

GAME THEORY: STATIC GAMES OF COMPLETE INFORMATION

Guillaume Pommey
guillaume.pommey@uniroma2.eu

BAE
University of Rome Tor Vergata

Academic year: 2021/2022

General information

Structure

- Part 1: Standard Approach.
- Part 2: Behavioral Approach (Stefano Papa).

Schedule for Part 1

- Class: Tuesday, Wednesday and Friday (3 weeks).
- Practice session: Tuesday 1, Friday 4 and 11.

General information

Exam

- Joint exam with an exercise for each part of the class.

References

- Osborne, M. (2004). An Introduction to Game Theory, Oxford University Press.
- Gibbons, R. (1992), Game theory for applied economists, Princeton University Press.

Table of Contents

1. **Introduction**
2. Normal-form representation
3. Iterated Elimination of Strictly Dominated Strategies
4. Dominant strategies
5. Nash equilibrium
6. Bertrand duopoly
7. Cournot duopoly
8. Mixed strategies

What is game theory?

The very first question is: **What is a game?**

Let us start with what you **usually** call a game:

- Card games: War, Poker, Bridge, ...
- Board games: Chess, Checkers, Go, ...
- Rock-paper-scissors, dice game, ...
- Sports: Football, handball, ...

What are the **common features** of all those examples?

What is game?

1. **Multiple players**

- At least two players interact.

2. **A set of rules**

- A description of which **actions** players can choose.

3. **Identified outcomes**

- A description of the **consequences** of players' actions.

4. **Conflict of interests**

- Most of the time players prefer different outcomes.
- Not always true (see PS1)!

Strategies

When you play a game against/with other players you **formulate strategies**.

You try to **anticipate** what other players might do:

- According to their **own interests** and as a **response** of what you might do.

And you will try find the **best way to respond** their actions.

Example. In rock-paper-scissors, two players:

- For instance, if you play R the other wants to play P .
- But then you want to play S .
- But now the other wants to play $R \dots$

Strategic interactions

The key idea is that players are aware of **strategic interactions**.

You know that other players **anticipate or react** to what you choose to do and *vice versa*.

Another example. In Chess,

- When you consider moving the bishop to D4 or to E5.
- You try to anticipate how the other will play in each scenario, according to **their own interest**.
- And you choose the scenario that seems **the most favorable to you**.

So what is game theory?

Game theory is a **toolbox to describe and analyze games** in a systematic way.

It relies on a **formal language** to precisely define what is

- A game
- Strategies
- An equilibrium

Mathematical formalism is helpful to

- Describe a wide range of possible games/results
- Rely on **rigorous logical reasoning**

Why is game theory relevant to economics?

What does **economics** have to do with Chess, Poker or Rock-Paper-Scissors?

Most economic problems **involve several players who interact and affect each other**:

- Firms competing by choosing their price.
- States deciding on pollution regulation policies.
- Bargaining over the sale of a car.
- ...

Game theory and economics

Game theory is **pervasive** in modern economic theory.

A vast majority of fields heavily rely on game theory such as

- **Industrial organization**
- Auction theory
- Behavioral economics
- Political Economy
- . . .

A brief look at standard microeconomics

Consider the following topics:

- Consumer theory
- Firm theory
- Market equilibrium

Several players interact with each other and take actions.

- How is that **different from a game?**

A brief look at standard microeconomics

Recall for instance that in standard micro, a consumer solves

$$\begin{aligned} \max_{x_1, \dots, x_n} \quad & u(x_1, \dots, x_n) \\ \text{s.t.} \quad & \sum_{i=1}^n p_i x_i \leq R. \end{aligned}$$

The consumer takes prices p_1, \dots, p_n and the revenue R **as given**.

- It is solely a **decision problem**.
- The agent does not consider their effect **on the rest of the economy**.

A brief look at standard microeconomics

And a firm solves

$$\max_{z_1, \dots, z_p} \quad pf(z_1, \dots, z_n) - \sum_{i=1}^p w_i x_i.$$

The firm takes the market price p and the input prices w_1, \dots, w_p **as given**.

- Also a simple **decision problem**.
- The firm acts as if producing more had no impact on the market price.

A brief look at standard microeconomics

In standard micro, all agents are **atomistic**.

None of them take into account their **impact on the aggregate demand/supply**.

Maybe a reasonable assumption in a **large economy**.

But is it still reasonable when you have **few players**?

- Network providers.
- Supermarkets.
- Airlines.

Goal of the course

Provide you with the **basic toolbox** to analyze games.

- **Static and dynamic games of complete information.**

Some **concrete** applications like Cournot, Bertrand and Stackelberg oligopolies.

But it is **beyond the scope of the course** to provide you with a wide range of economic applications of game theory.

You will apply game theory to more concrete problems in the Industrial Organization course.

Our first game: Informal description

Two friends, 1 and 2, have planned to hang out.

- Our players.

But they **have not decided where to meet** and 1's phone is dead.

They usually meet either at Pigneto or at San Giovanni.

- Their possible actions.

If they meet at the same place, they are happy, otherwise they are sad.

- Their **preferences** associated with each **action profile**.

Our first game: Players

Let us introduce a **more formal language** to describe this game.

1. Players

We call $N := \{1, 2\}$ the **set of players**.

- Describes who is participating in the game.

Our first game: Actions

2. Actions

We call $A_i := \{\text{Pigneto}, \text{San Giovanni}\}$ the **set of actions** of player $i = 1, 2$.

- Describes what each player can do in the game.
- Here, they have the same **action space** (not always the case).

Let $(a_1, a_2) \in A_1 \times A_2$ be an **action profile**.

- For instance (P, P) means both players chose Pigneto.
- And (P, S) means player 1 chose Pigneto and player 2 San Giovanni.

Our first game: Payoffs (1)

3. Preferences/Payoffs

We have to describe players' **preferences over each action profile**.

Let $u_i : A_1 \times A_2 \rightarrow \mathbb{R}$ be player i 's utility.

- For instance, $u_1(P, P)$ tells us how much player 1 values the situation in which they both went to Pigneto.

Our first game: Payoffs (2)

We want to represent the fact that players are happy if they meet at the same place and unhappy otherwise.

- The way to do this is not unique.

For instance, we could choose for $i = 1, 2$:

$$\begin{aligned}u_i(P, P) &= u_i(S, S) = 1 \\u_i(P, S) &= u_i(S, P) = 0.\end{aligned}$$

That is, they get 1 if they go the same place and 0 otherwise.

Our first game: Payoffs (3)

For games with **discrete** action spaces it is useful to introduce the **payoff matrix**:

1 \ 2	Pigneto	San Giovanni
Pigneto	1, 1	0, 0
San Giovanni	0, 0	1, 1

The **first column** (resp. **row**) represents player 1's actions (resp. player 2).

The **first number in each cell** represents player 1's payoff for that action profile. The second number is that of player 2.

Table of Contents

1. Introduction
2. Normal-form representation
3. Iterated Elimination of Strictly Dominated Strategies
4. Dominant strategies
5. Nash equilibrium
6. Bertrand duopoly
7. Cournot duopoly
8. Mixed strategies

Normal-form representation

What we have just done with our first game was to define its **normal-form representation**.

- 1. Players
- 2. Action spaces
- 3. Payoffs

Notice two important things:

- Players know everything about the game: **Complete information**.
- Only one period: **Static game**.

Let us now define more generally what we call a game.

Definition: Normal-form game

Definition. The normal-form representation of a static game of complete information specifies

- **1. A set of players:** $N := \{1, \dots, n\}$.
- **2. An action space for each player:** A_i for each $i \in N$.
 - Let $A := \times_{i \in N} A_i$.
- **3. Preferences of each player over action profiles:**
 $u_i : A \rightarrow \mathbb{R}$ for each $i \in N$.

