Game Theory (ECO 311)

Lecture Notes

Feda Mohammadi

Professor: Dr. Jarrod Burgh

Berea College

Spring 2026

These notes are for personal academic study. Examples and definitions follow standard game theory notation.

Contents

1 Introduction to Strategic Interaction 1
- 1.1 What is Game Theory? 1
  - 1.1.1 Strategic interdependence 1
  - 1.1.2 Conflict and cooperation 1
  - 1.1.3 A very short history (why the field exists) 2
  - 1.1.4 Where game theory is used 2
  - 1.1.5 How we represent a game: extensive form vs normal form 2
- 1.2 Mathematical Review 3
  - 1.2.1 Functions 3
  - 1.2.2 Maximization 4
  - 1.2.3 Sets and Strategies 4
  - 1.2.4 Basic Probability 5
- 1.3 Some Basic Mathematics That Will Come in Handy Later 6
  - 1.3.1 Basic Derivative Rules 6
  - 1.3.2 Derivative Rules for Combined Functions 6
  - 1.3.3 Concavity and Second Derivatives 7
  - 1.3.4 Local and Global Maxima 7
  - 1.3.5 Partial Derivatives 8
- 1.4 Choices 8
  - 1.4.1 Basic Decision Theory 8
  - 1.4.2 The Three Elements of a Decision Problem 9
  - 1.4.3 Rational Preferences 9
  - 1.4.4 Utility Representation 10
- 1.5 The Rational Choice Paradigm 11
- 1.6 Uncertainty, Probability, and Expectations 13
  - 1.6.1 Probability Distributions 14
  - 1.6.2 State-Dependent Utility 14
  - 1.6.3 Expected Value 14
  - 1.6.4 Risk Attitudes 16
  - 1.6.5 Risk and the Shape of Utility 16

2 Normal Form Games 17
- 2.1 Strategic Games 17
  - 2.1.1 The Formal Definition 17
  - 2.1.2 Prisoner's Dilemma 17

2.1.3 Bach or Stravinsky (Marriage Game) 18
2.1.4 Matching Pennies 19
2.2 Iterated Elimination of Strictly Dominated Strategies 19
2.3 Introducing Nash Equilibrium 21
2.4 Solving Games 22
2.5 Best Response and Nash Equilibria 24
2.5.1 Best Response Functions 24
2.5.2 Nash Equilibrium 24
2.5.3 Strict Nash Equilibrium 25
2.5.4 Pareto Dominance and Pareto Optimality 25
2.5.5 Dominated Strategies 26
2.5.6 Examples 27
2.5.7 Symmetric Games 29

3 Applications of Nash Equilibrium 30
3.1 Cournot's Oligopoly 30
3.1.1 Model Setup 30
3.1.2 Profit Maximization Problem 30
3.1.3 Best Response Function 31
3.1.4 Nash Equilibrium 31
3.1.5 Graphical Representation 32
3.2 Examples 33

4 Coming soon 39

Chapter 1: Introduction to Strategic Interaction

1.1 What is Game Theory?

Big idea

Game theory studies **interactive decisions**: situations where multiple players choose actions, and the outcome for each player depends on what *everyone* does.

Game theory builds models of situations where several entities (players) make choices, and the combination of their choices produces an outcome that affects all of them. Because different players can value the same outcome differently, game theory is also the study of **rational behavior** in these interactive settings. A central feature is **strategic interdependence**: my best action depends on what I think others will do.

1.1.1 Strategic interdependence

In many problems, you cannot choose well without thinking about others. This is what makes game theory different from standard single-person decision problems. An interaction has strategic interdependence **if each player's outcome depends on the actions of other players.**

Example 1: Competing coffee shops

Two coffee shops choose prices. If one shop lowers its price, it can attract customers away from the other shop. Each shop's profit depends on *both* prices, not just its own price. This is a game because each shop's best decision depends on what the other shop does.

1.1.2 Conflict and cooperation

Many strategic situations lie somewhere between pure conflict and pure cooperation. Some games look like competition (each player wants the other to do poorly), and some games allow mutually beneficial outcomes if players coordinate or cooperate.

Example 2: Roommates choosing quiet hours

Two roommates each prefer quiet time to study, but they also want flexibility. If both agree on quiet hours, both benefit. But if one ignores the agreement, they may gain personally while harming the other. This is a strategic situation mixing cooperation (agreement helps) and conflict (temptation to deviate).

1.1.3 A very short history (why the field exists)

Game theory has roots in early probability problems (for example, analyzing the probability of winning card games as play unfolds). Later work studied finite games with perfect information and strategic reasoning. Key foundational contributions include work by von Neumann (including the minimax theorem) and the book *Theory of Games and Economic Behavior* by von Neumann and Morgenstern (1944), which gave early formal descriptions of major classes of games and core principles of the theory. Nash later introduced the concept of strategic equilibrium and proved existence results that became central to modern game theory.

1.1.4 Where game theory is used

Game theory is heavily used in economics, but it is also widely applied in biology, computer science, political science, operations research, and other areas of mathematics.

Course focus

Games are often classified as **strategic (non-cooperative)** versus **coalitional (cooperative)**. But here, we mainly study **strategic (non-cooperative)** game theory.

1.1.5 How we represent a game: extensive form vs normal form

There are two common ways to describe the rules and structure of strategic interaction.

Extensive form (game as rules)

The extensive form describes the **rules of the game** explicitly: who moves when, what actions are available at each point, what information players have during play, and how play leads to an outcome.

A key concept in extensive-form games is a **strategy**. A strategy is a complete plan that specifies what a player would do at every situation where they might have to act.

Example 3: Strategy as a complete plan

In chess, a strategy is not just "move my queen now." It is a full plan describing what you would do for every possible configuration you might face later. That is why a strategy is a complete mapping from situations to actions.

Normal form (game as a mathematical object)

In normal form, we list each player's strategy set and specify a mapping from a profile of strategies to outcomes and payoffs. In this representation, an n-player game can be described by a map from a product of n spaces into $\mathbb{R}^n$ (one payoff for each player).

The extensive form is often simpler to work with and can lead to computations with lower complexity. The normal form is a universal and compact mathematical formulation once strategies and payoffs are specified.

