Nash Equilibrium: Strategy Across Games and Markets

Ernest Tang

February 2025

Contents

1 Introduction

Consider a market with two companies, Coca-Cola and Pepsi, competing by adjusting their selling prices to **optimise** production to **maximise** profit. How do these firms settle on a price that balances profit and competitiveness? With John F. Nash's discovery of the **Nash Equilibrium**, businesses can achieve a state where neither firm can increase its profit by unilaterally changing its strategy.

1.1 History of Nash Equilibrium

In the early development of game theory, the focus was predominantly on **two-player zero-sum games**, where one player's gain is another's loss, (such as in chess). This foundational perspective **limited** the study of **more complex interactions** where players have non-zero-sum interests, allowing for mutual benefit. In our earlier example, Coca-Cola and Pepsi do not face a simple **win-lose** scenario; instead, they seek an **optimal balance** that benefits both firms.

To address games involving **multiple players** and **non-zero-sum payoffs**, Nash introduced the concept of Nash Equilibrium in 1950, broadening game theory's application in **economics, politics, evolutionary biology and other sectors**. His theory offers a systematic way of studying and predicting strategic behaviour, ultimately leading to his well-deserved Nobel Prize in Economic Sciences in 1994, which recognised his equilibrium concept as a vital tool for understanding **real-world decision-making**.

Understanding Nash Equilibrium first requires a basic grasp of **game theory** and **strategic interactions**. [1]

1.2 What is Game Theory?

Game theory studies **strategic interactions** among rational decision-makers. It analyses how individuals behave in strategic games, where each player's outcome depends not only on their own decisions but also on the actions of others.

By creating **mathematical models**, we can **formulate, structure and analyse** the interactions between players and their decisions, allowing us to predict and influence outcomes to enhance our welfare.

Game theory is essential in **competitive situations**, as it helps explain existing strategies and optimise decision-making. Its impact extends across multiple sectors, including **economics, biology and sociology**. For example, it can help justify strategies of political candidates competing for votes or explain the behaviour of animals fighting over prey.

We now delve into game theory's core principle: the theory of rational choice.

1.3 Theory of Rational Choice

Game theory relies on the foundational **theory of rational choice**, which assumes that players act rationally to maximise their outcomes. To understand a **normal-form game** (a finite, n-person game), we define the following key variables:

- N : A finite set of n players, indexed by i.
- A_i : A finite set of actions available to Player i.
- $a = \{a_1, \cdots, a_n\}$: An **action profile**, representing all strategies chosen by the players.
- $u = \{u_1, \cdots, u_n\}$: A **payoff function**, where u_i assigns a real-valued payoff to player i.

The **payoff function** assigns a numerical value to each outcome based on specific actions. For example:

$$u(a_1) = 1, \quad u(a_2) = 0.$$

It means that the **first action** is preferred over the second, as it yields a **higher payoff**.

However, it is important to note that the payoff function only ranks preferences **rather than measuring the actual level of desirability**. For instance:

$$u(a_1) = 10, \quad u(a_2) = 3.$$

convey the same relative preference as:

$$u(a_1) = 1, \quad u(a_2) = 0.$$

Example: Rock-Paper-Scissors
Consider a two-player game of rock-paper-scissors:

- Players: $N = \{1, 2\}$.
- Action Sets: $A_1 = A_2 = \{rock, paper, scissors\}$.
- Action Profile: $a = \{paper, rock\}$, meaning Player 1 chooses *paper*, and Player 2 chooses *rock*.

We might assign the following payoffs:

$$u_1(paper) = 1, \quad u_2(rock) = 0.$$

Since paper beats rock, Player 1 receives a higher payoff. However, these payoffs only indicate a ranking of strategies and are not absolute numerical values - **different versions of the game may assign different payoffs**.

Rational Choice Assumption

The rational choice theory states that a decision-maker will always choose **the action that is at least as good as every other available action**. In other words, players select the action (a) that yields the highest payoff value (u).

How is this related to game theory?
In game theory, we assume that **all players are rational** and will always choose the best option to achieve the highest possible payoff. This assumption **leads to the concept of Nash Equilibrium**.

However, in real-world decision making:

1. Players may **lack complete information**, leading to suboptimal choices (especially for inexperienced players).
2. Some games are **too complex to compute** the best strategy (e.g., in chess, where skill and experience play a role).

Thus, rational choice theory explains how players select their best responses based on **known facts**, but it does not guarantee that players always find the best solution. [2]

1.4 Pure Strategy

In the **rock-paper-scissors** example, each player selects a single action and plays it without considering alternative choices. The outcome remains consistent - for instance, if Player 1 chooses *paper*, they will always win against *rock*. This type of strategy, **where a player consistently chooses the same action**, is known as a **pure strategy**.

Now that we have established the fundamentals of game theory, we begin defining and understanding **the Nash Equilibrium**, where decision-makers find optimal strategies in their interactions.

2 Nash Equilibrium

2.1 Definition

In a game, a player's best choice of action often depends on the actions of others. A Nash Equilibrium occurs **when no player can unilaterally change their strategy to improve their outcome**, assuming all other players keep their strategies unchanged. In this state, each player's strategy is optimal given the choices of others, meaning no player has an intention to deviate. Here is how players arrive at a Nash Equilibrium in a game:

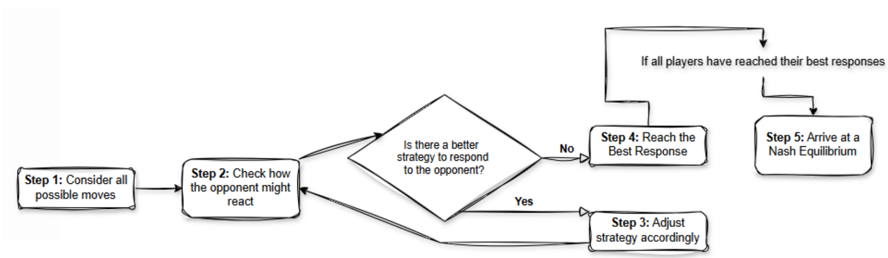

Figure 1: Flowchart Outlining the Process of Reaching a Nash Equilibrium

To express this mathematically, recall that an action profile a represents a specific combination of actions of all players. Let:

- a_i be **the action of player** i.
- a_{-i} be the actions of all players (excluding i).

For example, in the action profile (a_i, a'_{-i}), the notation a'_{-i} represents the actions chosen by all players except i, while Player i chooses a_i. We denote the **best response** of player i as a_i^*.

To compare different strategies, we use a **payoff function** $u_i(a)$ to evaluate the numerical outcomes of distinct actions. For example, if:

$$u(paper) > u(rock),$$

it indicates that the action *paper* is preferred over *rock* (ignoring *scissors* for now). More generally, the **best response condition** states that:

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

where a_i' represents **any alternative action** for Player i apart from a_i^*.

Returning to the definition of Nash Equilibrium, we can now state the formal condition:

A strategy profile a^* is a Nash Equilibrium **if no players can achieve a higher payoff by changing their strategy**, given that all other players maintain their best strategies:

$$u_i(a^*) \geq u_i(a_i, a_{-i}^*) \quad \forall a_i, \forall_i.$$

This equation means that for every player i ($\forall i$), the payoff from their equilibrium strategy a_i^* is at least as good as the payoff from any other possible strategy a_i ($\forall a_i$), given that the other players are using their equilibrium strategies a_{-i}^*.

