COSC-6590/GSCS-6390

Games: Theory and Applications

Lecture 02 - Policies

Luis Rodolfo Garcia Carrillo

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

September 6, 2018

Table of contents

Actions vs. Policies

Actions vs. Policies: Advertising Campaign

Action: a possible move, available to a player during a game.

Policy (strategy): a decision rule that a player uses to select actions, based on available information.

The distinction is only important when players acquire **information** during the game that can help them in making better decisions

- e.g., knowing the action of the other player.

Players select a **policy** ahead of time, but they only decide on the specific **actions** to take, as the game evolves.

The concept of Nash or **saddle-point equilibrium refers to policies**, not to specific actions.

A game in marketing: the advertising campaign

Rules: Game played by two companies P_1, P_2 competing for the same set of clients.

Each company has to decide between **two possible actions**:

- **S**: spend $30K in an advertising campaign.
- **N**: keep that money as profit.

Since both companies target the same set of clients the decision of one company will affect not only its own sales, but also the sales of the other company.

Revenue as a function of the actions.

Revenue for P_1

	P_1 chooses N	P_1 chooses S
P_2 chooses N	$169K	$200K
P_2 chooses S	$166K	$185K

Revenue for P_2

	P_1 chooses N	P_1 chooses S
P_2 chooses N	$260K	$261K
P_2 chooses S	$300K	$285K

A game in marketing: the advertising campaign

Revenue for P_1

	P_1 chooses N	P_1 chooses S
P_2 chooses N	\$169K	\$200K
P_2 chooses S	\$166K	\$185K

Revenue for P_2

	P_1 chooses N	P_1 chooses S
P_2 chooses N	\$260K	\$261K
P_2 chooses S	\$300K	\$285K

From **Revenue for P_1**:

- When P_1 advertises, it brings clients to itself.
- When P_2 advertises, it steals clients from P_1.

And a similiar efffect is observed from **Revenue for P_2**.

Objective: Each player wants to **maximize its profit**:

$$\text{profit} = \underbrace{\text{revenue}}_{\text{from Tables}} - \underbrace{\text{adverstising cost}}_{0 \text{ or } \$30\text{K}}$$

The advertising campaign: Simultaneous Play

A version of the game for which both players must decide on their actions without knowing the other player's decision.

The **policies** consist of deciding on a particular action, without extra information collected during the game.

Possible choices for both players, and corresponding profits.

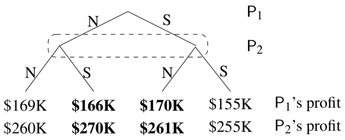

P_2 cannot distinguish between the two nodes inside the box. It must take the same action from all nodes inside the box.

Diagram is called **game representation in extensive form**.

The advertising campaign: Simultaneous Play

Solution (Nash): Simultaneous play has two Nash equilibria:

$$\begin{cases} \pi_1 \equiv P_1 \text{ selects } N \\ \pi_2 \equiv P_2 \text{ selects } S \end{cases} \quad \text{leading to} \quad \begin{cases} \pi_1\text{'s profit} = \$166K \\ \pi_2\text{'s profit} = \$270K \end{cases}$$

or

$$\begin{cases} \pi_1 \equiv P_1 \text{ selects } S \\ \pi_2 \equiv P_2 \text{ selects } N \end{cases} \quad \text{leading to} \quad \begin{cases} \pi_1\text{'s profit} = \$170K \\ \pi_2\text{'s profit} = \$261K \end{cases}$$

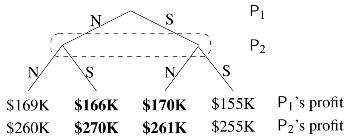

Both policies are Nash equilibria. . .
but they are not equally favorable to both players!

The advertising campaign: Simultaneous Play

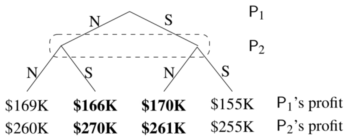

Both pairs of policies (π_1,π_2) are indeed Nash equilibria

- if any one of the players unilaterally **deviates** from their policy, **it will do strictly worse**.

These are **the only two Nash equilibria** for this game.

There is no information available to the players, then, policies consist of **unconditionally** selecting a particular action.

- such policies are often called **open loop**.

The advertising campaign: Simultaneous Play

Note.

