Strategic Interactions

Lecture Notes

Nuh Aygün Dalkıran

LAST UPDATED: December 19, 2025

These lecture notes introduce the foundations of game theory and provide an initial exposure to the economics of information, with an emphasis on strategic reasoning in economic environments. They were originally prepared for **Econ 204 – Microeconomics: Strategic Interactions**, a second-year undergraduate course at Bilkent University.

Recent advances in generative artificial intelligence played a supporting role in the preparation of these notes, contributing to drafting, exposition, and refinement. All content has been reviewed by the author, who remains fully responsible for its accuracy, interpretation, and pedagogical choices. The use of generative AI is intended solely as an aid to instruction and does not substitute for independent reasoning, verification, or scholarly judgment.[a]

Despite careful preparation, errors or omissions may remain. Readers are warmly encouraged to report any such issues by contacting the author at `dalkiran@gmail.com`.

[a] The generative AI models used, in descending order of frequency, are Gemini 3.0 Pro, ChatGPT 5.2, and Grok 4.1.

Contents

Chapter 1

von Neumann–Morgenstern Expected Utility

1.1 Motivation: From Strategic Certainty to Risk

In our study of strategic interactions, we often begin by analyzing games where players choose their actions, and the combination of actions directly determines a final outcome. Consider the following simple strategic environment involving two players. Player 1 chooses between Up (U) and Down (D), and Player 2 chooses between Left (L) and Right (R). The matrix below shows the final outcome resulting from each action profile:

		Player 2 L	Player 2 R
Player 1	U	x	y
	D	z	t

Here, x, y, z, t are abstract outcomes. They could be monetary payoffs, market share allocations, policy results, or any other consequence of the players' actions.

If Player 1 knew for certain what Player 2 would choose, her decision would be simple. If Player 2 chooses L, Player 1 compares outcome x to outcome z. If Player 2 chooses R, she compares y to t.

However, in many realistic scenarios, Player 1 may not know Player 2's action with certainty. Instead, she might have a belief about the likelihood of Player 2's choices. Let's suppose Player 1 believes that Player 2 will play L with probability $q \in [0, 1]$ and R with probability $1 - q$.

Now, Player 1's decision is no longer a simple comparison of deterministic outcomes.

- If she chooses **U**, she faces a *risky* prospect: she receives outcome x with probability q and outcome y with probability $1 - q$.

- If she chooses **D**, she faces another risky prospect: she receives outcome z with probability q and outcome t with probability $1 - q$.

We can think of these risky prospects as **lotteries**. Player 1's choice is now a choice between two lotteries. To analyze her optimal strategy, we must first have a theory that explains how a rational individual makes decisions under risk. How do we formally compare these lotteries? This is precisely the question addressed by the **von Neumann-Morgenstern (vNM) Expected Utility Theory**.

1.2 The Formal Framework of Choice Under Risk

Let's formalize the decision problem.

Definition 1.2.1 (Outcomes)**.** Let $X = \{x_1, x_2, \dots, x_n\}$ be a finite set of basic, mutually exclusive outcomes.[1]

Definition 1.2.2 (Lottery)**.** A **simple lottery** p on the set of outcomes X is a probability distribution over X. We can represent it as a vector $p = (p_1, p_2, \dots, p_n)$ where p_i is the probability of outcome x_i occurring, satisfying:

1. $p_i \geq 0$ for all $i = 1, \dots, n$.
2. $\sum_{i=1}^{n} p_i = 1$.

We denote the set of all possible lotteries over X by $\Delta(X)$.

For example, a lottery that gives outcome x_1 with certainty is denoted $p = (1, 0, \dots, 0)$. In our motivating game, if $X = \{x, y, z, t\}$, Player 1's choice of U induces the lottery $(q, 1-q, 0, 0)$, while choosing D induces the lottery $(0, 0, q, 1-q)$.

Our goal is to model a decision maker's preferences over these lotteries. We introduce a preference relation $\succsim$ on the set $\Delta(X)$. For any two lotteries $p, q \in \Delta(X)$, $p \succsim q$ is read as "lottery p is at least as good as lottery q." From this, we can derive:

- Strict preference: $p \succ q \iff p \succsim q$ and not $q \succsim p$.
- Indifference: $p \sim q \iff p \succsim q$ and $q \succsim p$.

1.3 Axioms of Rational Choice Under Risk

The vNM theory posits that a rational decision maker's preferences over lotteries should satisfy four key axioms.

Axiom 1.3.1 (Completeness)**.** For any two lotteries $p, q \in \Delta(X)$, either $p \succsim q$ or $q \succsim p$ (or both).

Interpretation: The decision maker can always compare any two lotteries and decide which one is preferred or if she is indifferent between them.

Axiom 1.3.2 (Transitivity)**.** For any three lotteries $p, q, r \in \Delta(X)$, if $p \succsim q$ and $q \succsim r$, then $p \succsim r$.

Interpretation: Preferences are internally consistent. If p is at least as good as q, and q is at least as good as r, then a chain of preference holds, and p must be at least as good as r.

A preference relation satisfying Completeness and Transitivity is called a *rational preference relation.* The next two axioms are specific to the context of choice under risk.

Axiom 1.3.3 (Continuity (Archimedean))**.** For any three lotteries $p, q, r \in \Delta(X)$ such that $p \succ q \succ r$, there exists a probability $\alpha \in (0, 1)$ such that the compound lottery $\alpha p + (1-\alpha) r \sim q$.

Interpretation: This axiom rules out "infinitely desirable" or "infinitely undesirable" outcomes. If you have a best lottery (p) and a worst lottery (r), any intermediate lottery (q) can be made indifferent to some probabilistic mix of the best and worst. There's always a tipping point.

[1] All of the results below extend to the cases where X is countably infinite or uncountable. But then, lotteries become probability measures and expected utility becomes an integral. To avoid technical complexities, we present only the finite case, which is simply a powerful and accessible starting point for teaching the theory.