Let us call this game $G := \langle N, (A_i)_{i \in N}, (u_i)_{i \in N} \rangle$.

Normal-form: Some examples

Rock-paper-scissors. Consider the 2-player game.

- **1. Players:** $N := \{1, 2\}$.
- **2. Action spaces:** $A_i = \{R, P, S\}$ for $i = 1, 2$.
- **3. Preferences:**

$1 \setminus 2$	R	P	S
R	0, 0	-1, 1	1, -1
P	1, -1	0, 0	-1, 1
S	-1, 1	1, -1	0, 0

Normal-form: Some examples

Prisoner's dilemma. Very famous example.

- **1. Players:** $N := \{1, 2\}$.
- **2. Action spaces:** $A_i = \{\text{Confess}, \text{Stay silent}\}$ for $i = 1, 2$.
- **3. Preferences:**

$1 \setminus 2$	C	S
C	$-5, -5$	$0, -10$
S	$-10, 0$	$-1, -1$

Normal-form: Some examples

Cournot duopoly. n firms compete by choosing how much quantity they produce.

- **1. Players:** $N := \{1, \dots, n\}$.
- **2. Action spaces:** $A_i = \mathbb{R}_+$, i.e. each firm chooses $q_i \geq 0$.
- **3. Preferences:** Profits (see later for details).

An outline of the course

We will study two types of games in this course.

Static games: The ones we have just described.

- Players choose an action simultaneously and then the game ends.

Dynamic games: Games in which players interact over more than one period and can observe the past play of their opponents.

Solution concept

Now we know how to represent a **wide range of possible games** with a concise and precise notation.

The next step is to **solve** those games.

- That is, find what players will play and what will be the outcome of the game.

There is **no unique way** to solve a game, it depends on our choice of a **solution concept**.

Here we will explore two of them:

- Iterated Elimination of Strictly Dominated Strategies (IESDS).
- **Nash equilibrium**: The most important one.

Table of Contents

1. Introduction
2. Normal-form representation
3. Iterated Elimination of Strictly Dominated Strategies
4. Dominant strategies
5. Nash equilibrium
6. Bertrand duopoly
7. Cournot duopoly
8. Mixed strategies

Iterated Elimination of Strictly Dominated Strategies

We begin with a very *natural* solution concept: IESDS.

The idea is the following:

- We look at each player's possible actions and their outcome.
- We identify what actions of this player are **strictly dominated**.
- We **eliminate** those actions for this player.
- We repeat the process.

For now, the words **action** and **strategy** are used interchangeably.

Strictly dominated strategies: Example 1

Consider the following game:

1 \ 2	L	R
U	2, 2	3, 1
D	0, 3	1, 4

- When *P2* plays *L*: *P1* is better-off playing *U* than *D* ($2 > 0$).
- When *P2* plays *R*: *P1* is better-off playing *U* than *D* ($3 > 1$).

No matter what *P2* is doing, playing *D* is always **strictly dominated** by *U* for *P1*.

Strictly dominated strategies: Example 2

Consider the following game:

1 \ 2	L	M	R
U	1, 0	1, 2	0, 1
D	0, 3	0, 1	2, 0

- When *P1* plays *U*: *P2* is better-off playing *M* than *R* ($2 > 1$).
- When *P1* plays *D*: *P2* is better-off playing *M* than *R* ($1 > 0$).

No matter what *P1* is doing, playing *R* is always **strictly dominated** by *M* for *P2*.

Definition: Strictly dominated strategies

Let us define formally what is a **strictly dominated strategy**.

Definition. Consider the game $G = \langle N, (A_i)_{i \in N}, (u_i)_{i \in N} \rangle$. For player i , we say that strategy $\underline{s}_i \in A_i$ is **strictly dominated strategy by** strategy $s_i \in A_i$ if

$$u_i(s_i, s_{-i}) > u_i(\underline{s}_i, s_{-i}),$$

for all $s_{-i} \in A_{-i}$.

Note 1: The notation $-i$ means "all players except player i ". Hence if we have three players, $A_{-2} = A_1 \times A_3$.

Note 2: The vector (s_i, s_{-i}) is just a convenient way to write player i 's strategy and all other players' strategies. The order of the arguments is implicitly not changed.

Strictly dominated strategies

Let us apply the definition to the second example.

1 \ 2	L	M	R
U	1, 0	1, 2	0, 1
D	0, 3	0, 1	2, 0

Let $\underline{s}_2 = R$ and $s_2 = M$, we have that

$$u_2(U, s_2) = 2 > u_2(U, \underline{s}_2) = 1,$$

$$u_2(D, s_2) = 1 > u_2(D, \underline{s}_2) = 0,$$

which satisfies our definition of a strictly dominated strategy.

Strictly dominated strategies

Why are we interested in this type of strategies?

It is clear that if some player i has a strictly dominated strategy, then they will surely **never use it**.

Hence, we could **eliminate** this strategy from the game **as if it did not exist** because no rational player would ever use it.

Repeating this process could help us identify what players will do in the end.

IESDS: Example 1 (1)

Let us solve the first example with IESDS.

1 \ 2	L	R
U	2, 2	3, 1
D	0, 3	1, 4

We have seen that strategy D is **strictly dominated** by strategy U for P1.

Let us **eliminate** strategy D .

IESDS: Example 1 (2)

Eliminating D we now have the following game.

1 \ 2	L	R
U	2, 2	3, 1

Only player 2 has to play now.

Clearly, R is **strictly dominated** by L for P2 now (1 ; 2).

- Be careful, this **was not true** before we eliminate D .

Hence, we can **eliminate** R for P2.

IESDS: Example 1 (3)

Eliminating R we finally obtain

$1 \setminus 2$	L
U	$2, 2$

Hence IESDS predicts that the **solution** of the game is the couple of strategies (U, L) .

We will comment more on that later.

IESDS: Example 2 (1)

Let us solve the second example with IESDS.

1 \ 2	L	M	R
U	1, 0	1, 2	0, 1
D	0, 3	0, 1	2, 0

We have seen that strategy R is **strictly dominated** by strategy M for P2.

Let us **eliminate** strategy R .

IESDS: Example 2 (2)

After elimination of R we obtain:

1 \ 2	L	M
U	1, 0	1, 2
D	0, 3	0, 1

Notice that now D is **strictly dominated** by U for P1.

Let us **eliminate** strategy D .

IESDS: Example 2 (3)

After elimination of D we obtain:

$1 \setminus 2$	L	M
U	1, 0	1, 2

Finally, it is clear that P2 will choose M over L .

Hence, we end up with

$1 \setminus 2$	M
U	1, 2

IESDS: Comments

It seems that our first solution concept, IESDS, works quite well.

- We were able to find a **unique solution** in both examples.
- These solutions are **rational** for each player as we have eliminated strategies that they would never want to play.
- Each player's choice is based only on their own payoff, it requires neither knowledge nor understanding of what the other player could do.

Unfortunately, it was a mere artifact of the chosen examples.

IESDS: Not always satisfying (1)

Consider the following example:

1 \ 2	L	M
U	2, 0	2, 3
D	0, 4	4, 2

- When P1 plays U , P2 is better-off playing M than L .
- When P1 plays D , P2 is better-off playing L than M .

Hence, **no strictly dominated strategy** for P2.

IESDS: Not always satisfying (2)

1 \ 2	L	M
U	2, 0	2, 3
D	0, 4	4, 2

- When P2 plays L , P1 is better-off playing U than D .
- When P2 plays M , P1 is better-off playing D than U .

Hence, **no strictly dominated strategy** for P1.

IESDS: Not always satisfying (3)

There is **no strictly dominated strategy to eliminate** in this game.

As a result, **IESDS predicts nothing** more than the initial game itself.