2.2 Prisoner's Dilemma (Example)

In Section 2.1, we explored Nash Equilibrium through best response strategies and payoff. To deepen our understanding, we now turn to the classic example of the **Prisoner's Dilemma**, a strategic game that illustrates the Nash Equilibrium.

Scenario

Two suspects are arrested and imprisoned separately. The police lack sufficient evidence to convict them of the major crime unless one of them confesses (defects). The possible outcomes are:

- **If both remain silent** ($Quiet$, $Quiet$) → each receives 1 year in prison for a minor offence
- **If one confesses while the other stays silent** ($Defect$, $Quiet$) → the confessor is freed, while the other serves 4 years.
- **If both confess** ($Defect$, $Defect$) → both receive 3 years in prison.

The scenario can be represented as **a strategic game**, where both suspects try to minimise their time in prison. The ranked action profiles for each suspect, from best to worst, are:

$$(Defect, Quiet) > (Quiet, Quiet) > (Defect, Defect) > (Quiet, Defect)$$

To simplify calculations, we assign numerical payoffs for both players (Player 1, Player 2):

$$u(Defect, Quiet) = (3,\ 0), \quad u(Quiet, Quiet) = (2,\ 2),$$
$$u(Defect, Defect) = (1,\ 1), \quad u(Quiet, Defect) = (0,\ 3)$$

Representing this visually:

		Player 2	
		Quiet	**Defect**
Player 1	**Defect**	3, 0	1, 1
	Quiet	2, 2	0, 3

Figure 2: Prisoner's Dilemma Payoff Matrix

Finding the Nash Equilibrium

To determine each player's best response, we analyse the payoff matrix:

1. To start, we focus on how Player 1 would respond rationally. If Player 2 chooses $Quiet$, we consider the vertical column $Quiet$ and can note that Player 1's best option is '$Defect$' (since $u(Defect)$ $> u(Quiet)$).

		Player 2	
		Quiet	**Defect**
Player 1	**Defect**	3, 0	1, 1
	Quiet	2, 2	0, 3

Figure 3: Player 1 Chooses $Defect$ Given That Player 2 Chooses $Quiet$

2. Similarly, if Player 2 chooses $Defect$, we read off the vertical $Defect$ column and notice that Player 1's best option remains $Defect$.

		Player 2	
		Quiet	**Defect**
Player 1	**Defect**	3, 0	1, 1
	Quiet	2, 2	0, 3

Figure 4: Player 1 Chooses $Defect$ Given That Player 2 Chooses $Defect$

3. Repeating this for Player 2, we find that $Defect$ is also their best choice, regardless of Player 1's decision.

		Player 2	
		Quiet	**Defect**
Player 1	**Defect**	3, 0	(1, 1)
	Quiet	2, 2	(0, 3)

Figure 5: Player 2 Chooses Defect Regardless of Player 1's Decision

Since **both player's best responses** (circled) **lead to** $(Defect, Defect)$, this must be the Nash Equilibrium.

Why $(Defect, Defect)$ is the Nash Equilibrium

Although $(Quiet, Quiet)$ gives both players **a better outcome**, it is not a Nash Equilibrium. If one of the players **deviates to** $Defect$, they receive a higher payoff. Since both players have an incentive to defect, $(Defect, Defect)$ remains the only stable outcome. Thus, the Prisoner's Dilemma demonstrates how individual rationality can lead to a **suboptimal** joint result.

2.3 Other Examples

The Stag Hunt

The Stag Hunt models cooperation, where players must trust each other to achieve the best outcome.

Scenario

A group of hunters must decide whether to:

- **Cooperate to hunt a stag** (which provides a large reward but requires full cooperation).
- **Hunt a hare alone** (a smaller but guaranteed reward)

If all the hunters choose $Stag$, they successfully hunt it together, yielding the highest payoff. However, **if one hunter chooses** $Hare$, the stag escapes, leaving the defector with a hare while the others get nothing.

Finding the Nash Equilibriums

Considering the scenario with only two hunters, we can summarise the preferences in a matrix table:

		Player 2	
		Stag	**Hare**
Player 1	**Stag**	2, 2	0, 1
	Hare	1, 0	1, 1

Figure 6: Stag Hunt Payoff Matrix

Applying Nash Equilibrium analysis, we find **two stable outcomes**:

		Player 2	
		Stag	**Hare**
Player 1	**Stag**	(2, 2)	0, 1
	Hare	1, 0	(1, 1)

Figure 7: Two Stable Outcomes in Stag Hunt

- $(Stag, Stag)$ → Both players choose *Stag*, leading to maximum payoff.
- $(Hare, Hare)$ → Both players choose *Hare*, ensuring a consistent but lower payoff.

Although (*Stag*, *Stag*) provides the highest payoff, **both equilibriums are stable** because no player is incentivised to switch strategies unilaterally.

Multiplayer Stag Hunt In a game with more than two hunters, the Nash Equilibriums remain:

$$(Stag, \cdots, Stag) \textbf{ and } (Hare, \cdots, Hare).$$

The intuition is that if one player defects, it breaks cooperation, forcing the others to choose $Hare$; to avoid getting nothing.

Rock-Paper-Scissors Returning to rock-paper-scissors, we analyse its Nash Equilibrium.

Rock beats *Scissors*, *Scissors* beats *Paper*, *Paper* beats *Rock*.

Using payoff values (1 for a win, 0 for a loss), we construct the payoff matrix and identify the best responses:

		Player 2		
		Rock	**Paper**	**Scissors**
Player 1	**Rock**	0, 0	-1, 1	1, -1
	Paper	1, -1	0, 0	-1, 1
	Scissors	-1, 1	1, -1	0, 0

Figure 8: Rock-Paper-Scissors Payoff Matrix

The following table includes circled best responses to the opponent's actions.

		Player 2		
		Rock	**Paper**	**Scissors**
Player 1	**Rock**	0, 0	-1, 1	1, -1
	Paper	1, -1	0, 0	-1, 1
	Scissors	-1, 1	1, -1	0, 0

Figure 9: Best Responses in Rock-Paper-Scissors

In this example, no pure strategy Nash Equilibrium exists.

- **No action pair** satisfies both simultaneously, leading to continuous deviations.
- Since a pure strategy always has a counter-strategy, the game never settles for a fixed outcome.

In this example, none of the action pairs satisfy both players, leading to **an endless loop** of continuous deviations. This means that no pure strategy Nash Equilibrium exists.

However, instead of choosing a single action, players can **randomise** their choices using a mixed strategy. To explore this concept further, we introduce Mixed Strategies and their role in Nash Equilibrium. [3]

2.4 Introduction to Mixed Strategy

Mixed strategies introduce **unpredictability** in games like rock-paper-scissors. By randomising choices, players prevent opponents from exploiting predictable patterns, maximising their chances of winning.

Key Notation

- $\Pi(A)$: The set of all probability distributions over the action set A
- Example:

$$\Pi(A) = \{0.4, 0.6\}$$

This indicates that the first action occurs with a probability of 0.4, while the second action occurs with a probability of 0.6.

Definition of Mixed Strategy

A mixed strategy S_i represents the set of all probability distributions over the actions available to Player i.