Axiom 1.3.4 (Independence)**.** For any three lotteries $p, q, r \in \Delta(X)$ and any probability $\alpha \in (0,1)$,

$$p \succsim q \iff \alpha p + (1-\alpha)r \succsim \alpha q + (1-\alpha)r$$

Interpretation: This is the most crucial (and sometimes controversial) axiom. It states that your preference between two lotteries (p and q) should not be reversed by mixing both of them with an irrelevant third lottery (r). If you prefer p to q, you should also prefer a ticket that gives you p with probability α and r with probability $1-\alpha$ to a ticket that gives you q with probability α and r with probability $1-\alpha$. The common part of the lotteries "cancels out."

1.4 The Expected Utility Theorem

The power of these four axioms lies in the fact that they are precisely the conditions needed to guarantee that a decision maker's preferences can be represented by a very convenient mathematical structure: the expectation of a utility function.

Theorem 1.4.1 (von Neumann-Morgenstern, 1944)**.** A preference relation $\succsim$ on the space of lotteries $\Delta(X)$ satisfies the above listed four axioms (completeness, transitivity, continuity, and independence) if and only if there exists a function $u : X \to \mathbb{R}$, typically referred to as a **Bernoulli utility function**, such that for any two lotteries $p = (p_1, \ldots, p_n)$ and $q = (q_1, \ldots, q_n)$,

$$p \succsim q \iff \sum_{i=1}^{n} p_i u(x_i) \geq \sum_{i=1}^{n} q_i u(x_i)$$

The value $U(p) = \mathbb{E}_p[u(x)] = \sum_{i=1}^{n} p_i u(x_i)$ is called the **vNM expected utility** of the lottery p.

Remark 1.4.1. This theorem is profound. It states that if a person's choices conform to the four axioms of rationality, then they behave *as if* they are assigning a numerical utility value $u(x_i)$ to each basic outcome x_i and then choosing the lottery that yields the highest probability-weighted average of these utilities.

An important feature of the vNM utility function is its uniqueness property. Unlike ordinal utility where any strictly increasing transformation represents the same preferences, vNM utility is cardinal in a specific sense.

Theorem 1.4.2 (Uniqueness up to Positive Affine Transformation)**.** Suppose a vNM utility function $u : X \to \mathbb{R}$ represents the preference relation $\succsim$. Then another utility function $v : X \to \mathbb{R}$ represents the same preferences if and only if there exist constants $a > 0$ and $b \in \mathbb{R}$ such that for all $x_i \in X$,

$$v(x_i) = a \cdot u(x_i) + b$$

Proof Sketch. Let's show that $v(x) = au(x) + b$ represents the same preferences. We need to check if $p \succsim q$ is equivalent to $\mathbb{E}_p[v(x)] \geq \mathbb{E}_q[v(x)]$.

$$\begin{aligned}
\mathbb{E}_p[v(x)] &= \sum_{i=1}^{n} p_i v(x_i) = \sum_{i=1}^{n} p_i (a \cdot u(x_i) + b) \\
&= a \sum_{i=1}^{n} p_i u(x_i) + b \sum_{i=1}^{n} p_i \\
&= a \cdot \mathbb{E}_p[u(x)] + b
\end{aligned}$$

Since $a > 0$, it follows that:

$$\mathbb{E}_p[v(x)] \geq \mathbb{E}_q[v(x)] \iff a \cdot \mathbb{E}_p[u(x)] + b \geq a \cdot \mathbb{E}_q[u(x)] + b \iff \mathbb{E}_p[u(x)] \geq \mathbb{E}_q[u(x)]$$

This confirms that $v(x)$ represents the same preferences over lotteries as $u(x)$.

Remark 1.4.2. This property implies that while the absolute values of utility are meaningless, and utility differences are not directly interpretable (e.g., $u(x) - u(y) = 2[u(z) - u(w)]$ does not mean the first preference gap is "twice as strong"), the ordering of expected utilities is what matters. This is a stronger form of measurement than ordinal utility but weaker than cardinal measures like temperature or weight.

1.5 Conclusion: Back to the Game

Let's return to our initial strategic problem. Player 1 must choose between U and D, where Player 2 plays L with probability q. This is a choice between two lotteries:

- p_U: Outcome x with probability q, outcome y with probability $1 - q$.
- p_D: Outcome z with probability q, outcome t with probability $1 - q$.

If Player 1's preferences satisfy the vNM axioms, then her behavior can be modeled by an expected utility maximizer. She has a vNM utility function $u(\cdot)$ over the set of outcomes $\{x, y, z, t\}$. Her decision rule is simple: she will compare the expected utility of the two lotteries.

The expected utility of choosing U is:

$$U(p_U) = q \cdot u(x) + (1 - q) \cdot u(y)$$

The expected utility of choosing D is:

$$U(p_D) = q \cdot u(z) + (1 - q) \cdot u(t)$$

Player 1 will choose U if $U(p_U) > U(p_D)$, choose D if $U(p_D) > U(p_U)$, and will be indifferent if $U(p_U) = U(p_D)$.

The vNM theory provides the essential foundation for analyzing games with uncertainty. It allows us to replace outcomes in a game matrix with utility values (payoffs) and analyze players' strategies over beliefs and randomized actions, which will be central to our concepts of mixed strategy Nash equilibrium and Bayesian games.

1.6 Utility over Money and Risk Attitudes

So far, our set of outcomes X has been abstract. A very common and important application is when the outcomes are different amounts of money. Let's denote an outcome of receiving w dollars as simply w. A lottery is then a probability distribution over monetary outcomes. A natural first guess might be that a rational agent should simply try to maximize their expected monetary value. This idea, however, runs into a famous problem.

Example 1.6.1 (The St. Petersburg Paradox)**.** Consider a game where a fair coin is tossed repeatedly until it comes up heads for the first time. If the first head appears on the k^{th} toss, the player receives a prize of $\$2^k$.

- The probability of the first head on the 1st toss is $\frac{1}{2}$, payoff is $\$2^1 = \$2$.
- The probability of the first head on the 2nd toss is $(\frac{1}{2})^2 = \frac{1}{4}$, payoff is $\$2^2 = \$4$.
- The probability of the first head on the k^{th} toss is $(\frac{1}{2})^k$, payoff is $\$2^k$.

What is the expected monetary value (EV) of this lottery?

$$EV = \sum_{k=1}^{\infty} \left(\frac{1}{2}\right)^k \cdot 2^k = \sum_{k=1}^{\infty} 1 = 1 + 1 + 1 + \cdots = \infty$$

An agent who maximizes expected value should be willing to pay any finite price to play this game. However, empirically, most people would only be willing to pay a small, finite amount (perhaps \$10-\$20). This paradox suggests that the expected value of money is not what people maximize.