1 \ 2	L	M
U	2, 0	2, 3
D	0, 4	4, 2

IESDS: Conclusion

IESDS offered us a **first look** at player's optimal choices.

But the concept is too strict.

- There exists another version in which you eliminate **weakly dominated strategies** but it suffers from other problems (see chapter 12.3 of Osborne if you are interested).

We can summarize IESDS as a solution concept in which *you eliminate strategies that players will never choose.*

Table of Contents

1. Introduction
2. Normal-form representation
3. Iterated Elimination of Strictly Dominated Strategies
4. Dominant strategies
5. Nash equilibrium
6. Bertrand duopoly
7. Cournot duopoly
8. Mixed strategies

Dominant strategies

A different but conceptually close idea to strictly dominated strategies is the one of **dominant strategy**.

A strategy \bar{s}_i is said to be a dominant strategy for player i if, regardless of what other players do, player i is always better-off playing \bar{s}_i than any other of their strategies.

Dominant strategies: An example

Recall the prisoner's dilemma.

1 \ 2	C	S
C	$-5, -5$	$0, -10$
S	$-10, 0$	$-1, -1$

P1 prefers C **regardless of** what the other player is doing.

- C is a dominant strategy for P1.

The same applies to P2.

Dominant strategies: Definition

Definition. Consider the game $G = \langle N, (A_i)_{i \in N}, (u_i)_{i \in N} \rangle$. We say that strategy $\bar{s}_i \in A_i$ is a **strictly dominant strategy** for player i if

$$u_i(\bar{s}_i, s_{-i}) \geq u_i(s_i, s_{-i}),$$

for all $s_i \in A_i$, all $s_{-i} \in A_{-i}$, and at least one inequality is strict.

Notice that a dominant strategy is not defined with respect to another strategy like it was the case for strictly dominated strategy.

A dominant strategy is **better than all others regardless of what other players do.**

Dominant strategies: Solution

Another possible solution concept can be derived with the notion of **dominant strategies**.

If all players have a dominant strategy, then we can readily **predict the outcome of the game**.

1 \ 2	C	S
C	-5, -5	0, -10
S	-10, 0	-1, -1

(C, C) in the prisoner's dilemma game.

Dominant strategies: Drawbacks

Using dominant strategies as a way to *solve* the game is very strict.

- If even one player has no dominant strategy, we cannot use it.

We are now going to investigate a more subtle solution concept, an **equilibrium** concept.

- Namely, the **Nash equilibrium** solution concept.

Table of Contents

1. Introduction
2. Normal-form representation
3. Iterated Elimination of Strictly Dominated Strategies
4. Dominant strategies
5. Nash equilibrium
6. Bertrand duopoly
7. Cournot duopoly
8. Mixed strategies

Nash equilibrium: Introduction

Nash equilibrium is the most famous **solution concept** for static games of complete information.

"Nash" stands for the name of his creator, the mathematician John Nash.

The important word is **equilibrium**.

- This solution concept will rely on a notion of **stability** of players' choices.

We will start with an informal description and some examples.

NE: Best responses

With IESDS, we were focusing on players' strategies that were **always** better regardless of what the other player was doing.

We will investigate a more subtle notion: **best responses**.

A player i formulates a best response to every possible strategy of another player j .

Informally, player i makes a *plan* of what actions to play in response to the other player's actions.

Best responses: An example

Recall the Rock-paper-scissors game with two players, 1 and 2.

For each player,

- R is a best response to S .
- S is a best response to P .
- P is a best response to R .

Be careful, it **does not mean** that one player is observing what the other one is doing before playing.

Here, the concept of best response only captures the idea that each player can think in advance about how they would act in each possible scenario.

Best responses: Another example

Consider the following game.

1 \ 2	L	M
U	2, 2	0, 1
D	1, 1	3, 0

What are each player's best responses?

Best responses: Another example

1 \ 2	L	M
U	2, 2	0, 1
D	1, 1	3, 0

For player 1,

- U is a best response to L .
- D is a best response to M .

For player 2,

- L is a best response to U .
- L is a best response to D .

Best-responses functions

Let us now informally define what we call **best-response functions** in a two-player game.

For player 1, $BR_1 : A_2 \rightarrow A_1$.

- $BR_1(\cdot)$ is a function that says what player 1 wants to play for each possible action of player 2.
- The function $BR_2 : A_1 \rightarrow A_2$ is defined symmetrically.

$BR_i(a_j) = a_i$ means that player i 's best-response to a_j is a_i .

Best-responses function: Example

1 \ 2	L	M
U	2, 2	0, 1
D	1, 1	3, 0

For player 1,

- U is a best response to L .
 - $BR_1(L) = U$.
- D is a best response to M .
 - $BR_1(M) = D$.

For player 2,

- L is a best response to U .
 - $BR_2(U) = L$.
- L is a best response to D .
 - $BR_2(D) = L$.

Best responses: Interpretations

As said previously, **best responses** help us identify how player i *optimally behaves* for each possible choice of player j .

Once again, **we do not assume** that player i observes player j 's choice to formulate their best response.

Think about you playing rock-paper-scissors or some card game.

- You do not see what the other player is doing before you play but you can still formulate best responses.

Stability and deviations

Together with best responses, let us introduce the concepts of **stability and deviations**.

We are investigating an **equilibrium concept**.

- Think about a physical object being in equilibrium.
- It is **stable**, it doesn't fall.

Similarly, our notion of equilibrium will rely on the idea that we have reached a **stable point** for each player.

- That is, at this point, no player is willing to move to another point.
- We will say that, at this point, **no player is willing to deviate**.

Stability and deviations: Example (1)

Let us come back to our previous example.

1 \ 2	L	M
U	2, 2	0, 1
D	1, 1	3, 0

Consider the point (or action profile) (D, L) .

Is it **stable**?

Stability and deviations: Example (2)

1 \ 2	L	M
U	2, 2	0, 1
D	1, 1	3, 0

For player 2, **it is stable**.

- Indeed, if P2 thinks P1 plays D , it is better for them to play L (1) than M (0).
- And remember we have already said that for player 2, L was a best response to D .

We say that, at (D, L) , player 2 **has no incentive to deviate**.

Stability and deviations: Example (3)

1 \ 2	L	M
U	2, 2	0, 1
D	1, 1	3, 0

For player 1, **it is not stable**.

- Indeed, if P1 thinks P2 plays L , it is better for them to play U (2) than D (1).
- Remember that we had found that U was a best response to L .

We say that, at (D, L) , player 1 **wants to deviate**.

Stability and deviations: Example (4)

1 \ 2	L	M
U	2, 2	0, 1
D	1, 1	3, 0

Hence, the idea is that (D, L) is **not stable**.

- That is, if we were to start at this point it will "fall" because one player (P1) is not willing to remain at this point.

This point **cannot be an equilibrium**.

Stability and deviations: Example (5)

What about (D, M) ? Is it stable?

1 \ 2	L	M
U	2, 2	0, 1
D	1, 1	3, 0

It is for P1.

But not for P2.

- P2 would like to **deviate** to L .

Stability and deviations: Example (5)

What about (U, M) ? Is it stable?

1 \ 2	L	M
U	2, 2	0, 1
D	1, 1	3, 0

It is not for P1.

And neither it is for P2.

Stability and deviations: Example (6)

What about (U, L) ? Is it stable?

1 \ 2	L	M
U	2, 2	0, 1
D	1, 1	3, 0

It is for P1 and it is for P2.

At (U, L) , no player is willing to deviate.

- It is a stable point.
- And we will see it is what we call a Nash equilibrium.

What have we learned?

The example illustrated some important aspects that we will see over and over.

First, we have investigated each "point" and look whether one of the players was better-off **deviating** than "sticking to the point".

- You can think about finding the physical equilibrium of an object once again.
- You would place it in a given position (a point) and see whether it moves or stay put.
- Then you repeat the operation for other positions and see which ones are stables.

