of pure strategies for player i .

Given a profile of mixed strategies $x = (x_1, \dots, x_n) \in X$, let

$$x_{-i} = (x_1, x_2, \dots, x_{i-1}, \text{empty}, x_{i+1}, \dots, x_n)$$

I.e., x_{-i} denotes everybody’s strategy except that of player i .

By abuse of notation, for a mixed strategy $y_i \in X_i$, let $(x_{-i}; y_i)$ denote the new profile:

$$(x_1, \dots, x_{i-1}, y_i, x_{i+1}, \dots, x_n)$$

In other words, $(x_{-i}; y_i)$ is the new profile where everybody’s strategy remains the same as in x , except for player i , who switches from mixed strategy x_i , to mixed strategy y_i .

Best Responses

Definition: A (mixed) strategy $z_i \in X_i$ is a **best response** for Player i to x_{-i} if for all $y_i \in X_i$,

$$U_i((x_{-i}; z_i)) \geq U_i((x_{-i}; y_i))$$

Clearly, if any player were given the opportunity to “cheat” and look at what other players have done, it would want to switch its strategy to a best response.

Of course, the rules in the strategic form game don’t allow that: all players pick their strategies simultaneously.

But it still makes sense to consider the following situation:

Suppose, somehow, the players “arrive” at a profile where everybody’s strategy is a best response to everybody else’s. Then no one has any incentive to change the situation.

We will be in a “stable” situation: an “*Equilibrium*”.

That’s what a “Nash Equilibrium” is.

Nash Equilibrium

Definition: For a strategic game Γ , a strategy profile $x = (x_1, \dots, x_n) \in X$ is a mixed Nash Equilibrium if for every player, i , x_i is a best response to x_{-i} .

In other words, for every Player $i = 1, \dots, n$, and for every mixed strategy $y_i \in X_i$,

$$U_i((x_{-i}; x_i)) \geq U_i((x_{-i}; y_i))$$

In other words, *no player can improve its own payoff by unilaterally deviating from the mixed strategy profile $x = (x_1, \dots, x_n)$.*

x is in addition called a pure Nash Equilibrium if every x_i is a pure strategy $\pi_{i,j}$, for some $j \in S_i$.

Brief Discussion: The many “interpretations” of a Nash equilibrium.

Nash's Theorem

This can, with agruable justification, be called

“The Fundamental Theorem of Game Theory”

Theorem(Nash 1950) Every finite n -person strategic game has a mixed Nash Equilibrium.

We will prove this theorem next time.

To prove it, we will “cheat” and use a fundamental result from topology: the Brouwer Fixed Point Theorem.

The crumpled sheet experiment

Let's all please conduct the following experiment:

1. Take two identical rectangular sheets of paper.
2. Make sure neither sheet has any holes in it, and that the sides are straight (not dimpled).
3. "Name" each point on both sheets by its " (x, y) -coordinates".
4. Crumple one of the two sheets any way you like, *but make sure you don't rip it in the process.*
5. Place the crumpled sheet completely on top of the other flat sheet.

Fact! There must be a point named (a, b) on the crumpled sheet that is directly above the same point (a, b) on the flat sheet. (Yes, really!)

As crazy as it sounds, this fact, in its more formal and general form, will be the key to why every game has a mixed Nash Equilibrium.