The resolution, first proposed by Daniel Bernoulli, is that people do not value money linearly. Instead, they have a utility function over money, $u(w)$, and they act to maximize their *expected utility*, not expected value. For example, if an agent's utility for wealth is $u(w) = \ln(w)$, the expected utility of the St. Petersburg lottery is finite, resolving the paradox.

This insight leads us to define an agent's attitude toward risk based on the shape of their vNM utility function over money (often called a Bernoulli utility function in this context).

Definition 1.6.1 (Fair Gamble)**.** A **fair gamble** is a lottery that has an expected monetary value of zero.

For an agent with initial wealth w, a simple fair gamble could be a 50-50 chance to win or lose $\epsilon > 0$. The agent's final wealth would be $w + \epsilon$ or $w - \epsilon$, each with probability $\frac{1}{2}$. The expected wealth after the gamble is $\frac{1}{2}(w + \epsilon) + \frac{1}{2}(w - \epsilon) = w$, which is the same as their wealth without the gamble.

Definition 1.6.2 (Risk Attitudes)**.** Let an agent have a vNM utility function $u(\cdot)$ over wealth.

1. The agent is **risk-averse** if they reject any fair gamble. This means they prefer the certainty-equivalent wealth to the gamble:
$$u(w) > \frac{1}{2}u(w + \epsilon) + \frac{1}{2}u(w - \epsilon) \quad \forall w, \epsilon > 0$$
This inequality is the definition of a **strictly concave** function. Thus, a risk-averse agent has a strictly concave utility function.

2. The agent is **risk-neutral** if they are indifferent to any fair gamble:
$$u(w) = \frac{1}{2}u(w + \epsilon) + \frac{1}{2}u(w - \epsilon) \quad \forall w, \epsilon > 0$$
This implies the utility function is **linear**.

3. The agent is **risk-loving** (or risk-seeking) if they accept any fair gamble:
$$u(w) < \frac{1}{2}u(w + \epsilon) + \frac{1}{2}u(w - \epsilon) \quad \forall w, \epsilon > 0$$
This inequality is the definition of a **strictly convex** function.

Remark 1.6.1 (Jensen's Inequality)**.** The connection between concavity and risk aversion is a direct consequence of Jensen's inequality. For any random variable $\tilde{w}$ and any strictly concave function u, the inequality states that $\mathbb{E}[u(\tilde{w})] < u(\mathbb{E}[\tilde{w}])$. For a risk-averse agent facing a fair gamble, their expected wealth is $\mathbb{E}[\tilde{w}] = w$. Their expected utility from the gamble is $\mathbb{E}[u(\tilde{w})]$. Jensen's inequality implies $\mathbb{E}[u(\tilde{w})] < u(w)$, so the agent prefers the certain wealth w to the gamble.

Remark 1.6.2 (Normalization for Risk-Neutrality)**.** If an agent is risk-neutral, their utility function must be linear, i.e., of the form $u(w) = aw + b$ for some constants $a > 0$ and b. By the uniqueness property of vNM utility functions, we can apply a positive affine transformation to simplify this function without changing the agent's preferences. Let's define a new utility function $v(w) = \frac{1}{a}u(w) - \frac{b}{a}$. This gives:
$$v(w) = \frac{1}{a}(aw + b) - \frac{b}{a} = w$$
Therefore, for a risk-neutral agent, we can assume without loss of generality that their utility function is the identity function, $u(w) = w$. Such an agent is an expected monetary value maximizer.

Chapter 2

Nash Equilibrium

2.1 The Normal Form Game

We begin with the formal definition of a normal (or strategic) form game, which is the basic building block for our analysis.

Definition 2.1.1 (Normal Form Game). A normal form game is a tuple $\langle N, (A_i)_{i \in N}, (u_i)_{i \in N} \rangle$, where:[1]

- $N = \{1, 2, \ldots, n\}$ is a finite set of players.
- A_i is a finite set of pure strategies (or actions) available to player i. Let $A = A_1 \times A_2 \times \cdots \times A_n$ be the set of all possible pure strategy profiles. An element $a = (a_1, \ldots, a_n) \in A$ is a strategy profile.
- $u_i : A \to \mathbb{R}$ is the payoff function for player i. This function assigns a utility value (a Bernoulli payoff) to each possible outcome of the game.

For a strategy profile $a = (a_1, \ldots, a_n)$, we often use the notation $a_{-i} = (a_1, \ldots, a_{i-1}, a_{i+1}, \ldots, a_n)$ to denote the strategy profile of all players *except* player i. Thus, we can write $a = (a_i, a_{-i})$.

2.2 Pure Strategy Nash Equilibrium

The central solution concept in game theory is the Nash Equilibrium. It describes a steady-state or stable outcome in a strategic interaction, where no single player has an incentive to change their strategy, given what everyone else is doing.

2.2.1 Best-Reply Correspondence

To formalize this idea, we first need to define what it means for a player to be playing optimally given the strategies of others.

Definition 2.2.1 (Best Reply). Player i's strategy $a_i^* \in A_i$ is a **best reply** (or best response) to the strategies of the other players, $a_{-i} \in A_{-i}$, if it maximizes player i's payoff. That is:

$$u_i(a_i^*, a_{-i}) \geq u_i(a_i, a_{-i}) \quad \text{for all } a_i \in A_i$$

The set of all such best replies for player i to a given a_{-i} is the **best-reply correspondence** $BR_i : A_{-i} \to 2^{A_i}$ which is defined formally as:

$$BR_i(a_{-i}) := \arg\max_{a_i \in A_i} u_i(a_i, a_{-i})$$

[1] A normal form game is also known as a strategic form game or a simultaneous-move game.

Here, 2^{A_i} denotes the power set of A_i, which is the set of all possible subsets of A_i. This indicates that the correspondence can return a set of best replies, not just a single one. The arg max operator returns the value(s) of the argument a_i that maximize the function $u_i(a_i, a_{-i})$.

2.2.2 Definitions of Nash Equilibrium

Using the concept of a best reply, we can now define a Pure Strategy Nash Equilibrium (PSNE) in several equivalent ways.