Example: If Player 1 has actions $\{rock, paper, scissors\}$ and chooses *paper* and *scissors* with equal probabilities, their mixed strategies are:

$$S1 = \{0, 0.5, 0.5\}$$

Mathematicians use the product symbol Π to reflect the structure of **probability sets**. Hence, a mixed strategy S_i is often denoted as:

$$S_i = \Pi(A_i)$$

Definition of Mixed Strategy Profile

A mixed strategy profile gives an overview of **all chosen strategies** in a game (similar to a pure strategy action profile). A mixed strategy profile is defined using the **Cartesian product** of the individual mixed strategy sets:

$$S_1 \times S_2 \times \cdots \times S_n.$$

Cartesian Product of Two Sets

The Cartesian product of two sets A and B is the set of all ordered pairs $(a,\, b)$, where a is from A and b is from B:

$$A \times B = \{(a,\, b) \,|\, a \in A,\, b \in B\}$$

Example:

- $A = \{x, y, z\}$
- $B = \{1, 2, 3\}$

We can use the Cartesian product to present the pairings of all the elements in the two sets visually:

		B		
		1	**2**	**3**
A	**x**	x, 1	x, 2	x, 3
	y	y, 1	y, 2	y, 3
	z	z, 1	1, -1	z, 3

Figure 10: Cartesian Product

Then:

$$A \times B = \{(x, 1),\, (x, 2),\, (x, 3),\, (y, 1),\, (y, 2),\, (y, 3),\, (z, 1),\, (z, 2),\, (z, 3)\}$$

Similarly, for three sets:

$$A \times B \times C = \{(a, b, c) \,|\, a \in A,\, b \in B,\, c \in C\}$$

Applying Cartesian Products to Mixed Strategy Profiles

For **two players**, the mixed strategy profile is:

$$S_1 \times S_2 = \{(s_1, s_2) \,|\, s_1 \in S_1,\, s_2 \in S_2\}.$$

Where: s_1 and s_2 are the **specific strategies** used by Player 1 and 2.

Example

If:

- $S_1 = (0.3,\, 0.7)$
- $S_2 = (0.4,\, 0.6)$

Then the Cartesian product is:

$$S_1 \times S_2 = \{(0.3, 0.4), (0.3, 0.6), (0.7, 0.4), (0.7, 0.6)\}$$

A player's **probability** of choosing action ai is denoted by:

$$s_i(a_i).$$

For instance, if:

$$s_2(a_1) = 0.3,$$

then Player 2 chooses a_1 with a 0.3 probability.

*A pure strategy is **a special case of a mixed strategy**, where one action has a 100% probability while all others have zero probability.

We also use the Cartesian product to represent the set of all outcomes A:

$$A = A_1 \times A_2 \times \cdots \times A_n.$$

Recall that A_i is the finite set of actions available to player i. By taking every possible combination of actions across all players, we construct the set A which represents all possible action profiles (i.e., game outcomes).

2.5 Calculating the Expected Payoff and Nash Equilibrium

To compute the **expected payoff of a mixed strategy**, follow these steps:

1. List **all possible action profiles**.
2. Calculate the **joint probabilities** of each profile occurring.
3. Multiply each **payoffs** by its corresponding probability.
4. Sum **all weighted payoffs** to obtain the total expected payoff.

Mathematical Definition

Given an **n-player normal-form game**, the expected payoff u_i for Player i is:

$$u_i(s) = \sum_{a \in A} u_i(a) \prod_{j=1}^{n} s_j(a_j).$$

Breaking this equation into components:

1. $\sum_{a \in A}$ (**Listing all possible outcomes**)

   The symbol $\sum_{a \in A}$ means 'sum over all elements a in the set A'. In the context of game theory, A represents the set of all action profiles, meaning we **sum up all possible combinations of actions chosen by players**. By doing so, we ensure every possible outcome contributes to the player's overall expected payoff.

2. $\prod_{j=1}^{n} s_j(a_j)$ (**Calculating joint probabilities**)

   The term $\prod_{j=1}^{n} s_j(a_j)$ represents **the probability of the action profile a occurring**. By multiplying the individual probabilities $s_j(a_j)$ of each player choosing their respective action a_j, we obtain the joint probability that the entire action profile is played in the game.

3. $u_i(a) \prod_{j=1}^{n} s_j(a_j)$ (**Multiplying by payoffs**)

   Each action profile's payoff is multiplied by its probability. By combining the payoff function $u_i(a)$ with the probability of a occurring, we calculate **the expected contribution of each action profile** to Player i's total payoff.

4. **Summing over all payoffs**

   The total expected payoff is computed by summing over all weighted payoffs.

Example Calculation

Consider a two-player game where:

- Player 1's actions $= \{a, b\}$,
- Player 2's actions $= \{x, y\}$,
- Player 1's payoff matrix (u_1):

		Player 2	
		x	y
Player 1	a	3	1
	b	0	2

Figure 11: Player 1's Payoff Matrix u_1

- Player 1's mixed strategy s_1: $(P(a), P(b)) = (0.3, 0.7)$.
- Player 2's mixed strategy s_2: $(P(x), P(y)) = (0.25, 0.75)$.

Step 1: Identify All Possible Action Profiles

$$A_1 \times A_2 = \{(a, x), (a, y), (b, x), (b, y)\}$$

Step 2: Calculate Joint Probabilities

$$\begin{aligned}
P(a, x) &= s_1(a) \times s_2(x) = 0.3 \times 0.25 = 0.075 \\
P(a, y) &= s_1(a) \times s_2(y) = 0.3 \times 0.75 = 0.225 \\
P(b, x) &= s_1(b) \times s_2(x) = 0.7 \times 0.25 = 0.175 \\
P(b, y) &= s_1(b) \times s_2(y) = 0.7 \times 0.75 = 0.525
\end{aligned}$$

Note that the probabilities add up to 1.

Step 3: Compute Weighted Payoffs

$$u_1(a,x) \times P(a,x) = 3 \times 0.075 = 0.225$$
$$u_1(a,y) \times P(a,y) = 1 \times 0.225 = 0.225$$
$$u_1(b,x) \times P(b,x) = 0 \times 0.175 = 0$$
$$u_1(b,y) \times P(b,y) = 2 \times 0.525 = 1.05$$

Step 4: Sum Weighted Payoff

$$u_1(s) = 0.225 + 0.225 + 0 + 1.05 = 1.5.$$

Result: Player 1's expected payoff under mixed strategy profile s is: 1.5.

So, how do we uncover a game's Nash Equilibrium in mixed strategies?

Finding the Nash Equilibrium

We previously denoted Player i's best response as a_i^* for **pure strategies**. Now, we extend this to mixed strategies, where:

- s_i represents Player i's **chosen strategy**.
- s_i^* represents Player i's **best response**, yielding the highest expected payoff.

Thus, the mixed strategy Nash Equilibrium condition is:

$$u_i(s_i^*, s_{-i}) \geq u_i(s_i, s_{-i}^*)$$

This states that Player i's best response payoff is **at least as good as any other mixed strategy**, given that the other players are using their best strategy.

Infinite Best Responses in Mixed Strategies

However, interestingly, unlike pure strategies, a mixed strategy may have **infinitely many best responses** unless a unique pure strategy best response exists.

This occurs because players must be indifferent between all actions in a mixed strategy (the player values chosen actions equally). Otherwise, they would rationally **reduce the probability of choosing one or more actions to zero**.

In other words, for a mixed strategy to be valid, the payoff for each action in the mix should be the same. Otherwise, the player would always choose the action with the higher payoff, making the strategy non-mixed.

Why Are There Infinite Best Responses?

The purpose of a mixed strategy is to introduce randomness into decision-making. This means a player can randomise their actions using any probability distribution that keeps their expected payoff unchanged.