For games with a small number of possible actions, one can find all Nash equilibria by **manually checking** which combination of policies satisfies the equilibria conditions.

For larger games, we will need to find **computational solutions**.

The advertising campaign: Simultaneous Play

Note 3 (Repeated games)
Even though both policies are Nash equilibria, they are not equally favorable to both players.

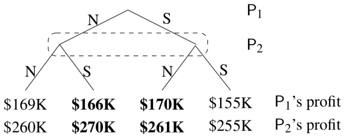

Clearly
- P_1 prefers policy 2.
- P_2 prefers policy 1.

If each player plays on its favorite equilibrium (both select **S**), the resulting policy is not a Nash equilibrium
- actually both players will profit from changing their policies.

The selection of **S** by both players is thus said to be **unstable** and will not remain in effect for long.

The advertising campaign: Simultaneous Play

The selection of **S** by both players is thus said to be **unstable** and will not remain in effect for long.

What will eventually happen is the subject of an area of game theory that investigates **repeated games**.

Repeated Games
The players face each other by playing the same game multiple times and adjust their policies based on the results of previous games.

Key question in this area:
- will the policies converge to a Nash equilibrium? and,
- in case there are multiple equilibria, to which one of them?

The advertising campaign: Alternate Play

Consider a version of the game in which P_1 **plays first** - it selects an **action** before P_2

Then, P_2 plays **knowing** P_1's choice.

Possible choices for this game and the corresponding profits.

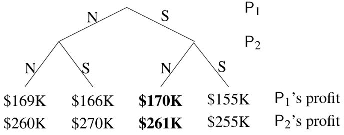

Notation.
This often called the P_1-P_2 version of the game, as opposed to the P_2-P_1 version in which P_2 plays first.

The advertising campaign: Alternate Play

Solution (Nash)

To find a reasonable solution, take the role of the second-playing player P_2 and hypothesize on what to do:

If P_1 chooses N, then selecting S maximizes P_2's profit

$$\begin{cases} P_1\text{'s profit} &= \$166K \\ P_2\text{'s profit} &= \$270K \end{cases}$$

If P_1 chooses S, then selecting N maximizes P_2's profit

$$\begin{cases} P_1\text{'s profit} &= \$170K \\ P_2\text{'s profit} &= \$261K \end{cases}$$

The advertising campaign: Alternate Play

This result motivates the following **rational policy** for P_2

$$\pi_2 \equiv P_2 \text{ selects } \begin{cases} S & \text{if } P_1 \text{ selected N} \\ N & \text{if } P_1 \text{ selected S} \end{cases}$$

If P_1 **assumes** that P_2 will **play rationally**, P_1 can **guess** what profit will result from their actions

- If I choose N, then P_2 will choose S, and my profit will be \$166K.
- If I choose S, then P_2 will choose N, and my profit will be \$170K.

In view of this, the **rational policy** for P_1 is

$$\pi_1 \equiv \text{ selects S}$$

The advertising campaign: Alternate Play

Attention!

The policy

$$\pi_2 \equiv P_2 \text{ selects } \begin{cases} S & \text{if } P_1 \text{ selected N} \\ N & \text{if } P_1 \text{ selected S} \end{cases}$$

is a **nontrivial decision rule** that a player uses to **select actions, based on the available information**.

Such a policy is often called **closed loop**.

The advertising campaign: Alternate Play

One can verify that the pair of policies (π_1,π_2)

$$\pi_1 \equiv \text{ selects S}$$

$$\pi_2 \equiv P_2 \text{ selects } \begin{cases} S & \text{if } P_1 \text{ selected N} \\ N & \text{if } P_1 \text{ selected S} \end{cases}$$

is a **Nash equilibrium**

If any player chooses **a different policy** (i.e., a **different decision rule** for its selection of actions) while the other one keeps its policy, then the first **player will do no better**

- potentially, it will do worse.

Multi-Stage Games

Multi-Stage Games: War of Attrition

The games considered so far were **single-stage**

- each player plays only once per game.

Multi-stage Games

- a **sequence** of rounds (or **stages**) and in each stage the players have the opportunity to take actions.

In **deciding the action** at one stage, the **players** usually **have available information collected from previous stages**.

Example of a multi-stage game

- war of attrition (also known as **the chicken game**)

Multi-Stage Games: War of Attrition

Rules: Two players fight repeatedly for a prize.