Definition 2.2.2 (Pure Nash Equilibrium)**.** A strategy profile $a^* = (a_1^*, \ldots, a_n^*) \in A$ is a **Pure Strategy Nash Equilibrium** (PSNE) if for every player $i \in N$, a_i^* is a best reply to a_{-i}^*.

This definition leads to the following equivalent formalizations:

1. **No Profitable Unilateral Deviation:** A strategy profile a^* is a PSNE if and only if no player i can increase her payoff by unilaterally deviating to another strategy a_i, given that all other players stick to a_{-i}^*. Formally:

$$\forall i \in N, \quad u_i(a_i^*, a_{-i}^*) \geq u_i(a_i, a_{-i}^*) \quad \text{for all } a_i \in A_i$$

2. **Intersection of Best Replies:** A strategy profile $a^* = (a_1^*, \ldots, a_n^*)$ is a PSNE if and only if each player's strategy is a best reply to the other players' strategies. Formally:

$$a_i^* \in BR_i(a_{-i}^*) \quad \text{for all } i \in N$$

2.3 Mixed Strategy Nash Equilibrium

In many games, a pure strategy equilibrium does not exist. We can extend our analysis by allowing players to randomize over their pure strategies.

Definition 2.3.1 (Mixed Strategy)**.** A **mixed strategy** for player i, denoted α_i, is a probability distribution over the set of pure strategies A_i. We denote the set of all mixed strategies for player i by $\Delta(A_i)$. A mixed strategy profile is a profile $\alpha = (\alpha_1, \ldots, \alpha_n)$.

When players play mixed strategies, the outcome is uncertain. We assume players are vNM expected utility maximizers. The payoff to a mixed strategy profile α for player i is their expected utility:

$$u_i(\alpha) = \sum_{a \in A} \left(\prod_{j=1}^{n} \alpha_j(a_j) \right) u_i(a)$$

where $\alpha_j(a_j)$ is the probability player j assigns to pure strategy a_j.

Definition 2.3.2 (Mixed Strategy Nash Equilibrium)**.** A mixed strategy profile $\alpha^* = (\alpha_1^*, \ldots, \alpha_n^*)$ is a **Mixed Strategy Nash Equilibrium** (MSNE) if for every player $i \in N$, α_i^* is a best reply to α_{-i}^*. That is:

$$u_i(\alpha_i^*, \alpha_{-i}^*) \geq u_i(\alpha_i, \alpha_{-i}^*) \quad \text{for all } \alpha_i \in \Delta(A_i)$$

Theorem 2.3.1 (Indifference Principle)**.** In a MSNE, any pure strategy played with positive probability by a player must yield that player the same expected payoff. If a pure strategy yields a strictly lower expected payoff, it must be played with zero probability.

This principle is the key tool for computing mixed strategy equilibria.

2.4 Examples and Analysis of Nash Equilibrium in 2x2 Games

We now analyze five canonical 2-player, 2-strategy games. Let Player 1's mixed strategy be to play their first action with probability p and their second with $1-p$. Let Player 2's mixed strategy be to play their first action with probability q and their second with $1-q$.

2.4.1 Prisoners' Dilemma

Two suspects are arrested and cannot communicate. If both Cooperate (C), they each get a small sentence. If one Defects (D) and the other cooperates, the defector goes free and the cooperator gets a long sentence. If both defect, they both get a medium sentence.

		Player 2	
		C (q)	D ($1-q$)
Player 1	C (p)	3, 3	1, 4
	D ($1-p$)	4, 1	2, 2

Pure Nash: For Player 1, D is a strictly dominant strategy ($4 > 3$ and $2 > 1$). Similarly, D is strictly dominant for Player 2 ($4 > 3$ and $2 > 1$). The unique PSNE is **(D, D)**.

Mixed Nash: Since D is a strictly dominant strategy for both players, they will never mix. The only MSNE is the pure one, where $p = 0$ and $q = 0$.

Best Replies: $BR_1(q) : p = 0$ for all $q \in [0, 1]$. $BR_2(p) : q = 0$ for all $p \in [0, 1]$. The correspondences are single-valued and constant. The unique intersection is at $(p = 0, q = 0)$.

Formally:

$$BR_1(q) = 0 \;\text{ for all }\; q \in [0,1] \qquad \& \qquad BR_2(p) = 0 \;\text{ for all }\; p \in [0,1].$$

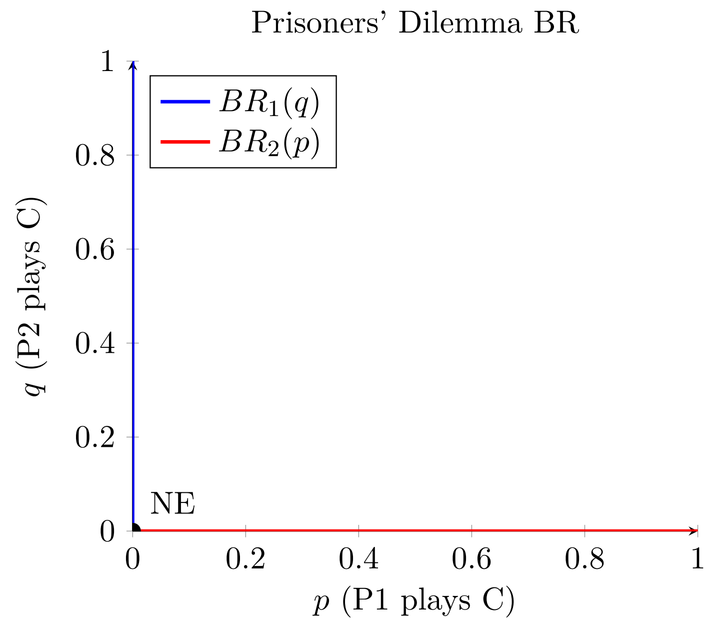

2.4.2 Battle of the Sexes

A couple wants to go out. One wants to go to the Football game (F), the other to the Ballet (B). They prefer going together to going alone.

		Player 2	
		F (q)	B ($1-q$)
Player 1	F (p)	3, 1	0, 0
	B ($1-p$)	0, 0	1, 3

Pure Nash: There are two PSNE: **(F, F)** and **(B, B)**. These are coordination equilibria.

