Simultaneous Games

Edicson Luna

August 15, 2025

Discrete Strategies

Pure Strategies in Simultaneous-Move Games

In a finite simultaneous-move game, a pure strategy for player i is a complete plan that prescribes *one* action from their finite action set S_i when the game is played.

$$S_i \;=\; \{s_{i1},\, s_{i2},\, \ldots,\, s_{ik}\}.$$

- **Prisoner's Dilemma:** $S_1 = S_2 = \{\text{Cooperate}, \text{Defect}\}$.
- **Rock–Paper–Scissors:** $S_i = \{\text{Rock}, \text{Paper}, \text{Scissors}\}$ for $i = 1, 2$.

Because moves are made *simultaneously*, each player must pick a pure strategy *without observing* the others' choices.
The central problem becomes: What do I believe the others will do, and which of my strategies is best compared to those beliefs?

Depicting Simultaneous-Move Games with Discrete Strategies

Simultaneous games with discrete strategies are often depicted using a **payoff matrix**. We will refer to it as the normal form representation of the game. See the following example (with nonsensical payoffs):

Table 1: Example

		P2		
		Left	**Middle**	**Right**
P1	**Top**	(3,1)	(2,3)	(10,2)
	High	(4,5)	(3,0)	(6,4)
	Low	(2,2)	(5,4)	(12,3)
	Bottom	(5,6)	(5,5)	(9,7)

Notation Clarification

In this example, we define the strategy sets as follows:
$S_1 = \{\text{Top, High, Low, Bottom}\}$ for Player 1 (P1), and
$S_2 = \{\text{Left, Middle, Right}\}$ for Player 2 (P2).

Payoffs are listed in the form (u_1, u_2), where u_1 is the payoff to P1 (the row player), and u_2 is the payoff to P2 (the column player). For instance, in the outcome (Top, Left), the payoffs are $(3, 1)$: 3 for P1 and 1 for P2.

We will also consider extensions to three-player games. In those cases, the third player is often referred to as the *matrix player*, who selects which matrix the other two players will play in. The third player's payoff will appear as the final entry in the triplet.

For games involving more than three players, we will rely on intuition and simplified representations, as full payoff matrices become impractical.

Definition: Nash Equilibrium

Recall the example in Table 1. Consider the strategy profile where P1 plays *Low* and P2 plays *Middle*. At this point, neither player has an incentive to unilaterally switch to a different strategy — each has no reason to change their choice on their own.

At (Low, Middle), players are making their choices independently, which is the definition of a noncooperative game.

- A **Nash Equilibrium (NE)** is a strategy profile such that no player can get a better payoff by switching to another available strategy, keeping the strategies of the others fixed.
- We can also say that none of the players has an incentive to *unilaterally deviate*.
- There can be games without a NE, with one NE, with multiple NE, or with infinitely many NE.

More on NE

Nash equilibria do not require a player's choice to yield a strictly better payoff than all other available options.

For example, (Low, Middle) is a NE even though P1 receives the same payoff as in (Bottom, Middle).

NEs are not necessarily jointly optimal for the players. Notice that both players would be better off at (Bottom, Right), yet this is not a NE. Why is that?

Dominance: Key Definitions

- A strategy is **strictly dominant** if it gives a player a strictly higher payoff than any other available strategy, no matter what the other players do.
- Unless otherwise specified, we will refer to strictly dominant strategies when using the term "dominant strategy."
- If a dominant strategy exists, it must be part of all Nash equilibria.
- A strategy is **dominated** if there exists another strategy that gives a strictly higher payoff in every possible scenario.
- Rational players never play dominated strategies.

Prisoners' Dilemma

		Prisoner B	
		Stay Silent	**Confess**
Prisoner A	**Stay Silent**	(–1, –1)	(–10, 0)
	Confess	(0, –10)	(–5, –5)

- *Confess* is a dominant strategy for both players.
- Result: Both confess, ending up at (–5, –5), even though (–1, –1) would be better for both.