At each stage $t \in \{0, 1, 2, \ldots\}$ each player must decide to either

1. **keep fighting** (action "F"), for which the player incurs a **cost of 1**; or
2. **quit** (action "Q"), which means that the **other player gets the prize** with value $v > 1$.

If both players quit at the same time no one gets the prize.

At each stage

- the players **decide on their actions simultaneously without knowing the other's decision at that stage**, but
- **knowing all the actions taken by both players at the previous stages**.

Multi-Stage Games: War of Attrition

Objective: Players want to maximize their discounted profits

$$P_1\text{'s profit} = \begin{cases} \underbrace{-1-\delta-\delta^2-\cdots-\delta^{T_1-1}}_{\text{cost for playing up to time } T_1} & \text{if} \quad \underbrace{T_1 \le T_2}_{P_1 \text{ quits first or tie}} \\ \underbrace{-1-\delta-\delta^2-\cdots-\delta^{T_2-1}}_{\text{cost for playing up to time } T_1} + \underbrace{\delta^{T_2} v}_{\text{prize collected when } P_2 \text{ quits}} & \text{if} \quad \underbrace{T_1 > T_2}_{P_2 \text{ quits first}} \end{cases}$$

$$P_2\text{'s profit} = \begin{cases} \underbrace{-1-\delta-\delta^2-\cdots-\delta^{T_1-1}}_{\text{cost for playing up to time } T_2} & \text{if} \quad \underbrace{T_2 \le T_1}_{P_2 \text{ quits first or tie}} \\ \underbrace{-1-\delta-\delta^2-\cdots-\delta^{T_1-1}}_{\text{cost for playing up to time } T_1} + \underbrace{\delta^{T_1} v}_{\text{prize collected when } P_1 \text{ quits}} & \text{if} \quad \underbrace{T_2 > T_1}_{P_1 \text{ quits first}} \end{cases}$$

T_1 and T_2: **times** at which P_1 and P_2, respectively quit.

$\delta \in (0, 1)$: a per-step **discount factor** that expresses the fact
- that paying in the future is less bad than paying now
- but receiving a prize in the future is less desirable.

Multi-Stage Games: War of Attrition

Note. Discount Factors are well established in games in economics to **take risk or inflation into account**.

They are also useful in engineering applications to **create incentives** for the players to accomplish a goal faster.

Extensive form of the War of Attrition Game

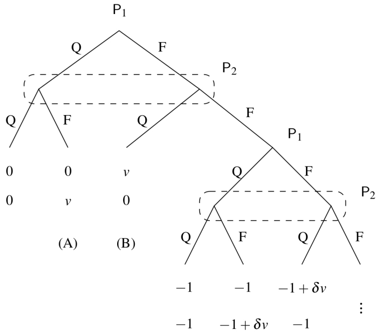

Multi-Stage Games: War of Attrition

Solution (Nash) This game has multiple Nash equilibria, some of them correspond to **open-loop policies**, whereas others correspond to **closed-loop policies**.

It is straightforward to verify that both pairs of (**open-loop**) policies satisfy the **Nash equilibrium** conditions:

$$\begin{cases} \pi_1 \equiv P_1 \text{ quits as soon as possible} \\ \pi_1 \equiv P_2 \text{ never quits} \end{cases} \quad \text{leading to} \quad \begin{cases} P_1\text{'s profit} & = 0 \\ P_2\text{'s profit} & = v \end{cases}$$

or

$$\begin{cases} \pi_1 \equiv P_1 \text{ never quits} \\ \pi_1 \equiv P_2 \text{ quits as soon as possible} \end{cases} \quad \text{leading to} \quad \begin{cases} P_1\text{'s profit} & = v \\ P_2\text{'s profit} & = 0 \end{cases}$$

The outcomes labeled as (A) and (B), respectively.