Mixed Nash:

- P1 is indifferent if $u_1(F, q) = u_1(B, q) \implies 3q + 0(1-q) = 0q + 1(1-q) \implies 3q = 1 - q \implies q = 1/4$.
- P2 is indifferent if $u_2(p, F) = u_2(p, B) \implies 1p + 0(1-p) = 0p + 3(1-p) \implies p = 3 - 3p \implies p = 3/4$.

The fully-mixed MSNE is $\alpha^* = (\alpha_1^*, \alpha_2^*) = (3/4F \oplus 1/4B, 1/4F \oplus 3/4B)$.

Best Replies:

- $BR_1(q)$: $p = 1$ if $q > 1/4$, $p = 0$ if $q < 1/4$, $p \in [0, 1]$ if $q = 1/4$.
- $BR_2(p)$: $q = 1$ if $p > 3/4$, $q = 0$ if $p < 3/4$, $q \in [0, 1]$ if $p = 3/4$.

Formally:

$$BR_1(q) = \begin{cases} 1 & \text{if } q > \frac{1}{4}, \\ [0,1] & \text{if } q = \frac{1}{4}, \\ 0 & \text{if } q < \frac{1}{4}, \end{cases} \qquad BR_2(p) = \begin{cases} 1 & \text{if } p > \frac{3}{4}, \\ [0,1] & \text{if } p = \frac{3}{4}, \\ 0 & \text{if } p < \frac{3}{4}. \end{cases}$$

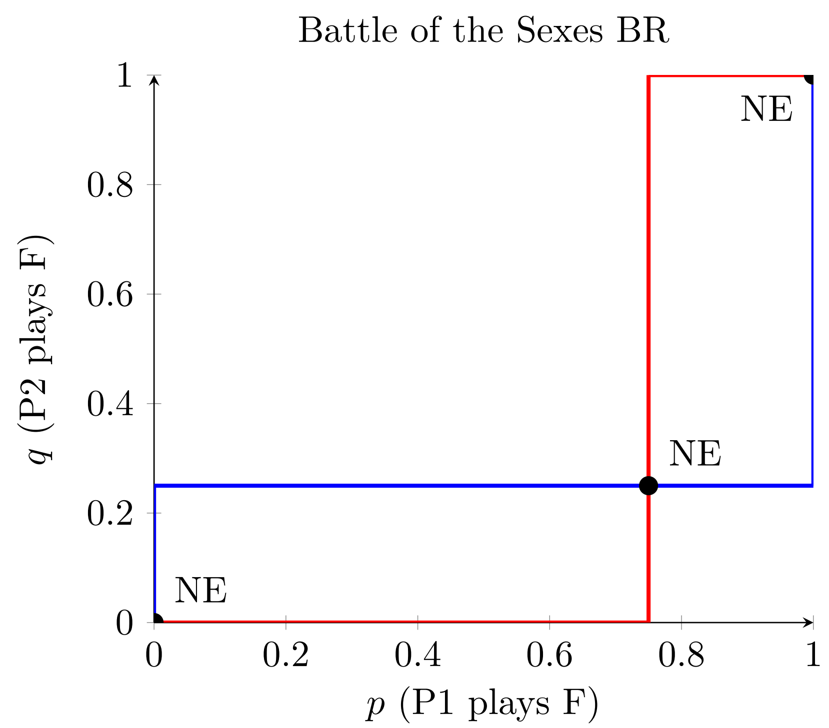

2.4.3 Stag Hunt

Two hunters can choose to hunt a Stag (S) or a Hare (H). Hunting a stag requires cooperation for success.

		Player 2	
		S (q)	H ($1-q$)
Player 1	S (p)	7, 7	0, 1
	H ($1-p$)	1, 0	1, 1

Pure Nash: Two PSNE: **(S, S)** (Pareto dominant) and **(H, H)** (risk dominant).

Mixed Nash:

- P1 is indifferent if $u_1(S,q) = u_1(H,q) \implies 7q + 0(1-q) = 1q + 1(1-q) \implies 7q = 1 \implies q = 1/7$.
- P2 is indifferent if $u_2(S,p) = u_2(H,p) \implies 7p + 0(1-p) = 1p + 1(1-p) \implies 7p = 1 \implies p = 1/7$.

The fully-mixed MSNE is $\alpha^* = (1/7S \oplus 6/7H, 1/7S \oplus 6/7H)$.

Best Replies:

- $BR_1(q)$: $p = 1$ if $q > 1/7$, $p = 0$ if $q < 1/7$, $p \in [0,1]$ if $q = 1/7$.
- $BR_2(p)$: $q = 1$ if $p > 1/7$, $q = 0$ if $p < 1/7$, $q \in [0,1]$ if $p = 1/7$.

Formally:

$$BR_1(q) = \begin{cases} 1 & \text{if } q > \frac{1}{7} \\ [0,1] & \text{if } q = \frac{1}{7} \\ 0 & \text{if } q < \frac{1}{7} \end{cases}, \qquad BR_2(p) = \begin{cases} 1 & \text{if } p > \frac{1}{7} \\ [0,1] & \text{if } p = \frac{1}{7} \\ 0 & \text{if } p < \frac{1}{7} \end{cases}.$$