Since probability distributions are continuous, a mixed strategy can have infinitely many possible best responses:

$$P(action1) = p, \quad P(action2) = 1 - p, \quad where\ 0 < p < 1.$$

Therefore, any probability mix that keeps the expected payoffs the same is a valid best response.

Example Calculation

Consider the same game from earlier, with mixed strategies:

$$u_1(a,x) = 3, \quad u_1(a,y) = 1, \quad u_1(b,x) = 0, \quad u_1(b,y) = 2.$$

$$s_1 = (p_a, p_b) = (0.3, 0.7).$$

$$s_2 = (q_x, q_y) = (0.25, 0.75).$$

Using our earlier payoff calculation, we found that Player 1's expected payoff is:

$$u_1(s) = 1.5.$$

Now, let's change Player 1's strategy to:

$$s_1 = (0.1, 0.9).$$

Using the same equation:

$$u_1(s) = (0.1 \times 0.25 \times 3) + (0.1 \times 0.75 \times 1) + (0.9 \times 0.25 \times 0) + (0.9 \times 0.75 \times 2) = 1.5$$

Since the payoff remains the same, this confirms that multiple mixed strategies yield the same expected payoff, proving that **there are infinitely many best responses**.

This specific example also implies that a mixed strategy approach would be sensible for Player 1 as they would be indifferent between actions a and b.

A More Efficient Method for Finding Mixed Strategy Equilibria

Instead of testing multiple probability distributions, we can compare the expected payoffs of individual actions.

Payoff for the Action a:

$$u_1(a) = P(x) \times u_1(a,x) + P(y) \times u_1(a,y)$$

$$u_1(a) = 0.25 \times 3 + 0.75 \times 1 = 1.5$$

Payoff for Action b:

$$u_1(b) = P(x) \times u_1(b,x) + P(y) \times u_1(b,y)$$

$$u_1(b) = 0.25 \times 0 + 0.75 \times 2 = 1.5$$

Since:

$$u_1(a) = u_1(b),$$

Player 1 is indifferent between these actions, confirming that they are **valid components** of a mixed strategy Nash Equilibrium.

Defining the Best Response Function

Since mixed strategies allow for multiple best responses, we define the **best response function** for Player i as:

$$B_i(s_{-i}) = \arg\max_{s_i \in S_i} \; u_i(s_i, s_{-i})$$

Where:

- $B_i(s_{-i})$ is Player i's best response function.
- It finds the strategy s_i^* that maximises Player i's expected payoff $u_i(s_i, s_{-i})$.
- $argmax$ stands for $theargumentsofthemaxima$, and means **finding the strategy** (not just the maximum value) **that achieves the highest payoff**.

Final Condition for Nash Equilibrium

A strategy profile $s = (s_1, \cdots, s_n)$ is a Nash Equilibrium if, for all players i:

$$s_i^* \in B_i(s_{-i}^*) \quad \forall i$$

This means that **every player is playing their best response** to the others' strategies.

Players reach a Nash Equilibrium when, for **all Player** i, their strategy yields the highest possible expected payoff, given the strategies of others.

Having calculated expected payoffs, we now demonstrate how the existence of Nash Equilibria is guaranteed in every finite game.

2.6 Proving the Existence of Nash Equilibria

Introduction

A Nash Equilibrium is a **strategy profile** where no player can improve their payoff by unilaterally deviating from their strategy. This concept is crucial in real-life decision-making, as it indicates stable strategic interactions, allowing us to predict rational behaviour.

A fundamental question arises: **Does a Nash Equilibrium always exist?**

In his Ph.D. thesis, John Nash proved that at least one equilibrium exists in every finite game, revolutionising game theory. His proof relies on **Brouwer's Fixed-Point Theorem**, a powerful tool that ensures the existence of fixed points under specific conditions.

What is a Fixed Point?

A fixed point of a function $F(x)$ is a point where **the input equals the output**:

$$F(x) = x.$$

The function $y = x^3$ has three fixed points where it intersects the line $y = x$ three times:

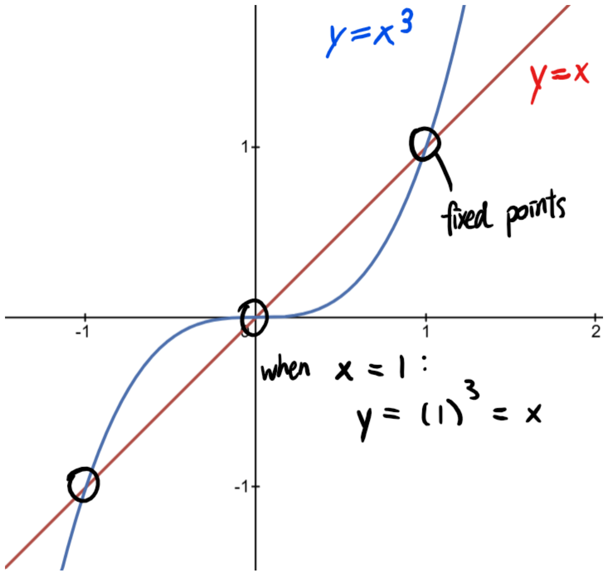

Figure 12: Illustration of Fixed Points: $y = x^3$ intersects $y = x$ three times

Example:

Imagine a function that maps study hours to test scores. A fixed point occurs when **the score equals the number of hours studied**.

Why Fixed Points Relate to Nash Equilibria

Intuitively, a Nash Equilibrium is a fixed point of best response functions.

- If we apply the **best response function** to **a strategy profile**, and the input equals the output, then no player wishes to change strategies.
- This aligns **exactly** with the definition of Nash Equilibrium:
  - Each player is playing a **best response** to the others' strategies.
  - **No player wants to deviate**.

Now that we have established this intuition, we analyse how mixed strategies satisfy **the conditions** for a fixed-point theorem.[4]

Brouwer's Fixed-Point Theorem

Brouwer's theorem guarantees that a continuous function mapping a compact, convex set to itself **must have** at least one fixed point. [5]

Conditions for the Theorem:

1. Compactness: The function's domain must be **closed** and **bounded**.
2. Convexity: Any line segment between two points in the set must also lie inside the set.

We now verify that **all** mixed strategy sets meet these conditions.

The Standard n-Simplex - Compact and Convex

The standard n-simplex generalises a triangle to higher dimensions in geometry:

- A 2D simplex is a triangle.
- A 3D simplex is a tetrahedron.

These are the only four simplexes that can be represented in 3D:

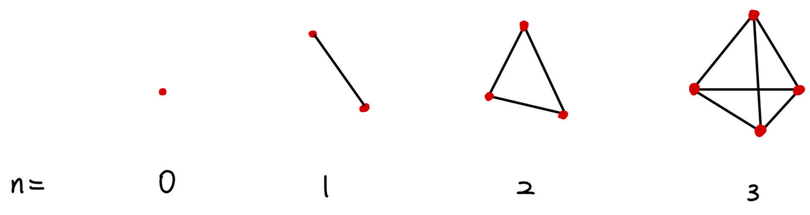

Figure 13: Four Simplexes That Can Be Represented in 3D

Mathematical Definition:

$$\Delta_n = \{y \in \mathbb{R}^{n+1} : \sum_{i=0}^{n} y_i = 1, \quad y_i \geq 0 \quad \forall i\}$$

Breaking Down the Equation:

- Δ_n: **The standard n-simplex**.
- $y \in \mathbb{R}^{n+1}$: Each strategy exists in **n+1 dimensional space**.
- $\sum_{i=0}^{n} y_i = 1$: **the probability constraint** (total probability of actions is 1).
- $y_i \geq 0$: Probabilities are **non-negative**.
- $\forall i$: For all values of i.