1.2 Mathematical Review

Game theory relies heavily on mathematical structure. Before studying strategic interaction, we need to review the mathematical tools that will appear repeatedly throughout this course.

1.2.1 Functions

A function assigns a number to each element in a set. In economics and game theory, functions usually represent preferences, payoffs, or beliefs.

If

$$f : X \to \mathbb{R},$$

then for each $x \in X$, the function assigns a real number $f(x)$.

Example: Payoff Function

Suppose a firm chooses quantity q. Its profit is

$$\pi(q) = 10q - q^2.$$

Here, q is the action and $\pi(q)$ is the payoff.
The function tells us how profitable each possible action is.

1.2.2 Maximization

A central idea in game theory is that players maximize something, usually utility or payoff.

Let $f(x)$ be a payoff function. A value x^* solves

$$x^* \in \arg\max_{x \in X} f(x)$$

if

$$f(x^*) \geq f(x) \quad \text{for all } x \in X.$$

This means x^* gives the highest attainable payoff.

Example: Simple Maximization

Let

$$f(x) = -x^2 + 4x.$$

Take the derivative:

$$f'(x) = -2x + 4.$$

Set it equal to zero:

$$-2x + 4 = 0 \quad \Rightarrow \quad x = 2.$$

So, $x = 2$ maximizes the function.

1.2.3 Sets and Strategies

Game theory describes choices using sets. If player i has a strategy set S_i, then:

$$S_i = \{\text{all strategies available to player } i\}.$$

If there are n players, the set of strategy profiles is

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

An element of S is written

$$s = (s_1, s_2, \ldots, s_n).$$

Example: Two-Player Strategy Set

Suppose two firms choose either High price (H) or Low price (L).
Then

$$S_1 = \{H, L\}, \quad S_2 = \{H, L\}.$$

The strategy profiles are:

$$(H, H), (H, L), (L, H), (L, L).$$

Each profile represents one possible outcome.

1.2.4 Basic Probability

Many games involve uncertainty. A probability distribution assigns likelihoods to possible events.

If $E_1, E_2, \ldots, E_k$ are mutually exclusive events, then

$$\sum_{j=1}^{k} P(E_j) = 1.$$

Expected value plays a crucial role in decision making under uncertainty. If a random variable X takes values x_j with probabilities p_j, then

$$E[X] = \sum_j p_j x_j.$$

Example: Expected Payoff

Suppose a player receives

$$100 \text{ with probability } 0.5 \quad \text{and} \quad 0 \text{ with probability } 0.5.$$

Then the expected payoff is

$$E[X] = 0.5(100) + 0.5(0) = 50.$$

A rational decision maker compares expected payoffs across actions.

1.3 Some Basic Mathematics That Will Come in Handy Later

In solving maximization problems and analyzing curvature properties of payoff functions, the following results will be used frequently.

1.3.1 Basic Derivative Rules

Let $f : C \to \mathbb{R}$, with $C \subseteq \mathbb{R}$.

The derivative of f at x is defined as

$$f'(x) = \lim_{h \to 0} \frac{f(x+h) - f(x)}{h}.$$

We also use the notation

$$f'(x) = \frac{df(x)}{dx}.$$

$->$ The derivative is a local concept. It measures the rate of change of a function at a point.

Useful Derivatives

$$\frac{d}{dx}(c) = 0$$

$$\frac{d}{dx}(a + bx) = b$$

$$\frac{d}{dx}(ax^n) = anx^{n-1}$$

$$\frac{d}{dx}(e^{ax}) = ae^{ax}$$

$$\frac{d}{dx}(\ln x) = \frac{1}{x}$$

1.3.2 Derivative Rules for Combined Functions

Let f and g be differentiable functions.

$$\frac{d}{dx}(f(x) + g(x)) = f'(x) + g'(x)$$

$$\frac{d}{dx}(f(x) - g(x)) = f'(x) - g'(x)$$

$$\frac{d}{dx}(f(x)g(x)) = f'(x)g(x) + g'(x)f(x)$$

If $g(x) \neq 0$,

$$\frac{d}{dx}\left(\frac{f(x)}{g(x)}\right) = \frac{f'(x)g(x) - g'(x)f(x)}{g(x)^2}$$

Chain rule:

$$\frac{d}{dx}f(g(x)) = f'(g(x))g'(x)$$

1.3.3 Concavity and Second Derivatives

Concavity plays a central role in optimization problems.

A function f is concave if

$$f(\lambda x_1 + (1-\lambda)x_2) \geq \lambda f(x_1) + (1-\lambda)f(x_2)$$

for all $x_1, x_2 \in C$ and $\lambda \in [0, 1]$.

If f is twice differentiable, then

$$f''(x) \leq 0 \quad \text{implies concavity.}$$

If $f''(x) < 0$, then f is strictly concave.

1.3.4 Local and Global Maxima

A point x^* is a local maximum if

$$f(x^*) \geq f(x) \quad \text{for all } x \text{ sufficiently close to } x^*.$$

It is a global maximum if

$$f(x^*) \geq f(x) \quad \text{for all } x \in C.$$

Strict concavity guarantees that a local maximum is automatically global.

1.3.5 Partial Derivatives

In many games, payoffs depend on several variables. For example,

$$\pi_i(q_i, q_j).$$

The partial derivative with respect to q_i is

$$\frac{\partial \pi_i}{\partial q_i}.$$

It measures how player i's payoff changes when q_i changes, holding other variables fixed.

Example: Two-Variable Function

Let

$$f(x, y) = 4x - x^2 - xy.$$

Then

$$\frac{\partial f}{\partial x} = 4 - 2x - y,$$

$$\frac{\partial f}{\partial y} = -x.$$

To find a critical point, we set both partial derivatives equal to zero.

Second partial derivatives help determine curvature in multivariable settings, which will be important when studying best responses and equilibrium.

1.4 Choices

Game theory begins with a simple idea: a decision maker must choose among alternatives, and each choice leads to consequences.

1.4.1 Basic Decision Theory

A decision maker (also called an agent or a player) faces a situation in which she must choose one of several alternatives.

Each choice leads to an outcome, and the consequences of that outcome are borne by the decision maker herself.

In order to analyze a decision problem, we must specify these three elements:

1. What actions are available?
2. What outcomes result from each action?
3. How does the decision maker rank those outcomes?