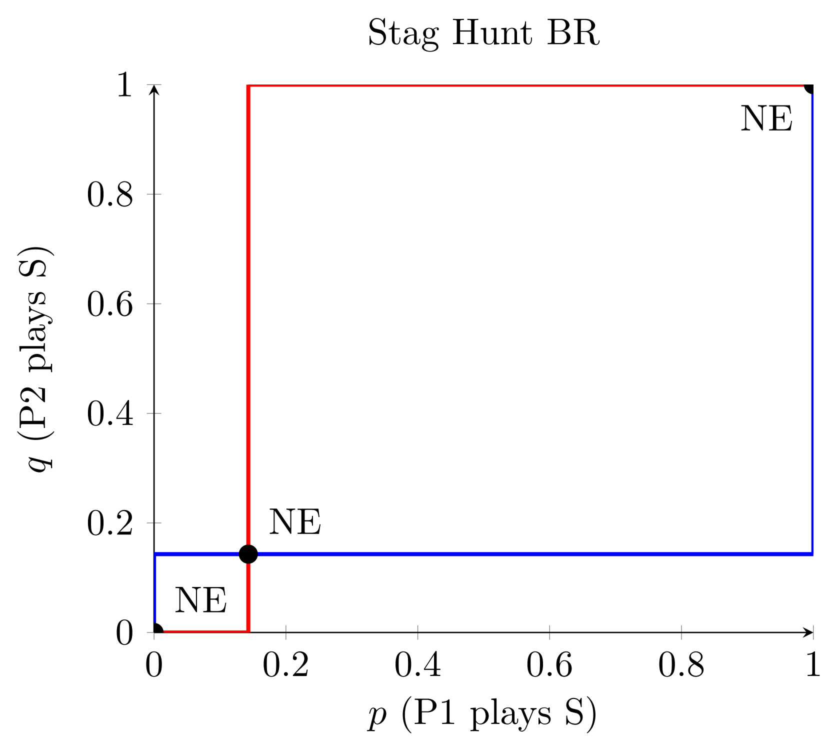

2.4.4 Hawk-Dove Game

Two animals contest a resource. They can be aggressive (Hawk, H) or passive (Dove, D). A Hawk fights other Hawks (costly), and exploits Doves.[2]

		Player 2	
		H (q)	D ($1-q$)
Player 1	H (p)	-10, -10	1, 0
	D ($1-p$)	0, 1	-1, -1

Pure Nash: Two PSNE: **(H, D)** and **(D, H)**. This is an anti-coordination game.

Mixed Nash:

- P1 is indifferent if $-10q + 1(1-q) = 0q - 1(1-q) \implies -11q + 1 = q - 1 \implies 2 = 12q \implies q = 1/6$.
- P2 is indifferent if $-10p + 1(1-p) = 0p - 1(1-p) \implies -11p + 1 = p - 1 \implies 2 = 12p \implies p = 1/6$.

The fully-mixed MSNE is $\alpha^* = (1/6H \oplus 5/6D, 1/6H \oplus 5/6D)$.

Best Replies:

- $BR_1(q)$: $p = 1$ if $q < 1/6$, $p = 0$ if $q > 1/6$, $p \in [0,1]$ if $q = 1/6$.
- $BR_2(p)$: $q = 1$ if $p < 1/6$, $q = 0$ if $p > 1/6$, $q \in [0,1]$ if $p = 1/6$.

Formally:

$$BR_1(q) = \begin{cases} 1 & \text{if } q < \frac{1}{6}, \\ [0,1] & \text{if } q = \frac{1}{6}, \\ 0 & \text{if } q > \frac{1}{6}, \end{cases} \qquad BR_2(p) = \begin{cases} 1 & \text{if } p < \frac{1}{6}, \\ [0,1] & \text{if } p = \frac{1}{6}, \\ 0 & \text{if } p > \frac{1}{6}. \end{cases}$$

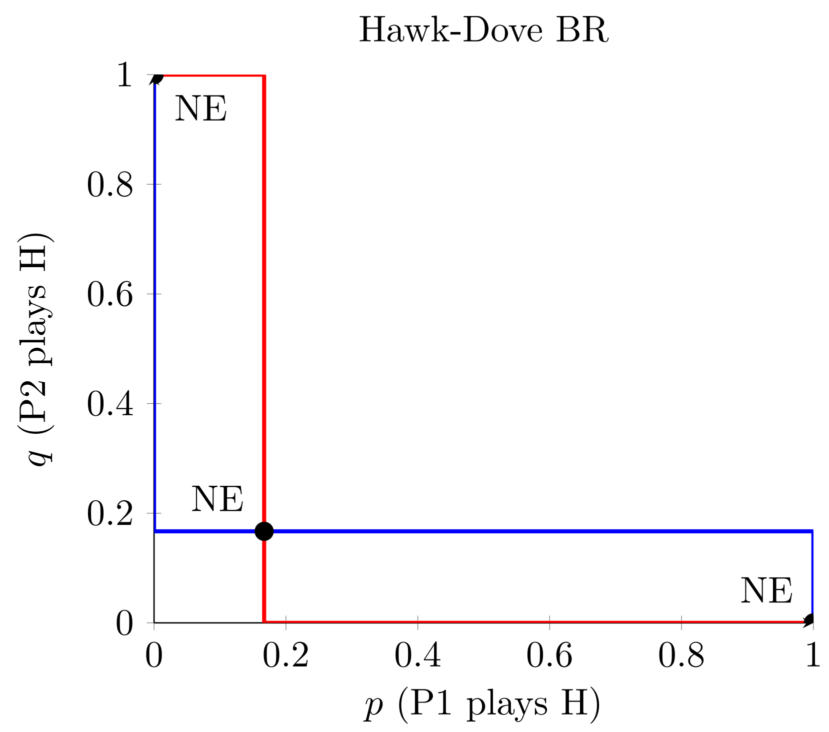

[2]The payoff structure here slightly departs from the standard Hawk–Dove formulation, which embodies the "Doves share with Doves" principle.

2.4.5 Matching Pennies

Player 1 wants to match Player 2's coin choice (Heads or Tails). Player 2 wants to mismatch.

		Player 2	
		H (q)	T ($1-q$)
Player 1	H (p)	1, -1	-1, 1
	T ($1-p$)	-1, 1	1, -1

Pure Nash: There is no PSNE. The best replies cycle: (H,H) → (H,T) → (T,T) → (T,H) → (H,H).

Mixed Nash:

- P1 is indifferent if $1q - 1(1-q) = -1q + 1(1-q) \implies 2q - 1 = -2q + 1 \implies 4q = 2 \implies q = 1/2$.
- P2 is indifferent if $-1p + 1(1-p) = 1p - 1(1-p) \implies -2p + 1 = 2p - 1 \implies 2 = 4p \implies p = 1/2$.

The unique MSNE is $\alpha^* = (1/2H \oplus 1/2T, 1/2H \oplus 1/2T)$.

Best Replies:

- $BR_1(q)$: $p = 1$ if $q > 1/2$, $p = 0$ if $q < 1/2$, $p \in [0, 1]$ if $q = 1/2$.
- $BR_2(p)$: $q = 0$ if $p > 1/2$, $q = 1$ if $p < 1/2$, $q \in [0, 1]$ if $p = 1/2$.