When Only One Player Has a Dominant Strategy

Example 2: U.S. Congress – Federal Reserve Game

		Federal Reserve (FED)	
		Low Rates	**High Rates**
Congress	**Balance**	(3, 4)	(1, 3)
	Deficit	(4, 1)	(2, 2)

- The Congress has a dominant strategy: *Deficit*.
- The FED does not have a dominant strategy.
- We eliminate Congress's dominated strategy (*Balance*) and be left with *Deficit*.
- Given this, the FED chooses *High rates*.
- This is an example of **iterated elimination of dominated strategies**.

IEDS

Definition. Iterated elimination of dominated strategies (IEDS) is the iterative procedure of removing strategies that are strictly dominated (i.e. always worse than some other strategy in every possible scenario) for one player, then re-examining the reduced game and repeating for the other player, until no more dominated strategies remain.

- If IEDS reduces each player to a single strategy, that surviving profile is the *unique* Nash equilibrium of the game.
- If the game cannot be reduced completely, then every Nash equilibrium must lie within the set of strategies that survive the elimination.
- Thus, IEDS provides a way to *simplify* the game, narrowing down the candidate profiles for further equilibrium analysis.

IEDS

		P2		
		X	**Y**	**Z**
P1	**A**	(2,2)	(1,3)	(0,1)
	B	(3,1)	(5,2)	(1,0)
	C	(1,3)	(4,1)	(2,–1)

- For P1: **A** is strictly dominated by **B** (since $2<3$, $1<5$, $0<1$) $\rightarrow$ eliminate **C**. Also, for P2 **Z** is strictly dominated by **X** (and Y) $\rightarrow$ eliminate **Z**. It is key that you drop everything that can be dropped at each stage.
- In the remaining 2x2 game, for P1 **C** is dominated by **B** $\rightarrow$ eliminate **C**.
- Finally, P2 chooses **Y** since $2>1$.
- The unique NE is (**B**,**Y**).

Best-Response

- A player's **best-response** to the other players' strategies is the one that yields the highest payoff, given what the others do.
- Intuitively: look at each possible strategy of your opponent(s), and pick the action(s) that give you the top payoff in that column (or row).
- A NE is an intersection of best responses.
- We'll use this to find NE by underlining best-responses in the payoff matrix.

Prisoners' Dilemma via Underlining

	Stay Silent	Confess
Stay Silent	$(-1, -1)$	$(-10, \underline{0})$
Confess	$(\underline{0}, -10)$	$(\underline{-5}, \underline{-5})$

- In the left column, A's best response to "Stay Silent" is <u>Confess</u> (0 $>$ –1).
- In the right column, A's best response to "Confess" is <u>Confess</u> (–5 $>$ –10).
- In the top row, B's best response to "Stay Silent" is <u>Confess</u> (0 $>$ –1).
- In the bottom row, B's best response to "Confess" is <u>Confess</u> (–5 $>$ -10).
- Both underlined in the *Confess–Confess* cell it is the unique Nash equilibrium.

Dominance Concepts (Intuition)

- **Superdominance.** A strategy "superdominates" another if, even in its worst case, it outperforms the best case of any other strategy. Superdominant strategies are dominant strategies, but the converse is not true.

- **Weak dominance.** A strategy "weakly dominates" another if it never does worse, and sometimes does strictly better. Intuition: You might be indifferent in some scenarios, but there's at least one situation where it's clearly better.

- Warning: If you apply the elimination process using weak dominance, you can accidentally throw away some valid Nash equilibria.

Example: Weak IEDS Can Drop an NE

Consider this 2×2 payoff matrix:

	Left	Right
Up	$(0, 0)$	$(1, 1)$
Down	$(1, 1)$	$(1, 1)$

- **P1's choice:** "Down" never gives less than "Up," and in one case gives more $\rightarrow$ so Down weakly dominates Up $\rightarrow$ eliminate Up.
- **P2's choice (in the reduced game):** Both "Right" and "Left" survive.
- Now only (Down, Left) and (Down, Right) remain. But originally, also (Up, Right) was a NE. Weak elimination has removed the "Up–Right" equilibrium even though it was valid.