1.4.2 The Three Elements of a Decision Problem

A decision problem consists of three basic components:

- **Actions.** Actions are the set of alternatives from which the decision maker can choose. We denote the set of possible actions by A.
- **Outcomes.** Outcomes are the possible consequences that result from actions. We denote the set of outcomes by X.
- **Preferences.** Preferences describe how the decision maker ranks the set of possible outcomes, from most desired to least desired. Preferences are typically written using a binary relation $\succeq$.

The preference relation $x \succeq y$ means that outcome x is weakly preferred to outcome y.

We also define:

$$x \succ y \quad \text{if } x \succeq y \text{ and not } y \succeq x,$$

$$x \sim y \quad \text{if } x \succeq y \text{ and } y \succeq x.$$

1.4.3 Rational Preferences

To model rational decision making, we impose two fundamental properties on preferences.

- **Completeness.** For any two outcomes $x, y \in X$, the decision maker can compare them. That is,

  $$x \succeq y, \quad y \succeq x, \quad \text{or both.}$$

  Completeness means the decision maker is never unable to rank two alternatives.
- **Transitivity.** For any $x, y, z \in X$,

  $$x \succeq y \quad \text{and} \quad y \succeq z \quad \Rightarrow \quad x \succeq z.$$

  Transitivity guarantees internal consistency in ranking.

Preferences satisfying completeness and transitivity are called **rational preferences**.

1.4.4 Utility Representation

If preferences are rational, we can represent them numerically.

A function

$$u : X \to \mathbb{R}$$

represents preferences if for any $x, y \in X$,

$$u(x) \geq u(y) \quad \text{if and only if} \quad x \succeq y.$$

The function $u(x)$ is called a payoff function or utility function. Utility does not measure happiness in a psychological sense. It is simply a numerical representation of preferences.

Only the ranking matters. If u represents preferences, then any strictly increasing transformation of u represents the same preferences.

Example

Suppose a decision maker ranks outcomes as

$$x_1 \succ x_2 \succ x_3.$$

One possible utility representation is

$$u(x_1) = 3, \quad u(x_2) = 2, \quad u(x_3) = 1.$$

Another valid representation is

$$v(x_1) = 100, \quad v(x_2) = 50, \quad v(x_3) = 10.$$

Both represent the same preference ordering.

1.5 The Rational Choice Paradigm

The rational choice paradigm assumes that a decision maker chooses an action that maximizes her payoff (utility), given what she knows about the available actions and the consequences of those actions.

A decision problem is fully described once the decision maker knows:

1. all possible actions, A;
2. all possible outcomes, X;
3. how actions map into outcomes (what outcome happens under each action);
4. her rational preferences (payoffs) over outcomes.

Definition (Rational Choice):

Let A be the set of feasible actions and let $u(\cdot)$ be a payoff (utility) function defined over actions.

A decision maker is **rational** if she chooses an action $a^* \in A$ that maximizes her payoff:

$$a^* \in \arg\max_{a \in A} u(a).$$

Equivalently, $a^* \in A$ is chosen if and only if

$$u(a^*) \geq u(a) \quad \text{for all } a \in A.$$

Example 1

Action space: $A = [0, 1]$
Utility function:

$$u(a) = 2a - 4a^2.$$

Question: What is the rational choice?

We maximize $u(a)$ over $a \in [0, 1]$. First, we compute the derivative:

$$u'(a) = 2 - 8a.$$

Then set $u'(a) = 0$:

$$2 - 8a = 0 \quad \Rightarrow \quad a = \frac{1}{4}.$$

Check curvature:

$$u''(a) = -8 < 0,$$

so u is strictly concave and this critical point is the unique maximizer (as long as it is feasible). Since $\frac{1}{4} \in [0, 1]$, it is feasible. For completeness, check endpoints:

$$u(0) = 0, \qquad u(1) = 2 - 4 = -2.$$

Therefore, the rational choice is

$$\boxed{a^* = \frac{1}{4}.}$$

Example 2

Action space: $p \in \mathbb{R}_{\geq 0}$
Market demand:

$$q = 40 - 2p.$$

Question: What is the monopolist's rational choice (profit-maximizing price)?

With no cost information given, profit equals revenue:

$$\pi(p) = pq = p(40 - 2p) = 40p - 2p^2.$$

Maximize $\pi(p)$ over $p \geq 0$.
Derivative:

$$\pi'(p) = 40 - 4p.$$

Set $\pi'(p) = 0$:

$$40 - 4p = 0 \quad \Rightarrow \quad p = 10.$$

Second derivative:

$$\pi''(p) = -4 < 0,$$

so this is a maximum. So the profit-maximizing (rational) price is

$$\boxed{p^* = 10.}$$

The corresponding quantity is

$$q^* = 40 - 2(10) = 20.$$

Example 3

Action space: $A = \mathbb{R}_{\geq 0}$
Utility function:

$$u(a) = \ln(2a) - \frac{1}{4}a.$$

Question: What is the rational choice?

Because of the logarithm, we must have $a > 0$.
Differentiate:

$$u'(a) = \frac{d}{da}\ln(2a) - \frac{1}{4} = \frac{1}{a} - \frac{1}{4}.$$

Set $u'(a) = 0$:

$$\frac{1}{a} - \frac{1}{4} = 0 \quad \Rightarrow \quad \frac{1}{a} = \frac{1}{4} \quad \Rightarrow \quad a = 4.$$

Second derivative:

$$u''(a) = -\frac{1}{a^2} < 0 \quad \text{for } a > 0,$$

so u is strictly concave and $a = 4$ is the unique maximizer.
Therefore, the rational choice is

$$\boxed{a^* = 4.}$$

1.6 Uncertainty, Probability, and Expectations

So far, we analyzed decision problems without uncertainty. However, in many situations, the decision maker does not know with certainty which outcome will occur.

Uncertainty plays a central role in:

- individual decision problems (uncertainty about the state of the world),
- strategic games (uncertainty about opponents or about nature).

Here, we focus on stochastic (probabilistic) uncertainty and primarily on risk neutrality.

1.6.1 Probability Distributions

Let Ω denote a finite set of states of the world.

A probability distribution over Ω is a function

$$p \in \Delta(\Omega)$$

such that

$$p(\omega) \geq 0 \quad \forall \omega \in \Omega,$$

$$\sum_{\omega \in \Omega} p(\omega) = 1.$$

Thus, probabilities are non-negative and sum to one.