For a mixed strategy to be valid, it must satisfy the two constraints:

$$\sum_{i=0}^{n} y_i = 1 \quad and \quad y_i \geq 0.$$

Hence, valid mixed strategies can be represented as **points in an n-simplex**.

Example: A Player with Three Actions ($n = 3$)

Consider a mixed strategy:

$$(0.1, 0.4, 0.5).$$

We verify the constraints:

$$\sum_{i=0}^{n} y_i = y_0 + y_1 + y_2$$

$$\sum_{i=0}^{n} y_i = 0.1 + 0.4 + 0.5 = 1,$$

$$0.1, 0.4, 0.5 \geq 0$$

Since both conditions hold, the strategy (0.1, 0.4, 0.5) lies inside a **2D simplex** (triangle) in **3D space**.

Why is the Simplex in Rn+1?

Since an n-simplex has n+1 vertices, each vertex is placed on a separate axis in $\mathbb{R}^{n+1}$ for easier calculations.

Here is an example of a mixed strategy represented on a 2-simplex in 3D:

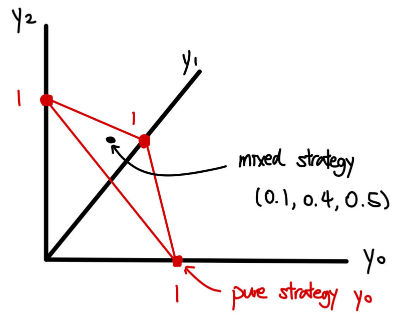

Figure 14: Example of a Mixed Strategy in 3D

Convexity of the Simplex?

A set is convex if, for any two points, **the line segment between them also lies inside the set**.

Mathematical Definition:

A set S is convex if:

$$\forall x, y \in S, \quad \lambda \in [0, 1], \quad \lambda x + (1 - \lambda) y \in S.$$

The intuition is that mixing two valid strategies with weightings produces **another valid strategy**.

Example in a 2-Simplex (Triangle)

Consider the two mixed strategies shown in Figure 15:

- (0.7, 0.2, 0.1) (blue dot)
- (0.2, 0.6, 0.2) (red dot)

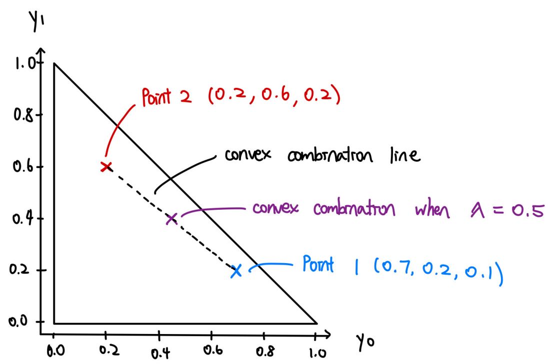

Figure 15: Convexity of a 2-Simplex: Mixing two strategies forms another valid strategy

(The example of a 2-simplex is in 2D instead of 3 for a more transparent explanation.)

Since we are dealing with a 2D region, we can **ignore the last action for now**, but the principle still applies.

The vertices (1, 0), (0, 1) and (0, 0) represent **pure strategies**, where the player has a probability of 1 playing the respective action. Any point inside the simplex represents a mixed strategy.

The **convex combination** connects the two mixed strategies (dashed black line), where $\lambda = 0.5$:

$$\lambda(0.7, 0.2) + (1 - \lambda)(0.2, 0.6)$$
$$= 0.5(0.7, 0.2) + 0.5(0.2, 0.6)$$
$$= (0.5 \times 0.7 + 0.5 \times 0.2,\ 0.5 \times 0.2 \times 0.5 \times 0.6)$$
$$= (0.45, 0.4)$$

Since (0.45, 0.4) (purple dot) is **inside the simplex**, the set is convex.

Applying this to a general simplex with n **pure strategies**, we can form the convex combination of two mixed strategies:

- $y = (y_0, y_1, \cdots, y_n)$.
- $z = (z_0, z_1, \cdots, z_n)$.

$$\lambda y + (1 - \lambda)z = (\lambda y_0 + (1 - \lambda)z_0, ..., \lambda y_n + (1 - \lambda)z_n)$$

Given that the two mixed strategies (y, z) are valid, we can use the probability constraint $\sum y_i = 1$ **and** $\sum z_i = 1$ to simplify the convex combination:

$$\sum_{i=0}^{n}(\lambda y_i + (1 - \lambda)z_i)$$
$$= \lambda \sum y_i + (1 - \lambda) \sum z_i$$
$$= \lambda(1) + (1 - \lambda)(1) = 1$$

We can use the fact that both y_i and z_i are non-negative ($y_i \geq 0$, $z_i \geq 0$) to deduce that their convex combination is also non-negative.

The two conditions are met, meaning that the convex combination is also a valid mixed strategy. Therefore, we have proved that the standard n-simplex is convex.

Boundedness and Closedness of the Simplex?

1. A set is bounded if all points lie within a **finite distance** of the origin.
2. A set is closed if it contains all its **limit points**.

For example, the function $y = x^3$ is unbounded because both its domain and range extend to infinity, whereas $y = x$ over the interval $\{-2 \leq x \leq 2\}$ is bounded:

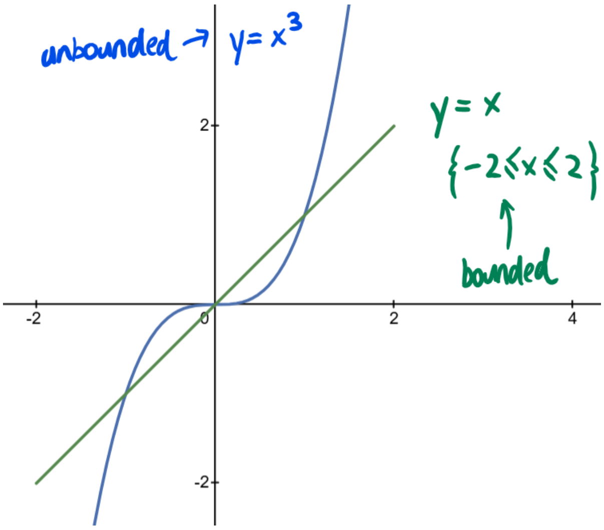

Figure 16: Bounded and Unbounded Functions

Recall that for a valid mixed strategy y, the probability constraint is met ($\sum y_i = 1$), and we can deduce that **no single coordinate** y_i **can exceed 1**:

$$0 \geq y_i \geq 1$$

Since every point in the n-simplex exists is **constrained**, we know that the simplex is bounded.

In a mixed strategy, players adjust strategies **over time**, converging to a limit.

If each term in a sequence $y(1), y(2), \cdots, y(k)$ represents a valid mixed strategy, we denote that the player is adjusting their strategy over time (k).

Therefore, as k approaches infinity, the sequence converges to a best response y^*, then:

$$\lim_{k \to \infty} y(k) = y^*.$$

Since the convergence can only occur with **valid mixed strategies**, all $y(k)$ must satisfy the two probability constraints, the limit y^* must also fulfil the conditions and the simplex is closed.

Applying Brouwer's Fixed-Point Theorem

For a Nash Equilibrium, we define a continuous function $F(s)$ that maps strategy profiles to the best responses:

$$F(s) = (B_1(s_{-1}), B_2(s_{-2}), ..., B_n(s_{-n})).$$

Essentially, $F(s)$ **takes in different strategy profiles** and **outputs the respective best response** that yields the player the highest payoff.