Example: Battle of the Sexes

		Girlfriend	
		Opera	**Soccer**
Boyfriend	**Opera**	(2,1)	(0,0)
	Soccer	(0,0)	(1,2)

Example: Stag Hunt

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

Example: Matching Pennies

		Player B	
		Heads	**Tails**
Player A	**Heads**	(1,–1)	(–1,1)
	Tails	(–1,1)	(1,–1)

Continuous Strategies

Motivation: When Strategies Are Not Discrete

So far, we have studied games with a finite set of choices (usually small).

But in many real-life situations, strategies are not just a small set of choices — they can take on any value in a range. Some examples include:

- A firm choosing a **price** for its product.
- A worker deciding how many **hours to work**.
- A person deciding how much **money to donate**.
- A company choosing how much to **advertise**.
- A government selecting a **tax rate**.

In these cases, we can't just draw a simple payoff table — we need new tools to analyze the game.

First-Order Derivatives

When players choose a value from a continuous range (like a price or quantity), we often want to find the value that gives the **highest possible payoff**.

A first-order derivative tells us how the payoff is changing as we slightly increase or decrease the strategy.

- If the derivative is **positive**, increasing the strategy increases the payoff.
- If the derivative is **negative**, increasing the strategy decreases the payoff.
- If the derivative is **zero**, the payoff is not changing — this could be a maximum, a minimum, or a flat point.

So: finding where the derivative equals zero helps us locate **potential optimal choices**.

Finding the "Top of the Hill"

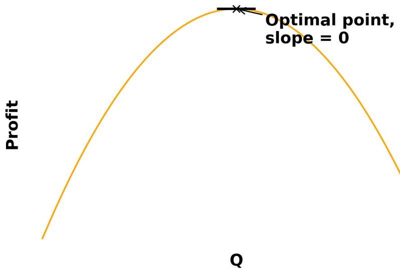

- Think of the curve as a smooth hill. The "slope" at any point tells you if you're still climbing (positive slope) or heading down (negative slope).
- The very top of the hill is where the slope is zero—there's no more "up" to climb.
- In our games, setting the slope to zero is just like finding that top: it gives the best possible choice (maximum payoff).

Taking First-Order Derivative

Suppose a function looks like

$$f(x) \;=\; a\,x^{\,b},$$

a and b are constants and x is the choice variable. $f'(x)$ means first-order derivative.

$$f'(x) = \frac{d}{dx}\left[a\,x^{\,b}\right] \;=\; abx^{\,b-1}.$$

<u>How to remember it</u>

1. Pull the exponent b down in front.
2. Keep the constant a along for the ride.
3. Subtract one from the old exponent to get the new one $(b-1)$.

Exercises

Example: If $f(x) = 5x^3 \implies f'(x) = 5 \times 3\,x^{2} = 15x^2$.

Take the first derivative with respect to x in the following functions

- $f(x) = 45x^{0.5}$
- $f(x, y) = xy$
- $f(x, y) = x^3y^2$

Example: A Simple Monopoly Problem

- Inverse demand: $P = 100 - Q$
- Constant marginal cost: $c = 20$

Writing the profit function, we get

$$\pi(Q) \;=\; (P - c)Q \;=\; \big(100 - Q - 20\big)\, Q \;=\; 80Q \;-\; Q^2.$$

We take the first-order derivative

$$\pi'(Q) \;=\; \frac{d}{dQ}\big(80Q - Q^2\big) \;=\; 80 \;-\; 2Q.$$

Then, setting the derivative to zero ("top of the hill")

$$80 - 2Q \;=\; 0 \quad \Longrightarrow \quad Q^* \;=\; 40.$$

$$P^* \;=\; 100 - 40 \;=\; 60.$$

So the monopolist produces 40 units and charges \$60. Total $\pi = (60 - 20) \times 40 = 1600$.

Quantity vs. Price Competition

- **Cournot competition (quantity game)**
  - Each firm chooses Q_i at the same time.
  - The market price adjusts to clear demand.
  - Good fit when plants must be scheduled in advance or output can't be changed quickly (steel, cement).
  - We will use first-order derivatives to get the solution.

- **Bertrand competition (price game)**
  - Each firm sets a *price* P_i simultaneously.
  - Consumers buy from the lowest-price seller (split if tied).
  - Good fit when firms can change prices faster than output (gas stations, airline tickets posted online).
  - We will get the solution via intuition.