1.6.2 State-Dependent Utility

Under uncertainty, payoffs depend on both the chosen action and the realized state.

We define a state-dependent payoff function:

$$u : A \times \Omega \to \mathbb{R}.$$

For each action a and state ω, the realized payoff is $u(a, \omega)$.

1.6.3 Expected Value

Given a probability distribution p, the expected value of action a is

$$\mathbb{E}_p(a) = \sum_{\omega \in \Omega} p(\omega) u(a, \omega).$$

The expectation is the probability-weighted average payoff.

A risk-neutral decision maker chooses

$$a^* \in \arg\max_{a \in A} \mathbb{E}_p(a).$$

Risk neutrality therefore means maximizing expected value.

Example: Coin Flip

Suppose a decision maker chooses action a and receives

$$u(a, H) = 10, \qquad u(a, T) = 4.$$

If the coin is fair,

$$p(H) = \frac{1}{2}, \qquad p(T) = \frac{1}{2}.$$

The expected payoff is

$$\mathbb{E}_p(a) = \frac{1}{2}(10) + \frac{1}{2}(4) = 7.$$

If heads is three times as likely as tails, then

$$p(H) = \frac{3}{4}, \qquad p(T) = \frac{1}{4},$$

and

$$\mathbb{E}_p(a) = \frac{3}{4}(10) + \frac{1}{4}(4) = 8.5.$$

Example: Fair Die

Suppose a decision maker rolls a fair die and receives payoff equal to the number rolled:

$$u(a, \omega) = \omega, \qquad \omega \in \{1, \ldots, 6\}.$$

Since the die is fair,

$$p(\omega) = \frac{1}{6}.$$

The expected payoff is

$$\mathbb{E}_p(a) = \frac{1}{6}(1 + 2 + 3 + 4 + 5 + 6) = \frac{21}{6} = 3.5.$$

1.6.4 Risk Attitudes

Risk attitudes describe how a decision maker evaluates uncertain payoffs.

- **Risk-neutral:** maximizes expected value.
- **Risk-averse:** prefers smoother, less extreme payoffs.
- **Risk-loving:** prefers more extreme payoffs.

1.6.5 Risk and the Shape of Utility

A decision maker's attitude toward risk is determined by the curvature of the utility function.

Theorem 1.1. *A decision maker is risk-averse if and only if the utility function u is concave.*

Theorem 1.2. *A decision maker is risk-loving if and only if the utility function u is convex.*

If $u''(x) < 0$, utility is concave and the decision maker dislikes risk. If $u''(x) > 0$, utility is convex and the decision maker likes risk.

Insurance Example

Suppose there is a 20% probability of incurring a \$100 loss.
The expected loss is

$$0.2 \times 100 = 20.$$

A risk-neutral decision maker would pay at most \$20 for full insurance.
A risk-averse decision maker would pay strictly more than \$20.

Lottery Example

Suppose a lottery ticket costs \$1 and pays \$1001 with probability p.
Expected value:

$$\mathbb{E} = 1001p - 1.$$

A risk-neutral agent buys the ticket if

$$1001p - 1 \geq 0 \quad \Rightarrow \quad p \geq \frac{1}{1001}.$$

A risk-loving agent may buy even if p is smaller.

Chapter 2: Normal Form Games

2.1 Strategic Games

A strategic game is a model of interaction between rational decision makers who choose their actions simultaneously. Each player has a set of possible actions, and the outcome of the game depends on the combination of all players' choices. The model specifies payoffs for every possible action profile, capturing each player's preferences over outcomes.

2.1.1 The Formal Definition

A **strategic game** consists of:

- A set of players.
- For each player, a set of actions.
- For each player, payoffs over the set of action profiles.

This definition applies to **static games**: players choose their actions **simultaneously**. We focus on **normal form games**.

2.1.2 Prisoner's Dilemma

- Two suspects in a major crime are held in separate cells.
- There is enough evidence to convict each of a minor offense.
- There is not enough evidence to convict either of the major crime unless one acts as an informer against the other (**Confess**).

Possible outcomes:

- If they both **Don't Confess**, each is convicted of a minor offense and spends 1 year in prison.
- If exactly one **Confesses**, the confessor is freed and used as a witness, while the other spends 3 years in prison.
- If they both **Confess**, each spends 2 years in prison.

Players, Actions, and Preferences

- Players: the two suspects.
- Actions: {Don't Confess, Confess}.

Define Don't Confess $= D$ and Confess $= C$. Then suspect 1's payoffs can be ordered as

$$u_1(C,D) > u_1(D,D) > u_1(C,C) > u_1(D,C),$$

and suspect 2's payoffs can be ordered as

$$u_2(D,C) > u_2(D,D) > u_2(C,C) > u_2(C,D).$$

Figure: Payoff Matrix for the prisoner's dilemma.

		Player 2	
		C	D
Player 1	C	2, 2	0, 3
	D	3, 0	1, 1

2.1.3 Bach or Stravinsky (Marriage Game)

Two people wish to go out together. Two concerts are available: one of music by Bach (B), and one of music by Stravinsky (S). One person prefers Bach and the other prefers Stravinsky. If they go to different concerts, each of them is equally unhappy listening to the music of either composer.

Figure: Payoff Matrix for Bach or Stravinsky.

		Player 2	
		B	S
Player 1	B	2, 1	0, 0
	S	0, 0	1, 2

2.1.4 Matching Pennies

Two people choose simultaneously whether to show the Head (H) or the Tail (T) of a coin. If they show the same side, person 2 pays person 1 a dollar; if they show different sides, person 1 pays person 2 a dollar. Each person cares only about the amount of money she receives, and prefers to receive more than less.

Figure: Payoff Matrix for the matching pennies game.

		Player 2	
		H	T
Player 1	H	1, -1	-1, 1
	T	-1, 1	1, -1

2.2 Iterated Elimination of Strictly Dominated Strategies

Having described how to represent a strategic game, we now introduce a first method for solving a normal form game.

The guiding idea is simple: a rational player will never play a strictly dominated strategy.