Since:

- $F(s)$ is continuous,
- The mixed strategy simplex is compact and convex,

Brouwer's theorem guarantees that $F(s)$ has **at least one fixed point**.

Conclusion

Since fixed points of $F(s)$ correspond to Nash Equilibria, we conclude:

$$s^* \in B(s^*),$$

will **always** be true.

Thus, at least one Nash Equilibrium exists in every finite game. [6]

3 Applications in Real Life

3.1 Introduction to Economic Applications

The Nash Equilibrium Theorem states that at least one Nash equilibrium exists in every finite game. But how does this concept apply to **real-life** strategic interactions?

Nash Equilibrium is widely used in:

- Economics: To model **competitive relationships** between firms
- Politics: To model **electoral strategies** of candidates.
- Biology: To model **evolutionary competition** between species.

This section explores how economists and businesses apply pricing strategies using **the principle of equilibrium**. Understanding the Nash Equilibrium allows economists to predict firm behaviour in competitive markets, assisting in strategic decision-making. [7] [8]

Revisiting our initial question, **how do rational firms apply strategies to maximise profit while considering their rivals?** To answer this, we construct a mathematical model to formulate the interdependent behaviour of firms in a competitive market.

3.2 Oligopoly

An oligopoly is a **market structure** where a small number of firms dominate, leading to strategic interdependence. This market structure is distinct as:

- Few firms exist, each with significant market power.
- High barriers to entry **prevent new firms from easily entering**.
- Products may be **identical or different**.

Since no single firm has absolute dominance, firms **anticipate and react** to their competitors' decisions.

In an oligopoly, since the firms exert roughly equal influence over a market, companies within the oligopoly collude with one another, keeping non-established players from invading the market. [9]

Example: The Airline Industry

Consider airlines operating direct flights between London Heathrow Airport and Hong Kong International Airport. This market is an oligopoly because:

- Few airlines operate direct flights due to strict regulations and high investment costs.
- Each airline's pricing strategy directly affects its competitors.
- Airlines react strategically to competitors' price changes.

If Airline A introduces **seasonal discounts**, it may attract more customers, but this forces other airlines to adjust their pricing to maintain their market share. This leads to a **cycle of strategic decision-making**, a defining characteristic of oligopoly markets.

Strategic Interaction in Oligopoly Markets

Firms in an oligopoly must predict and respond to their competitors' actions.

- If Airline A lowers its prices, other airlines will rationally reduce their prices to retain customers.
- If Airline A increases prices, competitors may either follow (if they have brand loyalty) or undercut (to steal market share)

This continuous counter-strategy process not only is **nothing like an endless loop**, but it also leads firms to a stable outcome - a Nash Equilibrium where no airline can unilaterally change its pricing to gain an advantage.

Thus, we can model oligopolistic markets as a normal-form game, where:

- Firms are the **players**.
- Pricing strategies are the **actions**.
- Profits are the **payoffs**.

In this game, at least one Nash Equilibrium exists, determining the stable price point that firms ultimately settle on.

To locate an equilibrium state in oligopolistic competition, we use **the classic Cournot Model**. The model:

- **Mathematically** represents firm interactions.
- Shows how firms choose optimal production levels **based on their rival's choices**.
- Leads to a Nash Equilibrium in output and pricing.

We now introduce the Cournot Model to quantify strategic decision-making in oligopolies.

3.3 The Cournot Duopoly Model

The Cournot Model, developed by Augustin Cournot, is a fundamental model in industrial organisation and microeconomics. It describes how firms compete by choosing production quantities rather than setting prices.

Key Assumptions of the Cournot Model

- Firms produce **identical goods**.
- Each firm **independently** decides its output, assuming its competitor's output remains constant.
- **The market price is determined by total output**, following a linear inverse demand function.

The Nash Equilibrium in this model represents an optimal state where neither firm has an incentive to adjust its production.

Consider a duopoly where two firms, Firm 1 and Firm 2, produce a homogeneous (identical) good. The total market output is:

$$Q = Q_1 + Q_2$$

where:

- Q_1 is Firm 1's output.
- Q_2 is Firm 2's output.

The market price is determined by an **inverse demand function**:

$$P(Q) = a - bQ$$

$$or \quad P(Q) = a - bQ_1 - bQ_2$$

where:

- a is the **maximum** price consumers are willing to pay for the good.
- b is the sensitivity of how the price changes with supply (both are positive constants).

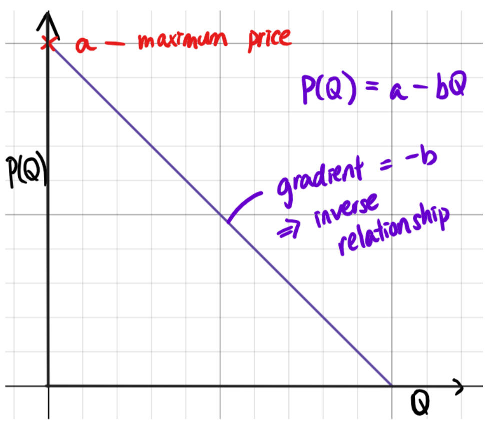

Figure 17: Market Demand Function $P(Q) = a - bQ$

From the graph above, when no goods are produced ($Q = 0$), the **price is at its maximum** ($P = a$). Conversely, when total production reaches a maximum limit, the **price falls to zero**, meaning the market is saturated.

Thus, the market demand function determines **the price** firms receive for their goods.

Economic Intuition: Why is Demand Inversely Related to Output?

- If total output increases, supply exceeds demand, causing the price to fall.
- If total output decreases, supply is limited, causing the price to rise.

Profit Maximisation in the Cournot Model

The profit for Firm i (Π_i), is calculated as the difference between revenue and costs:

$$\Pi_i = P(Q)Q_i - C(Q_i)$$

Where:

- $P(Q)$ is the market price.
- $C(Q_i)$ is the cost function of Firm i, representing costs required for production.

Assuming both firms have constant marginal costs (same cost for one product), the cost function of Firm i is directly dependent on the quantity produced. We can then simplify the cost function:

$$C(Q_i) = c_i Q_i$$

Where c_i is the marginal cost for Firm i.

Substituting $P(Q) = a - bQ_1 - bQ_2$ and $C(Q) = c_i Q_i$ into the profit equation:

$$\Pi_i = P(Q)Q_i - C(Q_i)$$

$$\Pi_i = (a - bQ_1 - bQ_2)Q_i - c_i Q_i$$

$$\Pi_1 = (a - bQ_1 - bQ_2)Q_1 - c_1 Q_1$$

$$\Pi_2 = (a - bQ_1 - bQ_2)Q_2 - c_2 Q_2$$

Finding the Best Response Functions

To maximise profit, each firm takes the partial derivative of its profit function with respect to its respective output.

By using the product rule or expanding the profit equation, we get:

$$\frac{\delta \Pi_1}{\delta Q_1} = a - 2bQ_1 - bQ_2 - c_1$$

$$\frac{\delta \Pi_2}{\delta Q_2} = a - 2bQ_2 - bQ_1 - c_2$$

Setting this equal to zero, we get the maximum profit (as the second derivative is $-2b$) :

$$0 = a - 2b(Q_1) - b(Q_2) - c_1$$

$$Q_1 = \frac{a - c_1 - bQ_2}{2b}$$

For Firm 2:

$$Q_2 = \frac{a - c_2 - bQ_1}{2b}$$

These are the **reaction functions**, showing how each firm adjusts its output in response to its competitor. The negative coefficients indicate the **inverse relationship** between the optimal production. In other words, the more one firm produces, the less the other should produce to maintain profitability.