Cournot

Market and cost structure

$$\text{Inverse demand:} \quad P = 100 - \big(Q_1 + Q_2\big), \qquad \text{(with } Q_i \text{ in units)}$$

$$\text{Constant marginal cost: } c = 20.$$

Each firm's profit

$$\pi_1 = (P - c)\, Q_1 = \big[\, 100 - (Q_1 + Q_2) - 20 \big]\, Q_1 = 80 Q_1 - Q_1^2 - Q_1 Q_2.$$

(Symmetric expression for π_2.)

Each firm chooses Q_i to *maximize* π_i taking the rival's quantity as given. We find the **best-response** functions and where they intersect (Cournot–Nash equilibrium).

Cournot

Cournot

Firm 1's first-order condition

$$\frac{\partial \pi_1}{\partial Q_1} \;=\; 80 - 2Q_1 - Q_2 \;=\; 0 \quad \Longrightarrow \quad \boxed{Q_1 = 40 - \tfrac{1}{2}Q_2}$$

Firm 2's first-order condition (symmetry)

$$\frac{\partial \pi_2}{\partial Q_2} \;=\; 80 - 2Q_2 - Q_1 \;=\; 0 \quad \Longrightarrow \quad \boxed{Q_2 = 40 - \tfrac{1}{2}Q_1}$$

These two linear equations are the firms' *best-response* curves. The Cournot equilibrium lies at their intersection.

Cournot

Substitute $Q_2 = 40 - \frac{1}{2}Q_1$ into Firm 1's best response:

$$Q_1 \;=\; 40 - \tfrac{1}{2}\Big(40 - \tfrac{1}{2}Q_1\Big) \;=\; 40 - 20 + \tfrac{1}{4}Q_1$$

$$\tfrac{3}{4}Q_1 = 20 \implies Q_1^* = \frac{80}{3} \approx 26.67.$$

By symmetry, $Q_2^* = 26.67$. Market price is

$$P^* \;=\; 100 - (Q_1^* + Q_2^*) \;=\; 100 - \frac{160}{3} \approx \$46.67.$$

Each firm's profit

$$\pi^* = (P^* - c)\,Q^* = (46.67 - 20) \times 26.67 \approx \$711.$$

Thus, in Cournot equilibrium each $Q_1^* = Q_2^* = 26.7$ units, $P = \$46.7$, and $\pi^* = \$711$ in profit.

Bertrand

- Two firms sell an identical product.
- Each **chooses a price** P_i simultaneously.
- Consumers buy entirely from the firm with the lower price (or split demand if prices tie).
- Marginal cost is constant and the same for both firms: c.

Intuition

- If my rival charges \$60, I charge \$59.99. I get *all* customers.
- The rival anticipates this and undercuts me by a penny.
- The undercutting race continues until price is pushed all the way down to cost c.

In this simplest Bertrand world, the outcome is brutal: **price = marginal cost** and profits are wiped out.

Bertrand

Let's go over a specific example.

- Demand if one firm is cheapest: $Q = 100 - P$.
- Marginal cost: $c = \$20$.
- Two firms, A and B, with identical costs.

Step-by-step logic

1. Suppose both post \$60. They split the market and each earns positive profit.
2. A spots an opportunity: cut to \$59.99, sell the whole market, and make more profit.
3. B reacts: cut to \$59.98 $\rightarrow$ capture all customers.
4. Repeat... prices cascade down until neither can profitably undercut.
5. The race stops only when $\mathbf{P} = \$\mathbf{20}$. Any lower price means producing at a loss.

Bertrand

$$\boxed{\text{Bertrand-Nash equilibrium: } P_A = P_B = \$20,\ \pi_A = \pi_B = 0.}$$

Take-away: With perfect substitutes and no capacity limits, price competition drives profit to zero—quite a contrast with Cournot!

Why undercutting is profitable?

Beyond the Basic Bertrand Story

The "price-equals-cost" outcome hinges on a few stark assumptions. Relax any of them and firms can sustain prices above cost.

- **Product differentiation** If products aren't perfect substitutes (think Pepsi vs. Coke), consumers tolerate small price gaps and undercutting is weaker.
- **Capacity constraints** A firm may not be able to serve the whole market even if it undercuts, softening the incentive to cut price.
- **Repeated interaction or price-matching guarantees** Future retaliation or posted promises can deter aggressive undercutting.