Strict Dominance

Definition. Consider a normal form game with two players. Let $a_i', a_i'' \in A_i$.
The action (strategy) a_i' is *strictly dominated* by strategy a_i'' if for every action $a_j \in A_j$ of player j, player i's payoff from playing a_i' is strictly less than the payoff from playing a_i'':

$$u_i(a_i', a_j) < u_i(a_i'', a_j) \quad \forall a_j \in A_j.$$

In words, strategy a_i'' yields a strictly higher payoff than a_i' no matter what the other player does. A rational player therefore has no reason to choose a_i'.

Application: Prisoner's Dilemma

Recall the Prisoner's Dilemma with actions C (Confess) and D (Don't Confess).

For prisoner i:

- If the other prisoner plays D, then i prefers C to avoid prison.
- If the other prisoner plays C, then i prefers C to receive a shorter sentence.

Thus, for each strategy of player j, the payoff to player i from playing D is strictly less than the payoff from playing C.

Therefore, for each prisoner, strategy D is strictly dominated by strategy C.

Iterated Elimination

The procedure is:

1. Remove any strictly dominated strategies.
2. Reduce the game accordingly.
3. Check again whether new strategies become strictly dominated.
4. Repeat until no strictly dominated strategies remain.

This process is called *Iterated Elimination of Strictly Dominated Strategies*.

Limitations

Although the method is intuitive, it has important limitations.

- Each step requires an assumption about what players know about each other's rationality.
- To apply the process repeatedly, we must assume *common knowledge* of rationality.

That is, we assume:

- All players are rational.
- All players know that all players are rational.
- All players know that all players know that all players are rational.
- And so on, *ad infinitum*.

A further drawback is that iterated elimination often produces imprecise predictions about play. Moreover, it becomes increasingly complex in games with more than two players.

2.3 Introducing Nash Equilibrium

What actions will be chosen by the players in a strategic game?

To answer this question, we need an equilibrium concept. In this course we use the notion of **Nash equilibrium** (and its extensions).

Intuition

Intuitively, a Nash equilibrium has two components:

1. Each player chooses her action according to rational choice, given her belief about the other players' actions.
2. Each player's belief about the other players' actions is correct.

Thus, in equilibrium, every player is best responding to the actions actually chosen by the others.

Informal Description

A Nash equilibrium is an action profile a^* with the property that no player can do better by unilaterally deviating, given that all other players adhere to a^*.

Formal Definition

Definition (Nash Equilibrium):
Let $a^* = (a_1^*, \ldots, a_n^*)$ be an action profile in a strategic game.
The action profile a^* is a Nash equilibrium if, for every player i and every action $a_i \in A_i$,

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

That is, given that every other player $j \neq i$ chooses a_j^*, player i cannot obtain a higher payoff by choosing any alternative action a_i.
Here u_i is the payoff function representing player i's preferences.

2.4 Solving Games

We now apply best response functions to solve games and identify Nash equilibria.

Bach or Stravinsky (Marriage Game)

Recall the payoff matrix:

Figure: Payoff Matrix for Bach or Stravinsky.

Player 1 \ Player 2	B	S
B	2, 1	0, 0
S	0, 0	1, 2

Best Responses: Player 1

- If Player 2 plays B, Player 1 prefers B since $2 > 0$.
- If Player 2 plays S, Player 1 prefers S since $1 > 0$.

Thus,

$$B_1(B) = \{B\}, \qquad B_1(S) = \{S\}.$$

Best Responses:Player 2

- If Player 1 plays B, Player 2 prefers B since $1 > 0$.
- If Player 1 plays S, Player 2 prefers S since $2 > 0$.

So,

$$B_2(B) = \{B\}, \qquad B_2(S) = \{S\}.$$

Nash Equilibria

The Nash equilibria are the action profiles where both players are playing best responses simultaneously.

Hence, there are two pure-strategy Nash equilibria:

$$(B, B) \quad \text{and} \quad (S, S).$$

This game has multiple equilibria.

Matching Pennies

Recall the payoff matrix:

Figure: Payoff Matrix for Matching Pennies.

		Player 2	
		H	T
Player 1	H	-1 1	1 -1
	T	1 -1	-1 1

Best Responses: Player 1

- If Player 2 plays H, Player 1 prefers H.
- If Player 2 plays T, Player 1 prefers T.

Best Responses: Player 2

- If Player 1 plays H, Player 2 prefers T.
- If Player 1 plays T, Player 2 prefers H.

Nash Equilibrium

There is no action profile in which both players are simultaneously best responding.
Therefore, **Matching Pennies has no pure-strategy Nash equilibrium.**

2.5 Best Response and Nash Equilibria

The concept of best responses allows us to systematically identify Nash equilibria in strategic games.

2.5.1 Best Response Functions

Definition (Best Response Function).
In an n-player strategic game, the best response function of player i is

$$B_i : \prod_{j \neq i} A_j \to A_i,$$

where for every action profile of the other players a_{-i},

$$B_i(a_{-i}) = \left\{ a_i \in A_i \;:\; u_i(a_i, a_{-i}) \geq u_i(a_i', a_{-i}) \text{ for all } a_i' \in A_i \right\}.$$

What it means.

The best response function tells us what player i should choose given what everyone else is doing. It lists the actions that maximize player i's payoff, holding the other players' actions fixed.

In a two-player game this simplifies to

$$B_1 : A_2 \to A_1 \qquad \text{and} \qquad B_2 : A_1 \to A_2.$$

Each player's best response depends only on the other player's action.

2.5.2 Nash Equilibrium

Definition (Nash Equilibrium).
An action profile $a^* = (a_1^*, \ldots, a_n^*)$ is a Nash equilibrium if

$$a_i^* \in B_i(a_{-i}^*) \quad \text{for every player } i.$$

What it means.

A Nash equilibrium is a situation where every player is best responding to the others. No one has an incentive to deviate unilaterally. If any single player changes their action while others keep theirs fixed, their payoff does not increase.

2.5.3 Strict Nash Equilibrium

Definition (Strict Nash Equilibrium).
An action profile a^* is a strict Nash equilibrium if for every player i,

$$u_i(a^*) > u_i(a_i, a^*_{-i}) \quad \text{for every } a_i \neq a^*_i.$$

What it means.

In a strict Nash equilibrium, each player's equilibrium action is strictly better than any other action. No ties are allowed. This makes the equilibrium more robust.