Since the reaction function $Q_1 = \frac{a-c_1-bQ_2}{2b}$ tells us Firm 1's optimal output, which varies with Firm 2's production, we can also use Q_1^* to represent the best response function.

Example: Cournot Competition Between Coca-Cola and Pepsi

Consider a duopoly market where Coca-Cola and Pepsi compete. Suppose:

$$a = 10, \quad b = 1, \quad c_1 = 1, \quad c_2 = 1.$$

Substituting these into the **reaction functions**:

$$Q_1 = \frac{10 - 1 - Q_2}{2(1)}$$

$$Q_2 = \frac{10 - 1 - Q_1}{2(1)}$$

$$Q_1 = \frac{9}{2} - \frac{1}{2}Q_2$$

$$Q_2 = \frac{9}{2} - \frac{1}{2}Q_1$$

Analyse the implications of each firm's output on the other's:

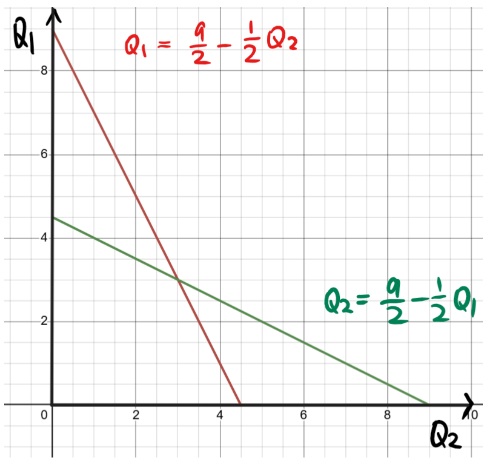

Figure 18: Reaction Functions of the Two Firms

From the graph, when Q_2 is zero (Pepsi produces nothing), Coca-Cola will rationally produce $\frac{9}{2}$ units to maximise profit ($Q_1 = \frac{9}{2} - \frac{1}{2}(0) = \frac{9}{2}$).

Now, Coca-Cola is earning the maximum profit while Pepsi suffers from a production of 0. Therefore, Pepsi responds and utilises the best response function to produce $\frac{9}{4}$ units ($Q_2 = \frac{9}{2} - \frac{1}{2}(\frac{9}{2}) = \frac{9}{4}$).

Then, Coca-Cola utilises its best response function to counter Pepsi's refinement. From: $Q_1 = \frac{9}{2} - \frac{1}{2}(\frac{9}{4}) = \frac{27}{8}$, we know that Coca-Cola will produce $\frac{27}{8}$ units to keep the competition alive.

This seemingly endless loop resolves at the equilibrium point, where neither firm has an incentive to deviate

The 'deviation path' continues and eventually terminates **when the two functions intercept**:

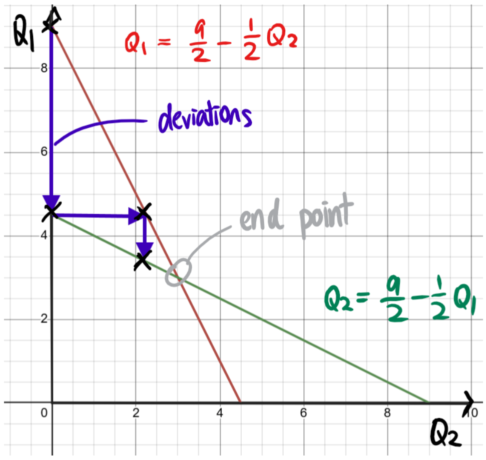

Figure 19: The Deviation Path in Cournot Competition Converges to Equilibrium

Finding the Nash Equilibrium

The interceptions between the two best response functions have a significant property - **none of the firms would want to deviate** from a different response. With this in mind, we can solve the simultaneous equations to find the optimal equilibrium between the outputs of Coca-Cola and Pepsi.

$$Q_1 = \frac{9}{2} - \frac{1}{2}Q_2$$

$$Q_2 = \frac{9}{2} - \frac{1}{2}Q_1$$

Substituting Q_2 into Q_1:

$$Q_1 = \frac{9}{2} - \frac{1}{2}(\frac{9}{2} - \frac{1}{2}Q_1)$$
$$Q_1 = \frac{9}{2} - \frac{9}{4} + \frac{1}{4}Q_1$$

$$\frac{3}{4}Q_1 = \frac{9}{4}$$
$$Q_1 = 3$$

Similarly:

$$Q_2 = \frac{9}{2} - \frac{1}{2}(3) = 3$$

Thus, the **Cournot-Nash Equilibrium** occurs when both firms produce 3 units each.

At this equilibrium:

- Neither firm has an incentive to deviate
- Both firms maximise profit given the competitors' strategy.
- Market price stabilises, **preventing a price war or oversupply**.

Generalising the Cournot-Nash Equilibrium

For any two firms, we derive a general formula:

$$Q_1^* = \frac{a - c_1 - bQ_2^*}{2b}$$
$$Q_2^* = \frac{a - c_2 - bQ_1^*}{2b}$$

Solving the simultaneous equations, we get:

$$Q_1^* = \frac{a - c_1 - bQ_2^*}{2b}$$
$$Q_1^* = \frac{a - c_1 - b(\frac{a-c_2-bQ_1^*}{2b})}{2b}$$
$$Q_1^* = \frac{2a - 2c_1 - a + c_2 + bQ_1^*}{4b}$$
$$Q_1^* = \frac{a - 2c_1 + c_2}{3b}$$

$$\text{Therefore:} \quad Q_2^* = \frac{a - 2c_2 + c_1}{3b}$$

This formula determines the **optimal production levels** for firms in a Cournot duopoly.

Conclusion: Understanding the Cournot Model in Business Strategy

The Cournot-Nash Equilibrium helps firms:

- Determine optimal production levels to maximise profit.
- Predict competitors' responses to production changes.
- Avoid overproduction and prevent price wars.

Next, we explore how firms use equilibrium models to analyse real-world market factors and competitive advantages.

3.4 Comparative Statics & Economic Implications

The discovery of the Nash Equilibrium enables us to analyse **how changes in market conditions affect equilibrium outcomes** in oligopolistic competition. By differentiating key parameters, we can examine the impact of **demand elasticity** (how sharply the price reacts with the change in supply), **costs**, and market structure on output and pricing decisions.

Total Market Output in Cournot-Nash Equilibrium

The total market output in a Cournot duopoly is:

$$Q^* = Q_1^* + Q_2^*$$

Substituting the previously derived equilibrium quantities:

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

$$Q^* = \frac{(a - c_1) + (a - c_2)}{3b}$$

$$Q^* = \frac{2a - (c_1 + c_2)}{3b}$$

Economic Interpretation

Effect of Demand Elasticity (b)

- If b **increases**, demand becomes more elastic, meaning that the price reacts more sharply to changes in output.
- Firms respond by reducing production, leading to **lower total output**.

Effect of Marginal Costs (c_1, c_2)

- If c_1 or c_2 **increases**, the respective firm's profitability declines, leading to lower individual output.
- This **reduces total market output**, increasing market price.