Best responses and stability

Second, notice that we can establish a link between **stability** and **best responses**.

1 \ 2	L	M
U	2, 2	0, 1
D	1, 1	3, 0

Best responses:

- $BR_1(L) = U$.
- $BR_1(M) = D$.
- $BR_2(U) = L$.
- $BR_2(D) = L$.

What is the common features among all non stable points (D, L) , (D, M) and (U, M) ?

Best responses and stability

Second, notice that we can establish a link between **stability** and **best responses**.

1 \ 2	L	M
U	2, 2	0, 1
D	1, 1	3, 0

Best responses:

- $BR_1(L) = U$.
- $BR_1(M) = D$.
- $BR_2(U) = L$.
- $BR_2(D) = L$.

What is the common features among all non stable points (D, L) , (D, M) and (U, M) ?

- There is always at least one player whose best response is different from the chosen point.

Best responses and stability

What about our stable point (U, L) ?

1 \ 2	L	M
U	2, 2	0, 1
D	1, 1	3, 0

Best responses:

- $BR_1(L) = U$.
 - $BR_1(M) = D$.
 - $BR_2(U) = L$.
 - $BR_2(D) = L$.
- For P1, U is a best response to L .
 - For P2, L is a best response to U .

The peculiar feature here is that each player's strategy is a best response to the other player's strategy.

Best responses and stability

Another way to think about this notion of "stability" is the following.

Recall the Rock-Paper-Scissors example and let R_i , P_i and S_i denote players' strategies with indexes to keep track of the identity of the player.

If we *conceptually* think about the *cycle* of players' best responses to each other, we have the following pattern:

$$R_1 \rightarrow P_2 \rightarrow S_1 \rightarrow R_2 \rightarrow P_1 \rightarrow S_2 \rightarrow R_1 \rightarrow P_2 \rightarrow \dots$$

You see that it is a *never ending* pattern in which players always want to change their strategy.

Best responses and stability

Now think about the same concept but in our last example. We have

$$U \rightarrow L \rightarrow U \rightarrow L \rightarrow U \rightarrow \dots$$

When P1 plays U , P2 wants to play L

- then P1 still wants to play U
- and then P2 still wants to play L
- ...

Here it is **stable** as one player action *reinforces* the action of the other player.

Another example

Consider another example.

1 \ 2	C	D
A	1, 1	0, 2
B	0, -1	2, 0

Identify best-response functions for each player and equilibria.

Another example

Solution to the example.

1 \ 2	C	D
A	1, 1	0, 2
B	0, -1	2, 0

We have $BR_1(C) = A$, $BR_1(D) = B$, $BR_2(A) = D$ and $BR_2(B) = D$.

Only equilibrium is (B, D) .

Nash equilibrium

What we have presented informally as *stable points* are in fact what we will call **Nash equilibria** from now on.

We will say that an action profile $a^* = (a_1^*, \dots, a_n^*) \in A$ is a Nash equilibrium **if no player i has an incentive to deviate and play some $a_i \neq a_i^*$ given that all other players play a_{-i}^* .**

The central idea is that of **unilateral deviation**.

To identify an equilibrium point we fix the strategy of all $j \neq i$ and we investigate only the incentive to deviate of i .

- And we repeat for all $i \in N$.

Nash equilibrium: Formal definition

Definition. Consider the game $G = \langle N, (A_i)_{i \in N}, (u_i)_{i \in N} \rangle$. We say that the action profile $a^* = (a_1^*, \dots, a_n^*) \in A$ is a **pure-strategy Nash equilibrium** of the game G if

$$u_i(a_i^*, a_{-i}^*) \geq u_i(a_i, a_{-i}^*),$$

for all $a_i \in A_i$ and for all $i \in N$.

Note 1: For each player i , we *fix* all other players' actions to a_{-i}^* and we only investigate whether a_i^* does better than any other a_i .

Note 2: Forget about the term "pure-strategy" for now, we will see later that there exists other types of Nash equilibria, namely "mixed-strategy" NE.

NE: Example in a two-player game

Consider a two-player game and say that (a_1^*, a_2^*) is a Nash equilibrium of this game.

Using the definition of a Nash it means that the following inequalities hold:

$$u_1(a_1^*, a_2^*) \geq u_1(a_1, a_2^*) \text{ for all } a_1 \in A_1$$

$$u_2(a_1^*, a_2^*) \geq u_2(a_1^*, a_2) \text{ for all } a_2 \in A_2.$$

They say that

- P1 is better-off playing a_1^* than any other a_1 when P2 plays a_2^* .
- P2 is better-off playing a_2^* than any other a_2 when P1 plays a_1^* .

NE: Example in a two-player game

You can also think about the definition of a Nash equilibrium in terms of best responses.

The inequality for player 1

$$u_1(a_1^*, a_2^*) \geq u_1(a_1, a_2^*) \text{ for all } a_1 \in A_1,$$

says that a_1^* is a **best response** to a_2^* .

NE: Example in a two-player game

The inequality for player 2

$$u_2(a_1^*, a_2^*) \geq u_2(a_1^*, a_2) \text{ for all } a_2 \in A_2,$$

says that a_2^* is a **best response** to a_1^* .

Hence we come back to our previous informal discussion:

$$a_1^* \rightarrow a_2^* \rightarrow a_1^* \rightarrow a_2^* \rightarrow \dots$$

NE: Another interpretation

It is useful to interpret the Nash equilibrium condition in another way.

For player i , the **equilibrium condition** is that

$$u_i(a_i^*, a_{-i}^*) \geq u_i(a_i, a_{-i}^*),$$

for all $a_i \in A_i$ and given a_{-i}^* .

In more *mathematical terms* what does $u_i(a_i^*, a_{-i}^*)$ represents for player i ?

NE: Another interpretation

Given actions a_{-i}^* , the action a_i^* is the action that **maximizes** player i 's utility.

In other words, a_i^* solves,

$$\max_{a_i \in A_i} u_i(a_i, a_{-i}^*),$$

or, we can also write

$$a_i^* = \arg \max u_i(a_i, a_{-i}^*).$$

- This last notation is not entirely correct (the "=" sign in particular), but we will see later why.

NE: Example in a two-player game

Let us apply our definition to our previous example:

1 \ 2	L	M
U	2, 2	0, 1
D	1, 1	3, 0

We have that

$$\begin{aligned}u_1(U, L) = 2 &\geq u_1(D, L) = 1 \\u_2(U, L) = 2 &\geq u_2(U, M) = 1.\end{aligned}$$

NE: Example in a two-player game

But for other action profiles it is not the case.

1 \ 2	L	M
U	2, 2	0, 1
D	1, 1	3, 0

For instance, for (D, L)

$$u_1(D, L) = 1 < u_1(U, L) = 2$$

$$u_2(D, L) = 1 \geq u_2(D, M) = 0.$$

Player 1's inequality is not satisfied.

NE: Example in a two-player game

But for other action profiles it is not the case.

1 \ 2	L	M
U	2, 2	0, 1
D	1, 1	3, 0

For instance, for (D, M)

$$u_1(D, M) = 3 \geq u_1(U, M) = 0$$

$$u_2(D, M) = 0 < u_2(D, L) = 1.$$

Player 2's inequality is not satisfied.

NE: A trick to solve simple games

Now that we know how NE works, here is a useful trick to solve games presented in the payoff matrix form.

Consider the following game.

1 \ 2	L	M
U	2, 2	0, 1
D	1, 1	3, 0

For each player, we are going to spot their best responses and underline the corresponding payoff.

NE: A trick to solve simple games

For instance, we know that for player 1, D is a best response to M .

Hence we will **underline the corresponding payoff** for this player, 3 in that case.