2.5.4 Pareto Dominance and Pareto Optimality

Definition (Pareto Dominance).
An action profile a^* Pareto dominates another profile a' if

$$u_i(a^*) \geq u_i(a') \quad \text{for all players } i,$$

and

$$u_j(a^*) > u_j(a') \quad \text{for at least one player } j.$$

What it means.

One outcome Pareto dominates another if everyone is at least as well off and someone is strictly better off. It represents an improvement that harms no one.

Definition (Pareto Optimality).
An action profile is Pareto optimal if it is not Pareto dominated by any other action profile.

What it means.

A Pareto optimal outcome is efficient. There is no way to make someone better off without making someone else worse off.

2.5.5 Dominated Strategies

Definition (Strict Dominance).
An action a_i'' strictly dominates a_i' if

$$u_i(a_i'', a_{-i}) > u_i(a_i', a_{-i}) \quad \text{for every } a_{-i}.$$

What it means.

A strictly dominated strategy is never optimal. It gives a lower payoff no matter what others do, so a rational player should never choose it.

Definition (Weak Dominance).
An action a_i'' weakly dominates a_i' if

$$u_i(a_i'', a_{-i}) \geq u_i(a_i', a_{-i}) \quad \text{for every } a_{-i},$$

and

$$u_i(a_i'', a_{-i}) > u_i(a_i', a_{-i}) \quad \text{for some } a_{-i}.$$

What it means.

A weakly dominated strategy is never better and sometimes strictly worse. It might tie in some cases, but there exists at least one situation where it performs strictly worse.

2.5.6 Examples

Example 1 (Two-player game).

Player 1 \ Player 2	X	Y	Z
A	$(4, 0)$	$(2, 5)$	$(1, 3)$
B	$(7, 6)$	$(0, 4)$	$(2, 2)$
C	$(3, 1)$	$(5, 2)$	$(1, 4)$

1) Best response functions.
For Player 1 (rows), compare Player 1's payoff across rows in each column:

$$B_1(X) = \{B\}, \qquad B_1(Y) = \{C\}, \qquad B_1(Z) = \{B\}.$$

For Player 2 (columns), compare Player 2's payoff across columns in each row:

$$B_2(A) = \{Y\}, \qquad B_2(B) = \{X\}, \qquad B_2(C) = \{Z\}.$$

2) Nash equilibrium.
A pure-strategy Nash equilibrium is a cell where each action is a best response to the other. The only mutual best response is:

$$(B, X),$$

with payoff $(7, 6)$.
3) Pareto optimality.
(B, X) is Pareto optimal. Player 2's payoff is 6, and there is no other outcome that gives Player 2 at least 6 while also making Player 1 at least 7 (or even keeping Player 1 at least as well off). So no other action profile Pareto dominates (B, X).

Example 2 (Three-player game).

Player 1 chooses U or D. Player 2 chooses L or R. Player 3 chooses A or B.

The payoff matrices below show (u_1, u_2, u_3). The left matrix is when Player 3 plays A, and the right matrix is when Player 3 plays B.

A

Player 1 \ Player 2	L	R
U	$(3,2,3)$	$(1,2,0)$
D	$(1,2,0)$	$(1,2,0)$

B

Player 1 \ Player 2	L	R
U	$(1,4,2)$	$(1,2,0)$
D	$(1,2,0)$	$(3,4,1)$

1) Best response functions.

Player 1: best response to (a_2, a_3).

$$B_1(L,A) = \{U\}, \quad B_1(R,A) = \{U,D\}, \quad B_1(L,B) = \{U,D\}, \quad B_1(R,B) = \{D\}.$$

Player 2: best response to (a_1, a_3).

$$B_2(U,A) = \{L,R\}, \quad B_2(D,A) = \{L,R\}, \quad B_2(U,B) = \{L\}, \quad B_2(D,B) = \{R\}.$$

Player 3: best response to (a_1, a_2).

$$B_3(U,L) = \{A\}, \quad B_3(U,R) = \{A,B\}, \quad B_3(D,L) = \{A,B\}, \quad B_3(D,R) = \{B\}.$$

2) Nash equilibria: A profile (a_1, a_2, a_3) is a Nash equilibrium if each $a_i \in B_i(a_{-i})$. The pure-strategy Nash equilibria are:

$$(U,L,A), \qquad (U,R,A), \qquad (D,R,B).$$

Their payoffs are: $(U,L,A) : (3,2,3)$, $(U,R,A) : (1,2,0)$, $(D,R,B) : (3,4,1)$.

3) Pareto optimality: (U,R,A) is *not* Pareto optimal because (U,L,A) yields $(3,2,3) \geq (1,2,0)$ componentwise, with strict improvements for Players 1 and 3.

(U,L,A) *is* Pareto optimal (no other profile makes all three weakly better and at least one strictly better).

(D,R,B) *is* Pareto optimal (there is no other profile that weakly increases all three payoffs relative to $(3,4,1)$ while strictly improving at least one).

2.5.7 Symmetric Games

Def (Symmetric game). A two-player strategic game is *symmetric* if the players' sets of actions are the same and the players' preferences are represented by payoff functions u_1 and u_2 such that

$$u_1(a_1, a_2) = u_2(a_2, a_1) \quad \text{for every pair of actions } (a_1, a_2).$$

This means the game treats the two players identically. If we swap the players' roles, the payoffs swap in the same way.

Def (Symmetric Nash equilibrium). In a strategic game in which each player has the same set of actions, an action profile a^* is a *symmetric Nash equilibrium* if it is a Nash equilibrium and

$$a_i^* = a_j^* \quad \text{for all players } i, j.$$

So a symmetric Nash equilibrium is a Nash equilibrium where everyone plays the same action. A symmetric game can still have *no* symmetric Nash equilibrium.