Equilibrium Market Price

Using the market demand function:

$$P^* = a - bQ^*$$

Substituting Q^*

$$P^* = a - b\frac{(a - c_1) + (a - c_2)}{3b}$$

$$P^* = \frac{a + (c_1 + c_2)}{3}$$

Economic Implications

- Demand elasticity (b) does **not** directly affect the equilibrium price
- Higher production costs (c_1, c_2) result in **higher equilibrium prices**, as firms produce less, reducing supply.

Thus, firms operating under the Cournot Model must **balance cost efficiency and production levels** to remain competitive.

3.5 Extensions to n Firms

The Cournot Model can be extended to markets with n firms, allowing for a broader analysis of oligopolistic competition.

Assumptions:

1. n firms compete, indexed as $i = 1, 2, \cdots, n$.
2. Firms produce identical goods.
3. **Each firm faces the same marginal cost** (c) - for easier calculations.

Market Output and Demand Function

Total market output:

$$Q = \sum_{i=1}^{n} Q_i$$

Market demand function:

$$P(Q) = a - bQ$$

Profit function for Firm i:

$$\Pi_i = P(Q)Q_i - c(Q_i)$$

Where c is the identical marginal cost for all firms.

Substituting $P(Q) = a - bQ$:

$$\Pi_i = (a - bQ)Q_i - c(Q_i)$$

Finding the Nash Equilibrium for n Firms

To maximise profit, we take the partial derivative:

$$\frac{\delta\, \Pi_i}{\delta\, Qi} = (a - bQ) - bQ_i - c$$

Setting the derivative to zero:

$$0 = a - bQ - bQ_i - c$$

Since all firms are identical, all firms produce the same quantity at equilibrium:

$$Q_1^* = Q_2^* = \cdots = Q_n^*$$

Thus, the total market output is:

$$Q^* = nQ_i^*$$

Substituting this into the equilibrium condition with profit maximised:

$$0 = a - b(nQ_i^*) - bQ_i^* - c$$

$$0 = a - c - b(n + 1)Q_i^*$$

For **equilibrium output per firm**:

$$Q_i^* = \frac{a-c}{(n+1)b}$$

As a result, the total market output in n-firm Cournot competition:

$$Q^* = nQ_i^*$$

$$Q^* = n\frac{a-c}{(n+1)b}$$

Equilibrium Market Price

Substituting Q^* into the demand function:

$$P^* = a - bQ^*$$

$$P^* = a - b\,(n\frac{a-c}{(n+1)b})$$

$$P^* = \frac{a+nc}{n+1}$$

Economic Implications

So, what can we denote from this?

Effect of Market Size (n)

- As n increases, each firm's **equilibrium output decreases**.
- Total market output increases as more firms enter the market.
- As n approaches infinity, **the equilibrium market price P^* converges to the marginal cost** c:

$$\lim_{n\to\infty} P^* = c\,,$$

reflecting a shift toward **perfect competition**:

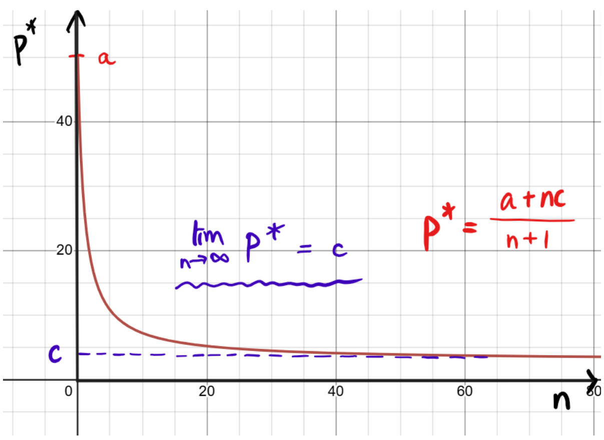

Figure 20: Market Price Convergence: As n increases, the price approaches marginal cost

Effect of Marginal Costs (c)

Higher c leads to **lower total output** and a **higher equilibrium price**.

Verification with Two-Firm Case

For $n = 2$:

$$Q_1^* + Q_2^* = \frac{(a-c)+(a-c)}{3b}$$

$$Q_1^* + Q_2^* = 2\frac{a-c}{(2+1)b}$$

This matches our two-firm Cournot model, confirming the general equation is valid.

To summarise,

- Firms adjust production based on market conditions.
- Higher costs lead to reduced output and higher prices.
- As the number of firms increases, competition drives prices closer to the marginal cost.

Economists and businesses use these models to **optimise production decisions**, **predict competitive behaviour**, and **design pricing strategies**.

4 Conclusion

The Nash Equilibrium stands as a cornerstone of game theory, demonstrating **how rational decision-makers navigate strategic interactions to optimise outcomes**. The concept is particularly relevant in everyday situations, such as **deciding whether to introduce discounts in response to competition or choosing a shortcut to avoid traffic during rush hour**. Just as taxi drivers adjust their routes of other vehicles to minimise their travel time, businesses must evaluate their rival's actions to maximise profits.

The Cournot Model provides a structured framework to explore these strategic interplay, particularly in oligopolistic markets. By analysing how firms determine optimal production levels, the model reveals key economic insights:

- The impact of **demand elasticity** on total output.
- The influence of **production costs** on pricing strategies.
- The effect of **market entry**, where new competitors force existing firms to recalculate their output.

Extending the Cournot framework to a larger number of firms reveals that as competition intensifies, market behaviour converges towards perfect competition, **where market prices reflect marginal costs rather than being dictated by a dominant firm**.

Ultimately, the Nash Equilibrium offers valuable insights into human interactions, guiding towards strategies that **maximise rewards while minimising risk**. Whether it's:

- Coordinating efforts in a **stag hunt**,
- Optimising moves in **rock-paper-scissors**, or
- Strategising production levels in a **competitive market**,

the principles of strategic thinking and mutual adaptation are embedded into our daily lives.

As we delve deeper into **economic behaviour** and **decision-making**, the Nash Equilibrium will remain a crucial tool for understanding stability, equilibrium states, and rational strategy selection. With its practical applications across multiple sectors, adopting Nash Equilibrium principles can help us **formulate decisions more effectively**, paving the way for **win-win scenarios.**

References

[1] Harold W Kuhn, John C Harsanyi, Reinhard Selten, J Weibull, and Eric Van Damme. The work of john nash in game theory. *journal of economic theory*, 69(1):153–185, 1996.

[2] Avinash K Dixit and Susan Skeath. *Games of strategy: Fourth international student edition.* WW Norton & Company, 2015.

[3] Martin J Osborne. An introduction to game theory. *Oxford University Press google schola*, 2:672–713, 2004.

[4] Kousha Etessami and Mihalis Yannakakis. On the complexity of nash equilibria and other fixed points. *SIAM Journal on Computing*, 39(6):2531–2597, 2010.

[5] David Gale. The game of hex and the brouwer fixed-point theorem. *The American mathematical monthly*, 86(10):818–827, 1979.

[6] Albert Xin Jiang and Kevin Leyton-Brown. A tutorial on the proof of the existence of nash equilibria. *University of British Columbia Technical Report TR-2007-25. pdf*, 14, 2009.

[7] Roger B Myerson. Nash equilibrium and the history of economic theory. *Journal of Economic Literature*, 37(3):1067–1082, 1999.

[8] Charles A Holt and Alvin E Roth. The nash equilibrium: A perspective. *Proceedings of the National Academy of Sciences*, 101(12):3999–4002, 2004.

[9] Avinash Dixit. Recent developments in oligopoly theory. *The American Economic Review*, 72(2):12–17, 1982.