We obtain:

1 \ 2	L	M
U	2, 2	0, 1
D	1, 1	<u>3</u> , 0

NE: A trick to solve simple games

If we repeat the process to find all best responses, this is what we obtain:

1 \ 2	L	M
U	<u>2</u> , <u>2</u>	0, 1
D	1, <u>1</u>	<u>3</u> , 0

Notice that the cell in which both players' payoffs are underlined is our Nash equilibrium (U, L).

- Consistent with the fact that a NE corresponds to the situation in which each player's best response "reinforce" themselves.

NE: Example 1

Find all pure-strategy Nash equilibria of this game.

1 \ 2	Pigneto	San Giovanni
Pigneto	1, 1	0, 0
San Giovanni	0, 0	1, 1

NE: Example 1

Find all pure-strategy Nash equilibria of this game.

1 \ 2	Pigneto	San Giovanni
Pigneto	<u>1</u> , <u>1</u>	0, 0
San Giovanni	0, 0	<u>1</u> , <u>1</u>

There are two (pure-strategy) Nash equilibria, (P, P) and (S, S) .

NE: Example 2

Find all pure-strategy Nash equilibria of this game.

1 \ 2	L	M	R
U	1, 0	1, 2	0, 1
D	0, 3	0, 1	2, 0

There is one (pure-strategy) Nash equilibrium, (U, M) .

NE: Example 2

Find all pure-strategy Nash equilibria of this game.

1 \ 2	L	M	R
U	$\underline{1}, 0$	$\underline{1}, \underline{2}$	$0, 1$
D	$0, \underline{3}$	$0, 1$	$\underline{2}, 0$

There is one (pure-strategy) Nash equilibrium, (U, M) .

NE: Example 3

Find all pure-strategy Nash equilibria of this game.

1 \ 2	R	P	S
R	0, 0	-1, 1	1, -1
P	1, -1	0, 0	-1, 1
S	-1, 1	1, -1	0, 0

NE: Example 3

Find all pure-strategy Nash equilibria of this game.

1 \ 2	R	P	S
R	0, 0	-1, <u>1</u>	<u>1</u> , -1
P	<u>1</u> , -1	0, 0	-1, <u>1</u>
S	-1, <u>1</u>	<u>1</u> , -1	0, 0

There is no (pure-strategy) Nash equilibrium in this game.

- We will see later than there is one in mixed strategy.

NE: Remarks

Multiplicity: Sometimes, there is more than one Nash equilibrium.

- In that case, our solution concept leaves us with multiple solutions.

No pure-strategy NE: Sometimes, there is no NE in pure-strategy (rock-paper-scissors).

- Hopefully, there are theorems that ensure that there always exists a NE (possibly in mixed strategies) in finite games (finite action spaces).

More on best responses

Earlier, we have defined **best responses as functions**.

For two players we have defined:

$$BR_1 : A_2 \rightarrow A_1$$

$$BR_2 : A_1 \rightarrow A_2.$$

By definition, a function $f : X \rightarrow Y$ is a mapping that assigns to every element of X , **one element** in Y .

Our best response function BR_i assigns to every action of player j , **one action** of player i .

More on best responses

Is it reasonable to do so? Consider the following game.

1 \ 2	L	R
U	1, 0	0, 1
D	1, 1	0, 0

- P1's best response to L is U .
- But D is also P1's best response to L .

So, what is $BR_1(L)$?

- $BR_1(L) = U$?
- $BR_1(L) = D$?

More on best responses

The problem in defining best responses as functions is that we do not know what to do in the previous case.

This is not just a technical point.

- Sometimes players are indifferent between two or more actions.
- We must take this into account otherwise we may miss some equilibria.

If we consider that $BR_1(L) = U$ in the previous example, then we would miss the pure-strategy NE: (D, L) .

Best-response correspondences

The solution is to think best responses as **correspondences** and not functions.

In a two-player game, we define

$$BR_1 : A_2 \rightrightarrows A_1$$

$$BR_2 : A_1 \rightrightarrows A_2.$$

The symbol \rightrightarrows says that BR_i is a mapping that assigns every element in A_{-i} to a subset of A_i .

- That is, a best response to a_j may not be unique.

In the example, we would have $BR_1(L) = \{U, D\}$.

Best responses as maximizers

Now that we understood that BR could be nonunique, it is useful to come back to our other interpretation of the NE conditions.

Recall that

$$\max_{a_i \in A_i} u_i(a_i, a_{-i}^*).$$

Hence

$$a_i^* \in \arg \max u_i(a_i, a_{-i}^*),$$

where the " \in " sign indicates that a_i^* is one among possibly other best responses against a_{-i}^* .

Indeed $\arg \max u_i(a_i, a_{-i}^*)$ is **a set**.

- The set of all actions that are BR to a_{-i}^* .

Best responses as maximizers

We can go even a bit further because it will be helpful later on.

In general, for any action profile a_{-i} , we can set the best response correspondence of player i as being

$$a_i(a_{-i}) = \arg \max u_i(a_i, a_{-i}).$$

That is, $a_i(a_{-i})$ gives us the set of all actions of player i that are best responses to a_{-i} .

Table of Contents

1. Introduction
2. Normal-form representation
3. Iterated Elimination of Strictly Dominated Strategies
4. Dominant strategies
5. Nash equilibrium
6. Bertrand duopoly
7. Cournot duopoly
8. Mixed strategies

Duopoly and competition

Now that we are equipped with the very first tool of game theory, Nash equilibrium, we are going to investigate more concrete applications.

We want to study **competition in a duopoly setting**.

Two firms produce an homogeneous good and compete for consumers.

It is reasonable to assume that firms have **market power**.

- They are **not price takers**.
- They are aware that their choices affect demand/price and the other competitor.

Bertrand and Cournot competition

There are two standard settings for oligopolistic competition in static games.

- **Bertrand:** Price competition.
- **Cournot:** Quantity competition.

We will see a third setting after introducing dynamic games:
Stackelberg oligopoly.

Bertrand duopoly

We start with the **Bertrand duopoly** setting.

Consider two firms, 1 and 2, that produce an **homogeneous good**.

Firms choose the price at which they sell the good.

- Let p_1 and p_2 denote those prices.

Both firms have a **marginal cost of production** equal to c .

Bertrand duopoly: Demand (1)

Consumers always buy to the **firm that sets the lowest price**.

When p_i is the **lowest price**, firm i faces demand $Q(p_i)$.

- With $Q' < 0$.

If $p_1 = p_2$ then consumers **split equally** between the two firms.

If p_i is the **highest price**, firm i faces zero demand.

Bertrand duopoly: Demand (2)

Hence we can write the demand faced by firm i as follows:

$$Q_i(p_i, p_j) = \begin{cases} Q(p_i) & \text{if } p_i < p_j, \\ \frac{1}{2}Q(p_i) & \text{if } p_i = p_j, \\ 0 & \text{if } p_i > p_j. \end{cases}$$

The profit (payoff) of firm i is given by

$$\pi_i(p_i, p_j) = (p_i - c)Q_i(p_i, p_j).$$

Notice that the profit of firm i **depends on both prices**.

Bertrand: Normal form

Let us write this game using the normal-form representation.

- **1. Players:** Firms, $N = \{1, 2\}$.
- **2. Actions:** Each firm chooses its price $p_i \in \mathbb{R}_+$.
 - Hence the action space of firm i is $A_i = \mathbb{R}_+$.
- **3. Payoffs: Profits** $\pi_i(p_i, p_j) = (p_i - c)Q_i(p_i, p_j)$.

It is a strategic game because each firm is aware that its choice of price affects the other firm's profits and thus its choice of price as well.

Bertrand: Nash equilibrium

How do we find the **Nash equilibrium** of this game?

We cannot use a payoff matrix to find best responses as we did for finite games.

Hence, we will have to think otherwise.