These extensions give richer, more realistic predictions—but the simple Bertrand "race to cost" remains a powerful benchmark for thinking about aggressive price competition.

Stackelberg

Our math tools worked for *simultaneous* decisions (like Cournot). We can also apply them to *sequential* decisions, where one firm moves first.

- One firm (the *leader*) chooses a quantity first.
- The other firm (the *follower*) observes that decision and responds with its own quantity.

We solve it using the rollback method.

When does Stackelberg fit well?

- **Markets with dominant players** — one firm or group has enough power to influence the behavior of others.
- **Example:** OPEC acts as a leader in global oil markets, setting production levels that smaller oil-producing countries respond to.

Stackelberg

- Inverse demand: $P = 100 - (Q_1 + Q_2)$
- Marginal cost for both firms: $c = 20$

We start solving the followers problem.

$$\pi_2 = (P - c)Q_2 \quad \Rightarrow \quad \frac{\partial \pi_2}{\partial Q_2} = 0 \implies Q_2 = 40 - \tfrac{1}{2}Q_1$$

Then, we insert $Q_2(Q_1)$ into Firm 1's profit:

$$\begin{aligned}\pi_1(Q_1) &= \left(100 - Q_1 - \left(40 - \tfrac{1}{2}Q_1\right) - 20\right) Q_1 \\ &= \left(40 - \tfrac{1}{2}Q_1\right) Q_1 = 40Q_1 - \tfrac{1}{2}Q_1^2.\end{aligned}$$

Taking first-order derivatives, $\pi_1'(Q_1) = 40 - Q_1 = 0$; and hence $Q_1^* = 40$. The Follower reacts with $Q_2^* = 40 - \frac{1}{2}(40) = 20$.

Stackelberg

Results

Total output: 60 (Cournot was $\frac{160}{3} \approx 53$).

$P^* = 100 - 60 = 40$ (Cournot price ≈ 46.7).

$\pi_1^* = (40 - 20) \times 40 = 800 \quad > \ \pi_2^* = (40 - 20) \times 20 = 400.$

Take-away: The leader produces more and earns higher profit than in Cournot, while the follower earns less—first-mover advantage in action.

Different Market Equilibrium

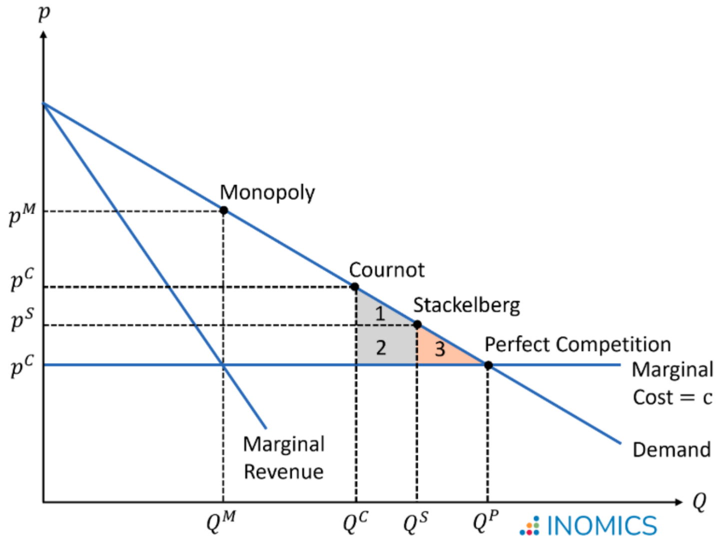

Discussion: Cartels and Collusion

A **cartel** is a formal or informal agreement between firms to avoid competition — usually by coordinating prices or quantities. The goal is simple: behave like a monopolist and increase joint profits.

Why is this important?

- In a cartel, firms agree to restrict output or fix high prices.
- This leads to **higher profits** for firms, but **worse outcomes** for consumers (higher prices, less choice).
- Cartels are **illegal** in most countries because they reduce market competition and harm social welfare.

Cartels help firms, hurt consumers, and are against competition laws in most of the world.