Formally:

$$BR_1(q) = \begin{cases} 1 & \text{if } q > \frac{1}{2} \\ [0,1] & \text{if } q = \frac{1}{2} \\ 0 & \text{if } q < \frac{1}{2} \end{cases}, \qquad BR_2(p) = \begin{cases} 0 & \text{if } p > \frac{1}{2} \\ [0,1] & \text{if } p = \frac{1}{2} \\ 1 & \text{if } p < \frac{1}{2} \end{cases}.$$

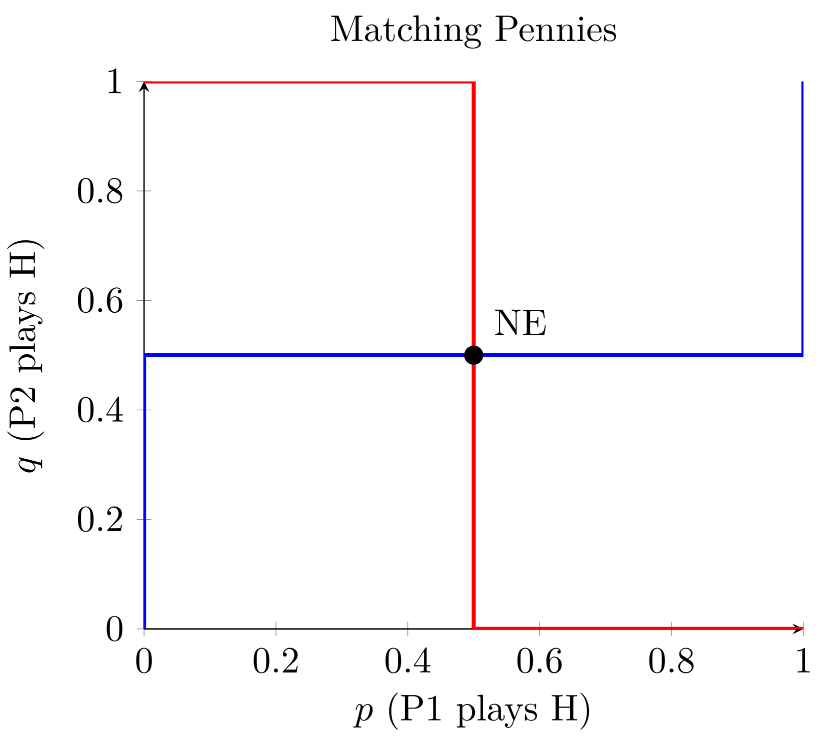

2.5 Existence of Nash Equilibrium

A natural and fundamental question is whether a Nash Equilibrium is guaranteed to exist in any game. In his 1951 dissertation, John Nash showed that every finite game has at least one mixed strategy Nash Equilibrium. His proof is a landmark achievement and a beautiful application of a mathematical tool known as a **fixed point theorem**.

Theorem 2.5.1 (Nash's Existence Theorem)**.** Every finite normal form game has at least one mixed strategy Nash equilibrium.

Remark 2.5.1. John Nash's original proof, a mere page long, was published in the *Proceedings of the National Academy of Sciences (PNAS)* in 1950. It elegantly demonstrated that a game's best-reply correspondence satisfies the conditions of Kakutani's Fixed Point Theorem, a major result in its own right. The core idea is to show that a Nash Equilibrium is equivalent to a "fixed point" of the best-reply correspondence.

2.5.1 Fixed Point Theorems

As highlighted above, Nash's proof uses **Kakutani's Fixed Point Theorem**, which is a generalization of the more intuitive **Brouwer's Fixed Point Theorem** to set-valued functions (correspondences). Below, we define the properties that are required to state these theorems.

Remark 2.5.2 (Topological Properties)**.** Let $S \subset \mathbb{R}^n$.

- **Bounded:** A set S is *bounded* if there exists a real number $r > 0$ such that S is contained in a ball of radius r centered at the origin.
- **Closed:** A set S is *closed* if it contains all of its limit points. Intuitively, this means that if you have a sequence of points in S that converges to a point, that limit point must also be in S.
- **Compact:** In Euclidean space ($\mathbb{R}^n$), a set S is *compact* if and only if it is both closed and bounded (this is the Heine-Borel Theorem). Compactness is a crucial property for ensuring the existence of maxima or minima for continuous functions.
- **Convex:** A set S is *convex* if for any two points $x, y \in S$, the line segment connecting them is also entirely contained within S. Formally, for any $x, y \in S$ and any $\lambda \in [0, 1]$, we have $\lambda x + (1 - \lambda)y \in S$.

Theorem 2.5.2 (Brouwer's Fixed Point Theorem)**.** Let $S \subset \mathbb{R}^n$ be a non-empty, compact, and convex set. If $f : S \to S$ is a continuous function, then there exists a point $x^* \in S$ such that $f(x^*) = x^*$. Such a point is called a fixed point of f.

The intuition can be grasped in one dimension. Let $S = [a, b]$ be a closed interval on the real line. A continuous function $f : [a, b] \to [a, b]$ must cross the 45-degree line ($y = x$) at least once. The point of intersection is the fixed point.

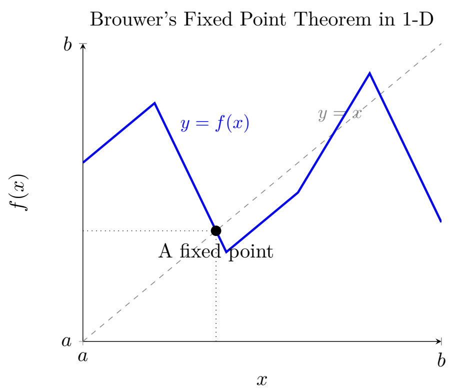

For game theory, we need a more powerful theorem because best replies are not always single points; they can be sets (e.g., when a player is indifferent between multiple strategies). This requires generalizing from a function f to a correspondence Φ.

Theorem 2.5.3 (Kakutani's Fixed Point Theorem)**.** Let $S \subset \mathbb{R}^n$ be a non-empty, compact, and convex set. Let $\Phi : S \to 2^S$ be a correspondence that satisfies:

1. For every $x \in S$, the set $\Phi(x)$ is non-empty and convex.

2. The correspondence Φ is upper hemi-continuous (has a closed graph).

Then there exists a point $x^* \in S$ such that $x^* \in \Phi(x^*)$.