- This is where the best response reasoning will be helpful

Bertrand: Nash equilibrium

First, how do we **formally define a Nash equilibrium** of this game?

A Nash equilibrium is a couple of prices (p_1^*, p_2^*) such that **no firm is willing to change its price given the price of the other firm.**

That is, prices (p_1^*, p_2^*) must be such that

$$\begin{aligned}\pi_1(p_1^*, p_2^*) &\geq \pi_1(p_1, p_2^*) \text{ for all } p_1 \in \mathbb{R}_+, \\ \pi_2(p_1^*, p_2^*) &\geq \pi_2(p_1^*, p_2) \text{ for all } p_2 \in \mathbb{R}_+.\end{aligned}$$

Once again, this means that p_1^* must be a best response to p_2^* and *vice versa*.

Bertrand: Best response functions

We are going to construct **best response functions** of each firm.

- Sometimes called *reaction functions*.

Let $BR_i(p_j)$ denote firm i 's best response to firm j 's price.

Let us first try to determine a **lower and an upper bound** on the price set by firm i .

- That is, prices that firm i will never find it optimal to set.

Bertrand: Lower bound on prices

A clear **lower bound** for firm i 's profit can be easily found by inspecting its profit function:

$$\pi_i(p_i, p_j) = (p_i - c)Q_i(p_i, p_j).$$

It is clear that setting $p_i < c$ would entail **negative profits** if this price were to attract consumers to firm i .

Bertrand: Upper bound on prices

Finding an **upper bound** is a bit more subtle.

For that matter, we have to think about what firm i would do, if it were a **monopolist**?

If firm i were the only firm, it would choose its price, p^m , so as to solve

$$\max_{p \geq 0} (p - c)Q(p).$$

And then, it is clear that firm i would never set $p_i > p^m$ even if it were to attract all consumers.

Bertrand: Best-response function (1)

Now let us think about firm i 's **best response** to p_j .

First assume that $p_j \in [0, c]$.

In that case, if firm i wants to attract any consumer, it has to set price below p_j and thus below c .

But we have seen that this would entail **negative profits**.

Hence, the best firm i could do is to set $p_i = c$ and makes zero profit.

- Actually, firm i could set any price $p_i > p_j$, but we assume it chooses $p_i = c$ for simplicity.

Bertrand: Best-response function (2)

Now assume that $p_j \in (p^m, +\infty)$.

Then firm i can choose $p_i = p^m$, attract all consumers and achieve the highest possible profits, the monopoly profit.

It is important to understand that setting $p^m < p_i < p_j$ is not optimal for firm i .

- Because setting $p_i > p^m$ is suboptimal for firm i , by definition of the monopoly price.

Bertrand: Best-response function (3)

Finally assume that $p_j \in (c, p^m]$.

In that region, firm i can make positive profits by setting $c < p_i < p_j$ and thus attract all consumers.

Also, as $p_j < p^m$, firm i would like to increase p_i as close as possible to p_j but always arbitrarily lower so as to attract all consumers.

Informally, we could say that firm i is willing to set $p_i = p_j - \epsilon$ with $\epsilon > 0$ and as small as possible.

- *Note:* Formally, it is not a very well-defined problem but we will ignore this for simplicity.

Bertrand: Best-response function (4)

To summarize, the best response function of firm i could be written as follows:

$$BR_i(p_j) = \begin{cases} c & \text{if } p_j \leq c, \\ p_j - \epsilon & \text{if } p_j \in (c, p^m], \\ p^m & \text{if } p_j > p^m. \end{cases}$$

The two firms' best-response functions are the same as firms are symmetric (same marginal cost).

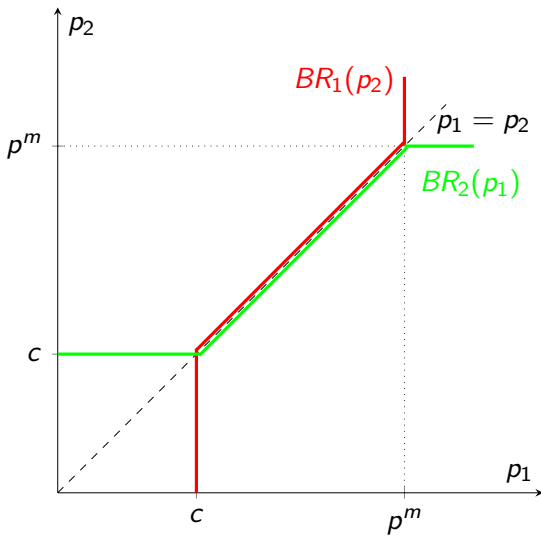
Bertrand: Nash equilibrium

We know that, at a Nash equilibrium, each firm's BR must be a BR to the other firm's BR.

That is $BR_i(BR_j(p_j)) = p_i$ for $i = 1, 2$ and $i \neq j$.

In our case, the best is to draw the best-response functions in the space (p_1, p_2) to see things more clearly.

Bertrand: Plotting BR



Bertrand: Nash equilibrium

To find the Nash equilibrium, we have to find **where the best responses intersect**.

Be careful, it may seem that they intersect all along the interval $(c, p^m]$.

But remember that each firm's best response in this interval is to **undercut** its competitor by setting a price that is slightly lower to get the whole demand.

Bertrand: Nash equilibrium

There is only one intersection point, when $p_1 = p_2 = c$.

We can check that no firm wants to deviate:

- Let $p_j = c$.
- Then if firm i chooses $p_i < p_j = c$, it makes negative profits.
- If otherwise firm i chooses $p_i > p_j = c$, it makes zero profits.
- While if firm i chooses $p_i = p_j = c$, it makes positive profits (gets half the demand).

Hence, we can conclude that $p_1 = p_2 = c$ is the **Nash equilibrium of this Bertrand duopoly game**.

Bertrand: Remarks

This result is sometimes called the **Bertrand paradox**. Why?

It suggests that moving from one firm (monopoly) to only two firms is enough to **fully restore the competitive equilibrium** (price = marginal cost).

- However, this result relies on some extreme assumptions, notably that firms have the same marginal costs.

Another important remark: In Bertrand, we say that prices are **strategic complements**.

- Best-response functions are increasing.
- If firm j sets a higher price, firm i best responds by increasing its own price as well.

Table of Contents

1. Introduction
2. Normal-form representation
3. Iterated Elimination of Strictly Dominated Strategies
4. Dominant strategies
5. Nash equilibrium
6. Bertrand duopoly
7. Cournot duopoly
8. Mixed strategies

Cournot duopoly

We now investigate the **Cournot duopoly** setting.

Consider two firms, 1 and 2, that produce an **homogeneous good**.

Firms choose how much quantity of the good they produce.

- Let q_1 and q_2 denote those quantities.
- The market price is defined by $P(q_1 + q_2)$ with $P' < 0$ and P continuously differentiable.

Each firm has a **marginal cost of production** equal to c_i .

Cournot: Normal form

Let us write this game using the normal-form representation.

- **1. Players:** Firms, $N = \{1, 2\}$.
- **2. Actions:** Each firm chooses its quantity $q_i \in \mathbb{R}_+$.
 - Hence the action space of firm i is $A_i = \mathbb{R}_+$.
- **3. Payoffs:** Profits $\pi_i(q_i, q_j) = (P(q_i + q_j) - c_i)q_i$.

Each firm's choice of quantity affects the market price $P(\cdot)$ and thus affects both firms' profits.

- Realistic? Models should never be taken *literally*.
- See Kreps and Scheinkman (1983).

Cournot: Nash equilibrium

How do we **formally define a Nash equilibrium** of this game?

A Nash equilibrium is a couple of quantities (q_1^*, q_2^*) such that **no firm has an incentive to deviate**.