Multi-Stage Games: War of Attrition

It is less straightforward to verify that the following pairs of policies also satisfy the **Nash equilibrium conditions**:

$$\begin{cases} \pi_1 \equiv \text{ if } P_1 \text{ did not quit before } t \text{, quit at time } t \text{ with probability } p := \frac{1}{1+v} \\ \pi_2 \equiv \text{ if } P_2 \text{ did not quit before } t \text{, quit at time } t \text{ with probability } p := \frac{1}{1+v} \end{cases}$$

or (same conditions, different formulation):

$$\begin{cases} \pi_1 \equiv P_1 \text{ quits at time } t \text{ with probability } (1-p)^t p, \quad p := \frac{1}{1+v} \\ \pi_2 \equiv P_2 \text{ quits at time } t \text{ with probability } (1-p)^t p, \quad p := \frac{1}{1+v} \end{cases}$$

For these pairs of policies, the **profit collected by each player** will be a random variable with expected values

$$E[P_1\text{'s profit}] = E[P_2\text{'s profit}] = 0.$$

Multi-Stage Games: War of Attrition

These solutions raise several questions:

Question 1

Why don't both players quit at time zero, which is much simpler, and gives the same expected profit of 0?

This is **not a Nash equilibrium**: any player that deviates from this will improve their own profit to v.

The second pair of policies yield the same (average) reward as quitting at time zero, but are **safer** to unilateral deviations by one of the players.

Multi-Stage Games: War of Attrition

These solutions raise several questions:

Question 2

Why don't both players never quit?

This would also not be Nash, since they would both **collect** a profit of

$$-\sum_{t=0}^{\infty} \delta^k = -\frac{1}{1-\delta}$$

whereas any one of them could collect a larger profit by quitting at any point.

Multi-Stage Games: War of Attrition

These solutions raise several questions:

Question 3

How to play this game?

This is probably the wrong question to ask!

A more reasonable question may be:

Why should one play a game with such a pathological Nash structure?

The contribution of noncooperative game theory to this problem is precisely to reveal that **the rewards structure of this game can lead to problematic solutions**.

Open vs. Closed-Loop

Open vs. Closed-Loop

Open-loop Game: players select their actions without access to information, other than what is available before game starts.

Policies consist of selecting a particular action, as in the simultaneous play advertising campaign game.

Closed-loop Game: players decide on their actions based on information collected after the game started.

Policies are truly decision laws, as we saw for the alternate plays advertising campaign game.

Note. Some games do not make sense in an open-loop fashion.

Open vs. Closed-Loop: Zebra in the Lake

This game falls under the class of **pursuit-evasion games**.

Rules: This game has two players:

- Player P_1: a zebra that swims in a circular lake with a maximum speed of v_{zebra}.
- Player P_2: a (hydrophobic) lion that runs along the perimeter of the lake with a max speed of $v_{\text{lion}} > v_{\text{zebra}}$.

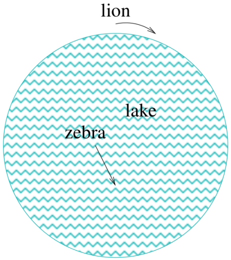

Open vs. Closed-Loop: Zebra in the Lake

Objective: The two players have opposite objectives:

- Zebra: wants to get to the shore of the lake without being caught coming out of the water.
- Lion: wants to be at the precise position where the zebra leaves the lake.

Policies: It is assumed that each player

- constantly sees the other, and
- can react instantaneously to the current position/velocity of the other.

This game only makes sense in **closed loop** if the lion can see the zebra and uses this to **decide where to run**
- otherwise the lion has no chance of ever catching the zebra.

Open vs. Closed-Loop: Zebra in the Lake

Solution: (solved in detail in [1, Chapter 8.5])

When $v_{\mathbf{zebra}} > 0.217 v_{\mathbf{lion}}$, the zebra can always escape.

This is achieved by a **three stage policy**

1.- The zebra starts by swimming to the center of the lake.

2.- Then swims towards the point in the shore opposite (furthest) from the current position of the lion. While she is close to the center of the lake she will be able to keep the (running) lion at the point in the shore furthest from her current position.

3.- At some point the zebra follows a curvy path to the shore that, while not keeping the lion on the opposite side of the shore, suffices to get her to the shore safely.

Open vs. Closed-Loop: Zebra in the Lake

The zebra follows a curvy path to the shore ([1, Equation 8.38])

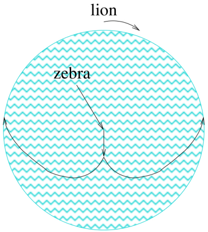

The zebra's actions for the two last stages are a function of the current position of the lion.