Example (A symmetric game with many Nash equilibria)

Let $S = [0, 1]$ and define for each player i:

$$u_i(s_1, s_2) = \begin{cases} \max\{s_1, s_2\} & \text{if } (s_1, s_2) \neq (1, 1), \\ 0 & \text{if } (s_1, s_2) = (1, 1). \end{cases}$$

Claim: The Nash equilibrium set is

$$\Big\{(s_1, s_2) \in [0, 1]^2 : s_1 = 1 \text{ or } s_2 = 1\Big\} \setminus \{(1, 1)\}.$$

If $s_2 < 1$, player 1 can improve by choosing $s_1 = 1$, since then $\max\{s_1, s_2\} = 1$ and the outcome is not $(1, 1)$. So in any Nash equilibrium with $s_2 < 1$, we must have $s_1 = 1$. Symmetrically, if $s_1 < 1$, we must have $s_2 = 1$. Finally, $(1, 1)$ is not a Nash equilibrium because it gives payoff 0, but a player can deviate to (say) $1 - \varepsilon$ and get payoff 1.

Chapter 3: Applications of Nash Equilibrium

3.1 Cournot's Oligopoly

3.1.1 Model Setup

Consider an industry with two firms competing in quantities (Cournot competition).

- Firms are symmetric.
- Constant marginal cost: $c > 0$.
- No fixed costs.
- Linear inverse demand:

  $$p(Q) = a - bQ,$$

  where $a > c$, $b > 0$, and

  $$Q = q_i + q_j$$

  is total industry output.

> The condition $a > c$ ensures that producing a positive quantity is profitable, since the highest willingness to pay exceeds marginal cost.

3.1.2 Profit Maximization Problem

Firm i's profit is:

$$\pi_i(q_i, q_j) = p(Q)q_i - cq_i$$

Substituting demand:

$$\pi_i(q_i, q_j) = (a - b(q_i + q_j))q_i - cq_i$$

Each firm chooses $q_i \geq 0$, taking q_j as given:

$$\max_{q_i \geq 0} \quad (a - bq_i - bq_j)q_i - cq_i$$

3.1.3 Best Response Function

Take the first-order condition with respect to q_i:

$$\frac{\partial \pi_i}{\partial q_i} = a - 2bq_i - bq_j - c = 0$$

Solving for q_i:

$$BR_i(q_j) = q_i = \frac{a - c - bq_j}{2b}$$

By symmetry:

$$BR_j(q_i) = \frac{a - c - bq_i}{2b}$$

Each firm's best response is downward sloping in the rival's output:

$$\frac{\partial q_i}{\partial q_j} = -\frac{1}{2}$$

Thus quantities are **strategic substitutes**.

3.1.4 Nash Equilibrium

The Nash equilibrium is where best response functions intersect.

Solve:

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

$$q_j = \frac{a - c - bq_i}{2b}$$

By symmetry, let $q_i = q_j = q^*$:

$$q^* = \frac{a - c - bq^*}{2b}$$

Multiply both sides by $2b$:

$$2bq^* = a - c - bq^*$$

$$3bq^* = a - c$$

$$\boxed{q_i^* = q_j^* = \frac{a-c}{3b}}$$

Total output:

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

Price:

$$p^* = a - bQ^* = a - \frac{2(a-c)}{3} = \frac{a+2c}{3}$$

3.1.5 Graphical Representation

The Nash equilibrium is the intersection of the two best response functions.

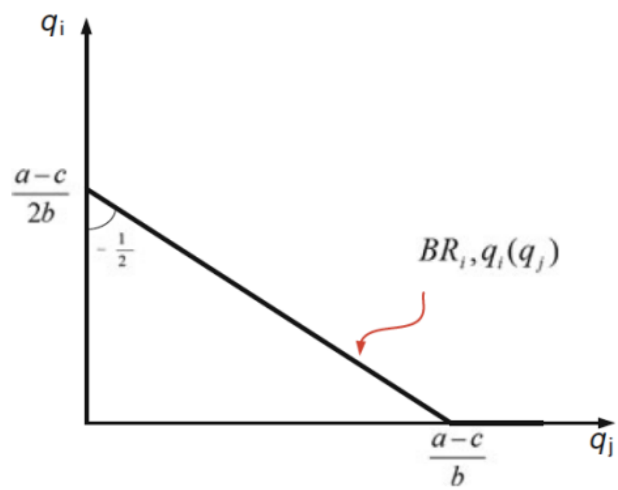

Figure 3.1: Best response functions and Nash equilibrium.

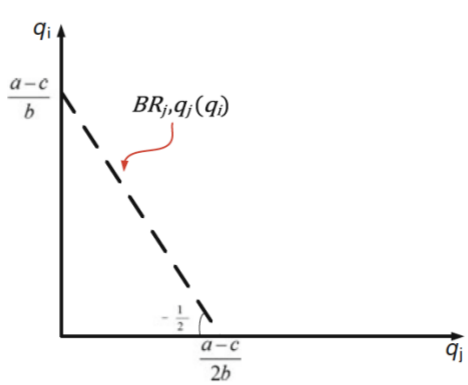

Figure 3.2: Best response of firm j.

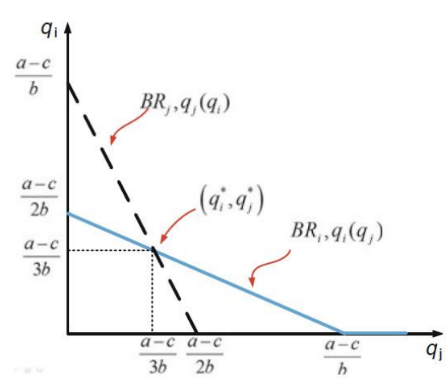

Figure 3.3: Best response of firm i.

The intersection point represents a Nash equilibrium because:

- Given q_j^*, firm i does not want to deviate.
- Given q_i^*, firm j does not want to deviate.

3.2 Examples

Congestion Game

Setup.
There are 1000 drivers traveling from A to B.

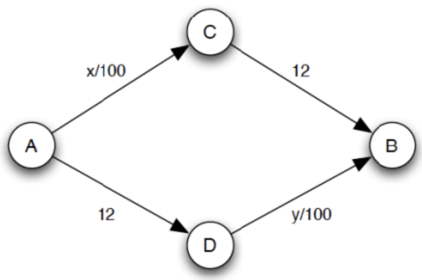

There are two possible routes:

- Upper route: $A \to C \to B$
- Lower route: $A \to D \to B$

Let:

$$x = \text{number of cars using edge } AC, \quad y = \text{number of cars using edge } DB.$$

Travel times:

$$\text{Time on } AC = \frac{x}{100}, \quad \text{Time on } DB = \frac{y}{100}.$$

$$\text{Time on } CB = 12, \quad \text{Time on } AD = 12.$$

Since all cars must travel from A to B,

$$x + y = 1000.$$

Each driver chooses a route to minimize travel time.