That is, quantities (q_1^*, q_2^*) must be such that

$$\pi_1(q_1^*, q_2^*) \geq \pi_1(q_1, q_2^*) \text{ for all } q_1 \in \mathbb{R}_+,$$

$$\pi_2(q_1^*, q_2^*) \geq \pi_2(q_1^*, q_2) \text{ for all } q_2 \in \mathbb{R}_+.$$

Cournot: Best-response functions

As in the Bertrand case, we are going to work with **best-response functions** to identify our Nash equilibrium.

Our profit functions $\pi_i(q_i, q_j) = (P(q_i + q_j) - c_i)q_i$ are differentiable in q_i .

- We can use calculus to determine best responses.

Firm i 's best response to a given q_j is the q_i that solves

$$\max_{q_i \geq 0} (P(q_i + q_j) - c_i)q_i.$$

Cournot: Best-response functions

Notice that the problem $\max_{q_i \geq 0} (P(q_i + q_j) - c_i)q_i$ will give a different q_i for each possible q_j .

Alternatively, we can write that

$$q_i(q_j) = \arg \max_{q_i \geq 0} (P(q_i + q_j) - c_i)q_i.$$

That is, $q_i(q_j)$ will tell us what is the best q_i firm i can set against each q_j .

- In other words, it's firm i 's best response function.

Cournot: Linear demand

We are not going to provide a general solution to the Cournot problem.

We will simply focus on the **linear demand case**.

- $P(q_i + q_j) = a - b(q_1 + q_2)$, where $(a, b) \in \mathbb{R}_+^2$.

Once again, notice that the market price is determined by the choice of quantity of both firms. Each of them is aware of the impact of their choice.

Cournot: Best responses

Firm i 's best response to a given q_j is the q_i that solves

$$\max_{q_i \geq 0} \pi_i(q_i, q_j) = (a - b(q_i + q_j) - c_i)q_i.$$

The first-order condition of this problem writes

$$a - b(q_i + q_j) - c_i - bq_i = 0.$$

Solving for q_i we get

$$q_i = \frac{a - c_i - bq_j}{2b}.$$

Cournot: Best responses

We could also write, for clarity,

$$q_i(q_j) = \frac{a - c_i - bq_j}{2b}.$$

This is the best-response function of firm i to firm j 's quantity.

- Notice that $q_i(q_j)$ is a decreasing function of q_j .
- We say that quantities are **strategic substitutes**.

Also notice that we could write that

$$q_i(q_j) = \arg \max_{q_i \geq 0} (a - b(q_i + q_j) - c_i)q_i.$$

Cournot: Nash Equilibrium

We have our best-response functions for each firm,

$$q_i(q_j) = \frac{a - c_i - bq_j}{2b}.$$

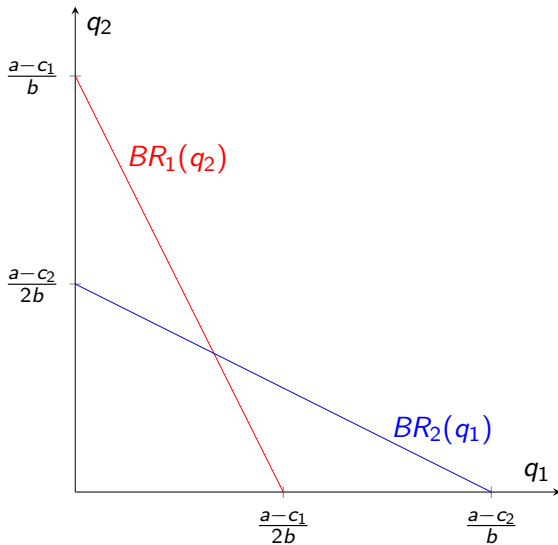
A Nash equilibrium (q_1^*, q_2^*) will be such that each firm's best response corresponds to the other firm's best response.

We could write it as follows:

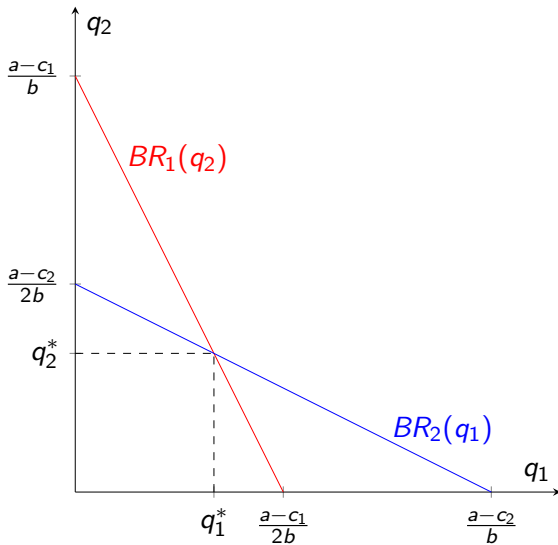
$$\begin{aligned}q_1(q_2(q_1^*)) &= q_1^*, \\q_2(q_1(q_2^*)) &= q_2^*.\end{aligned}$$

The first equality reads: The best response of firm 1 to the best response of firm 2 against q_1^* is q_1^* itself.

Cournot: Plotting BR



Cournot: Plotting BR



Cournot: Nash Equilibrium

Recall that our equilibrium quantities must solve

$$\begin{aligned}q_1(q_2(q_1^*)) &= q_1^*, \\q_2(q_1(q_2^*)) &= q_2^*.\end{aligned}$$

Using the expression for BR, it can be rewritten as:

$$\begin{aligned}\frac{a - c_1 - bq_2^*}{2b} &= q_1^*, \\ \frac{a - c_2 - bq_1^*}{2b} &= q_2^*.\end{aligned}$$

This is simply a **linear system, with two equations and two unknowns**, q_1^* and q_2^* .

- Solving it will give us the intersection point of the plot.

Cournot: Nash Equilibrium

Start with

$$\frac{a - c_1 - bq_2^*}{2b} = q_1^*,$$
$$\frac{a - c_2 - bq_1^*}{2b} = q_2^*.$$

We can for instance, **plug the first equation into the second one**:

$$\frac{a - c_2 - b \frac{a - c_1 - bq_2^*}{2b}}{2b} = q_2^*.$$

And now that this equation only depends on q_2^* , solving it for q_2^* will give us the equilibrium quantity of firm 2.

Cournot: Equilibrium quantities

We have

$$\frac{a - c_2 - b \frac{a - c_1 - bq_2^*}{2b}}{2b} = q_2^*,$$

that simplifies to

$$\frac{a - c_2}{2b} - \frac{a - c_1 - bq_2^*}{4b} = q_2^*,$$

and finally, solving for q_2^* :

$$q_2^* = \frac{a + c_1 - 2c_2}{3b}.$$

Cournot: Equilibrium quantities

Now plugging $q_2^* = \frac{a+c_1-2c_2}{3b}$ into the first equation we get

$$q_1^* = \frac{a + c_2 - 2c_1}{3b}.$$

Hence, we have found the **two equilibrium quantities**.

We could also write that

$$q_i^* = \frac{a + c_j - 2c_i}{3b}.$$

for $i = 1, 2$ and $i \neq j$.

Cournot: Equilibrium profits

We can also compute the equilibrium profits.

Recall that $\pi_i(q_i, q_j) = (a - b(q_i + q_j) - c_i)q_i$.

The equilibrium profit of firm i is given by

$$\begin{aligned}\pi_i(q_i^*, q_j^*) &= (a - b(q_i^* + q_j^*) - c_i)q_i^* \\ &= \left(a - b\frac{2a - c_i - c_j}{3b} - c_i\right)\frac{a + c_j - 2c_i}{3b} \\ &= \frac{(a + c_j - 2c_i)^2}{9b}.\end{aligned}$$

Cournot: Equilibrium price

We can also compute the equilibrium price $P(q_1^*, q_2^*)$. We have

$$\begin{aligned} P(q_1^*, q_2^*) &= a - b(q_1^* + q_2^*) \\ &= a - b \frac{2a - c_i - c_j}{3b} \\ &= \frac{a + c_i + c_j}{3}. \end{aligned}$$

Cournot: Comments

Under quantity competition, **firms make positive profits.**

We do not reach a *competitive outcome* like it is the case in Bertrand.