When $v_{\textbf{zebra}} < 0.217 v_{\textbf{lion}}$ the zebra may or may not be able to escape, depending on the initial positions of the lion and the zebra

Practice Exercises

Practice Exercises: Multiple Nash equilibrium

Exercise 2.1. Find **all Nash equilibria** for the alternate plays advertising campaign game

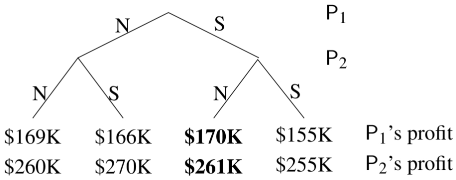

Solution to Exercise 2.1.

There are (only) **three Nash equilibria** for the game.

Practice Exercises: Multiple Nash equilibrium

Solution 1. One option is

$$\pi_1 \equiv P_1 \text{ selects } S, \qquad \pi_2 \equiv P_2 \text{ selects } \begin{cases} S & \text{if } P_1 \text{ selected } N \\ N & \text{if } P_1 \text{ selected } S \end{cases}$$

This is a **Nash equilibrium**

If P_1 selects π_1 (i.e., S), then there is no policy for P_2 that does strictly better than π_2.

Conversely, if P_2 selects π_2, then π_1 is definitely the best choice for P_1.

Practice Exercises: Multiple Nash equilibrium

Solution 2. Another (simpler) option is

$$\pi_1 \equiv P_1 \text{ selects } S, \qquad \pi_2 \equiv P_2 \text{ selects } N$$

This is also a **Nash equilibrium**

If P_1 selects π_1 (i.e., S), then there is no policy for P_2 that does strictly better than π_2, i.e., always chosing N.

Conversely, if P_2 selects π_2 (i.e., N), then π_1 is still the best choice for P_1.

Practice Exercises: Multiple Nash equilibrium

Solution 3. A final option is

$$\pi_1 \equiv P_1 \text{ selects } N, \qquad \pi_2 \equiv P_2 \text{ selects } S$$

This is also a **Nash equilibrium**

If P_1 selects π_1 (i.e., N), then there is no policy for P_2 that does strictly better than π_2, i.e., always choosing S.

Conversely, if P_2 selects π_2 (i.e., S), then π_1 is still the best choice for P_1.

Practice Exercises: Nash equilibrium vs. order of play

Exercise 2.2. Consider the game defined in extensive form in

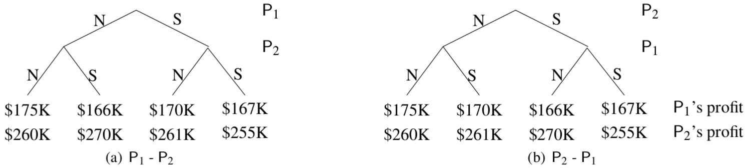

Find Nash equilibria for the following three orders of play

1. P_1 plays first ($P_1 - P_2$)
2. P_2 plays first ($P_2 - P_1$)
3. Simultaneous play.

Practice Exercises: Nash equilibrium vs. order of play

Solution to Exercise 2.2.

1. For P_1 plays first $(P_1 - P_2)$

$$\begin{cases} \pi_1 \equiv P_1 \text{ selects } S \\ \pi_2 \equiv P_2 \text{ selects } \begin{cases} S & \text{if } P_2 \text{ selected } N \\ N & \text{if } P_2 \text{ selected } S \end{cases} \end{cases} \text{leading to } \begin{cases} P_1\text{'s profit} = \$170K \\ P_2\text{'s profit} = \$261K \end{cases}$$

2. For P_2 plays first $(P_2 - P_1)$

$$\begin{cases} \pi_1 \equiv P_2 \text{ selects } N \\ \pi_1 \equiv P_1 \text{ selects } \begin{cases} S & \text{if } P_2 \text{ selected } S \\ N & \text{if } P_2 \text{ selected } N \end{cases} \end{cases} \text{leading to } \begin{cases} P_1\text{'s profit} = \$175K \\ P_2\text{'s profit} = \$260K \end{cases}$$

3. For simultaneous play there is no Nash equilibrium.

Practice Exercises: Chicken game with alternate play

Exercise 2.3. Find a Nash equilibrium policy for the chicken game with alternate play, in which P_1 plays first at each stage.

Solution to Exercise 2.3.

The following policies are a Nash equilibrium

- $\pi_1 \equiv P_1$ quits as soon as possible
- $\pi_2 \equiv P_2$ never quits.

End of Lecture

02 - Policies

Questions?