1) Nash equilibrium without the new road

Upper route time:

$$T_U = \frac{x}{100} + 12.$$

Lower route time:

$$T_L = 12 + \frac{y}{100}.$$

In equilibrium, any used routes must have equal travel times. Thus,

$$\frac{x}{100} + 12 = 12 + \frac{y}{100}.$$

Cancel 12:

$$\frac{x}{100} = \frac{y}{100}.$$

Thus:

$$x = y.$$

Since $x + y = 1000$,

$$x = y = 500.$$

Equilibrium travel time:

$$T = \frac{500}{100} + 12 = 5 + 12 = 17.$$

Total travel cost:

$$1000 \times 17 = 17{,}000.$$

2) After adding the new road $C \to D$ (zero cost)

Now a new route exists:

$$A \to C \to D \to B.$$

Its travel time is:

$$T_{ACDB} = \frac{x}{100} + 0 + \frac{y}{100}.$$

Now drivers compare:

$$T_{ACB} = \frac{x}{100} + 12,$$

$$T_{ADB} = 12 + \frac{y}{100},$$

$$T_{ACDB} = \frac{x}{100} + \frac{y}{100}.$$

In equilibrium, no driver can improve by switching routes.

Suppose all drivers use the new route. Then:

$$x = 1000, \quad y = 1000.$$

Then:

$$T_{ACDB} = \frac{1000}{100} + \frac{1000}{100} = 10 + 10 = 20.$$

Check deviations:

Upper route:

$$\frac{1000}{100} + 12 = 10 + 12 = 22.$$

Lower route:

$$12 + \frac{1000}{100} = 12 + 10 = 22.$$

Thus, 20 is strictly lower than 22, so no one deviates.

New Nash equilibrium:

$$x = 1000, \quad y = 1000.$$

Equilibrium travel time:

$$T = 20.$$

Total travel cost:

$$1000 \times 20 = 20{,}000.$$

Conclusion (Braess Paradox).

Adding the zero-cost road increases total travel time:

$$17{,}000 \quad \longrightarrow \quad 20{,}000.$$

The network becomes strictly worse after adding a new road.

War of Attrition

The War of Attrition is a strategic game in which two players compete over an indivisible object. Each player chooses a *concession time*, and waiting is costly.

Model

- **Players:** $i = 1, 2$.
- **Strategies:** Each player chooses a concession time $t_i \in [0, \infty)$.
- **Valuations:** Player i values the object at $v_i > 0$.
- **Cost of waiting:** Each unit of time costs 1.

Payoffs are given by

$$u_i(t_i, t_j) = \begin{cases} -t_i & \text{if } t_i < t_j, \\ \frac{1}{2}v_i - t_i & \text{if } t_i = t_j, \\ v_i - t_j & \text{if } t_i > t_j, \end{cases}$$

where $j \neq i$.

If player i concedes first, she pays the waiting cost and receives nothing. If the opponent concedes first, player i obtains the object but still incurs the waiting cost. If both concede at the same time, they split the object equally.

Best Responses

Fix t_j. Player i's optimal choice depends on the comparison between t_j and v_i:

- If $t_j < v_i$, it is optimal to wait longer than t_j.
- If $t_j = v_i$, player i is indifferent between conceding immediately and waiting longer than t_j.
- If $t_j > v_i$, it is optimal to concede immediately.

The best response correspondence is therefore

$$B_i(t_j) = \begin{cases} \{t_i : t_i > t_j\} & \text{if } t_j < v_i, \\ \{0\} \cup \{t_i : t_i > t_j\} & \text{if } t_j = v_i, \\ \{0\} & \text{if } t_j > v_i. \end{cases}$$

Nash Equilibria

A profile (t_1, t_2) is a Nash equilibrium if and only if either

$$t_1 = 0 \quad \text{and} \quad t_2 \geq v_1,$$

or

$$t_2 = 0 \quad \text{and} \quad t_1 \geq v_2.$$

In every equilibrium, one player concedes immediately and the other waits at least as long as the opponent's valuation. There is no symmetric pure strategy equilibrium.

Tournaments

Many strategic environments can be modeled as tournaments, where the probability of winning a prize depends on relative effort rather than absolute performance.

Model

Two firms compete in an R&D race.

- **Prize:** The winner receives 36.
- **Strategies:** Each firm chooses expenditure $x_i \in [0, 25]$.
- **Winning probability:**
$$\Pr(i \text{ wins}) = \frac{x_i}{x_1 + x_2}.$$
- **Cost:** Expenditure equals cost.

Firm 1's profit is

$$\pi_1(x_1, x_2) = 36\frac{x_1}{x_1 + x_2} - x_1,$$

and symmetrically,

$$\pi_2(x_1, x_2) = 36\frac{x_2}{x_1 + x_2} - x_2.$$

Best Response Functions

Firm 1 maximizes π_1 with respect to x_1. The first order condition is

$$\frac{\partial \pi_1}{\partial x_1} = 36\frac{x_2}{(x_1 + x_2)^2} - 1 = 0.$$

Rearranging,

$$36x_2 = (x_1 + x_2)^2.$$

Taking square roots,

$$6\sqrt{x_2} = x_1 + x_2.$$

Solving for x_1 gives firm 1's best response:

$$B_1(x_2) = 6\sqrt{x_2} - x_2.$$

By symmetry, firm 2's best response is

$$B_2(x_1) = 6\sqrt{x_1} - x_1.$$

Symmetric Nash Equilibrium

In a symmetric Nash equilibrium $x_1 = x_2 = x^*$. Substituting into the best response function:

$$x^* = 6\sqrt{x^*} - x^*.$$

Rearranging,

$$2x^* = 6\sqrt{x^*}.$$

Squaring both sides,

$$4(x^*)^2 = 36x^*,$$

which implies

$$x^* = 9.$$

Equilibrium Outcome

In equilibrium, both firms spend

$$x_1^* = x_2^* = 9.$$

Each firm wins with probability $\frac{1}{2}$ and earns expected profit

$$\pi_i = 36\frac{9}{18} - 9 = 18 - 9 = 9.$$

Chapter 4: Coming soon