Assume $c_1 = c_2$ to allow for comparison. It is possible to show that the joint production of Cournot firms is higher than the quantity a monopolist would choose.

- The presence of an additional firm **creates some competition.**
- But there is still **market power.**

Table of Contents

1. Introduction
2. Normal-form representation
3. Iterated Elimination of Strictly Dominated Strategies
4. Dominant strategies
5. Nash equilibrium
6. Bertrand duopoly
7. Cournot duopoly
8. Mixed strategies

Pure and mixed strategies

Before introducing the notion of mixed strategy, let us define more precisely the one we have used implicitly: Pure strategy.

So far, we said that players were choosing an action $a_i \in A_i$ to play in the game.

The implicit assumption here is that we assumed that players had to choose **one and only one action to play in the game**.

What if player i could **randomize over their possible actions**?

Mixed strategies

Now it becomes important to make the **distinction between actions and strategies**.

Because now, a **mixed strategy** will be a different object from an action.

Assume that $A_i = \{L, R\}$.

- Then L and R are actions.
- But “playing L half of the time and R the other half” is a (mixed) strategy.

Mixed strategies: Definition

Definition. Consider the game $G = \langle N, (A_i)_{i \in N}, (u_i)_{i \in N} \rangle$. A mixed strategy for player i is a probability distribution $(p_1, p_2, \dots, p_{|A_i|})$ over the possible actions of player i . By definition, each $p_k \in [0, 1]$ and $\sum_{k=1}^{|A_i|} p_k = 1$

Example. Assume $A_i = \{R, P, S\}$. Then we write that $(p_1, p_2, p_3) = (0.5, 0.5, 0)$ is a mixed strategy such that player i plays R and P with probability 0.5 and never plays S .

We could also say that the mixed-strategy $(p_1, p_2, p_3) = (0, 1, 0)$ corresponds to the pure-strategy “playing P ”.

Mixed strategies: An example

Let us try to find an equilibrium such that players choose mixed strategies.

Consider the following game.

1 \ 2	H	T
H	-1, 1	1, -1
T	1, -1	-1, 1

Notice that there is no NE in pure strategy.

Mixed strategies: Randomizing

First, notice that it is clear that if P1 were sure that P2 would play H (resp. T) then they would play T (resp. H) for sure, i.e., a pure strategy.

Hence, if we want P1 to randomize over their actions, it must be the case that P2 also randomizes over their actions.

Let us call $q \in [0, 1]$ and $1 - q$, the probabilities that P2 plays H and L respectively.

Mixed strategies: Payoffs

Given q , we can compute P1's **expected payoff** when playing H and T :

$$\nu_1(H, q) = -1q + 1(1 - q) = 1 - 2q,$$

$$\nu_1(T, q) = 1q + (-1)(1 - q) = 2q - 1.$$

What can we say about P1's strategy when $\nu_1(H, q) > \nu_1(T, q)$?

- That against q , P1 strictly prefers H to L .
- Symmetric reasoning for $\nu_1(H, q) < \nu_1(T, q)$.

But we know that if P1 plays a pure strategy then P2 will want to play a pure strategy as well.

Mixed strategies: Payoffs

Hence, the only possibility is that P1 is indifferent between H and T .

That is, we must have

$$\begin{aligned}\nu_1(H, q) &= \nu_1(T, q) \\ \Leftrightarrow 1 - 2q &= 2q - 1 \\ \Leftrightarrow q &= \frac{1}{2}.\end{aligned}$$

By symmetry, the same applies to P2.

- P2 is indifferent between H and T if and only if P1 plays each of their action with probability $\frac{1}{2}$.

Mixed strategies: An example

Hence, there exists one mixed-strategy Nash equilibrium of this game:

1 \ 2	H	T
H	-1, 1	1, -1
T	1, -1	-1, 1

And it is when both players **randomize equally over their actions**.

Mixed strategies: Another example

Consider now the following game.

1 \ 2	L	R
U	2, 1	0, 0
D	0, 0	1, 2

Let p and q denote the probabilities that P1 plays U and P2 plays L , respectively.

Mixed strategies: Another example

When is player 1 **indifferent** between playing U and D ?

Let us compute

$$\nu_1(U, q) = 2q + 0(1 - q) = 2q,$$

$$\nu_1(D, q) = 0q + 1(1 - q) = 1 - q.$$

Hence P1 is indifferent whenever $2q = 1 - q$ which corresponds to

$$q = \frac{1}{3}.$$

Mixed strategies: Another example

When is player 2 **indifferent** between playing L and R ?

Let us compute

$$\nu_2(p, L) = 1p + 0(1 - p) = p,$$

$$\nu_2(p, R) = 0p + 2(1 - p) = 2 - 2p.$$

Hence P2 is indifferent whenever $p = 2 - 2p$ which corresponds to

$$p = \frac{2}{3}.$$

Mixed strategies: Another example

Hence for the game:

1 \ 2	L	R
U	2, 1	0, 0
D	0, 0	1, 2

There exists a **mixed-strategy Nash equilibrium** such that

- P1 plays U with probability $\frac{2}{3}$ and D with probability $\frac{1}{3}$.
- P2 plays L with probability $\frac{1}{3}$ and R with probability $\frac{2}{3}$.

Mixed strategies: Remarks

Notice that it is not always the case that both players randomize.

Consider the following example.

1 \ 2	L	R
U	2, 1	0, 0
D	2, 3	1, 0

Mixed strategies: Remarks

1 \ 2	L	R
U	2, 1	0, 0
D	2, 3	1, 0

First, notice that L is a **dominant strategy** for P2.

Second, notice that P1 is **indifferent** between U and D when P2 plays L .

Hence, what could be a strategy for P1?

Mixed strategies: Remarks

1 \ 2	L	R
U	2, 1	0, 0
D	2, 3	1, 0

P1 could play U for sure, or D for sure and get a payoff of 2 in both cases.

But P1 could choose to **randomize** and play U with any probability $p \in [0, 1]$, they would also get a payoff of 2.

Hence, here P2 is playing a pure-strategy meanwhile P1 could play **any mixed or pure strategy**.

Mixed strategies: Remarks

Consider the slightly modified game:

1 \ 2	L	R
U	2, 0	0, 1
D	2, 3	1, 0

The difference is that now L is not anymore a dominant strategy for P2.

Can we still have an equilibrium in which P2 plays L with probability one while P1 randomizes between U and L ?

Mixed strategies: Remarks

1 \ 2	L	R
U	2, 0	0, 1
D	2, 3	1, 0

Let us denote by p the probability that P1 plays U .

The expected payoffs of P2 when playing L and R are, respectively:

$$\nu_2(p, L) = 0p + 3(1 - p) = 3 - 3p,$$

$$\nu_2(p, R) = 1p + 0(1 - p) = p.$$

Mixed strategies: Remarks

Hence from,

$$\nu_2(p, L) = 0p + 3(1 - p) = 3 - 3p,$$

$$\nu_2(p, R) = 1p + 0(1 - p) = p,$$

we have that P1 strictly prefers L to R if $\nu_2(p, L) > \nu_2(p, R)$ or equivalently

$$3 - 3p > p \Leftrightarrow p < \frac{3}{4}.$$

It means that P2 will play L for sure against $(p, 1 - p)$ only if P2 does not play U too often.

Because for P2, playing L when P1 plays U is worse than playing R now but playing L against D is better than R .