Kakutani's theorem generalizes Brouwer's. A continuous function is a special case of an upper hemi-continuous correspondence where the output sets $\Phi(x)$ are always singletons. The requirement $\Phi(x)$ be convex is trivially satisfied for a function (a single point is a convex set).

2.5.2 Application to Nash Equilibrium

To apply Kakutani's theorem, we define the game's overall best-reply correspondence. Let $\mathcal{A} = \times_{i \in N} \Delta(A_i)$ be the set of all mixed strategy profiles. The best-reply correspondence for the whole game, $BR : \mathcal{A} \to 2^{\mathcal{A}}$, is defined as:

$$BR(\alpha) := BR_1(\alpha_{-1}) \times BR_2(\alpha_{-2}) \times \cdots \times BR_n(\alpha_{-n})$$

Equivalently, $BR(\alpha) := \{\beta \in \times_{i \in N} \Delta(A_i) \mid \beta_i \in BR_i(\alpha_{-i}) \text{ for all } i \in N\}$. This correspondence takes a strategy profile α as input and returns the set of strategy profiles where each player's strategy is a best reply to the others.

A **fixed point** of this correspondence is a strategy profile $\alpha^* \in \mathcal{A}$ such that $\alpha^* \in BR(\alpha^*)$. By definition, this means $\alpha_i^* \in BR_i(\alpha_{-i}^*)$ for all players i. This is precisely our definition of a Mixed Strategy Nash Equilibrium.

Nash's proof consists of showing that the strategy space $\mathcal{A}$ and the best-reply correspondence BR satisfy all the conditions of Kakutani's Theorem:

- The set of mixed strategy profiles $\mathcal{A} = \times_{i \in N} \Delta(A_i)$ is itself non-empty, compact, and convex —being a cartesian product of such sets.

- The correspondence $BR(\alpha)$ always returns a non-empty, convex set. (This is because a player's best reply to a mixed strategy is either a pure strategy or any mix of the pure strategies that yield the same maximal expected payoff).

- The correspondence BR is upper hemi-continuous.

Since all conditions are met, Kakutani's theorem guarantees the existence of a fixed point, which proves that every finite game has at least one MSNE.

Chapter 3

Dominance and Iterative Elimination

3.1 Strictly Dominant and Dominated Strategies

Before analyzing Nash equilibria, we can sometimes simplify a game by identifying strategies that are unambiguously good or bad, regardless of what other players do.

Definition 3.1.1 (Strictly Dominant Strategy)**.** A pure strategy $a_i \in A_i$ is a **strictly dominant strategy** for player i if it gives a strictly higher payoff than any other strategy $a_i' \in A_i$, for any combination of strategies $a_{-i} \in A_{-i}$ played by the opponents. Formally, for all $a_i' \neq a_i$:

$$u_i(a_i, a_{-i}) > u_i(a_i', a_{-i}) \quad \text{for all } a_{-i} \in A_{-i}$$

Proposition 3.1.1. If a rational player has a strictly dominant strategy, they will always play it. If all players have strictly dominant strategies, the resulting strategy profile is the unique Nash Equilibrium of the game.

The flip side of a dominant strategy is a dominated one.

Definition 3.1.2 (Strictly Dominated Strategy)**.** A pure strategy $a_i \in A_i$ is **strictly dominated** if there exists another strategy $a_i' \in A_i$ (pure or mixed) that gives a strictly higher payoff than a_i, for any combination of strategies $a_{-i} \in A_{-i}$ played by the opponents. We say that a_i' strictly dominates a_i. Formally, there exists $\alpha_i' \in \Delta(A_i)$ such that:

$$u_i(\alpha_i', a_{-i}) > u_i(a_i, a_{-i}) \quad \text{for all } a_{-i} \in A_{-i}$$

A rational player should never play a strictly dominated strategy, because there is another strategy that will yield a better payoff no matter what others do. This implies that a strictly dominated strategy can never be a best reply to any belief about the opponents' actions.

Proposition 3.1.2. A strictly dominated strategy is not played with positive probability in any Nash Equilibrium.

Example 3.1.1 (Prisoners' Dilemma)**.** In the Prisoners' Dilemma, 'Defect' is a strictly dominant strategy for both players.

		Player 2	
		Cooperate	Defect
Player 1	Cooperate	3, 3	1, 4
	Defect	4, 1	2, 2

For Player 1, 'Cooperate' is strictly dominated by 'Defect' because $4 > 3$ (if P2 cooperates) and $2 > 1$ (if P2 defects). The same logic applies to Player 2. Rationality dictates that both players will choose 'Defect', leading to the unique NE of (Defect, Defect).

Example 3.1.2 (Domination by a Mixed Strategy)**.** Consider the following game.

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

For Player 1, no pure strategy strictly dominates another. For instance, U is not better than M if Player 2 plays R ($0 < 1$). However, strategy M is strictly dominated by a mixed strategy. Let's consider a mixed strategy that puts probability $p = 1/2$ on U and probability $1 - p = 1/2$ on D. The expected payoff from this mix is:

- If Player 2 plays L: $E[u_1] = \frac{1}{2}(3) + \frac{1}{2}(0) = 1.5$. This is strictly greater than the payoff of 1 from playing M.
- If Player 2 plays R: $E[u_1] = \frac{1}{2}(0) + \frac{1}{2}(3) = 1.5$. This is strictly greater than the payoff of 1 from playing M.

Since the mixed strategy (1/2 U $\oplus$ 1/2 D) gives a strictly higher payoff than the pure strategy M for any action Player 2 might take, M is a strictly dominated strategy and a rational Player 1 would never play it.

3.2 Never-Best Replies

The idea that dominated strategies are never played can be formalized by defining strategies that are never a best reply.

Definition 3.2.1 (Belief)**.** A **belief** for player i is a probability distribution $\mu_{-i} \in \Delta(A_{-i})$ over the strategy profiles of the other players. This belief represents player i's uncertainty about what the others will do.

Definition 3.2.2 (Never-Best Reply)**.** A pure strategy $a_i \in A_i$ is a **never-best reply** if it is not a best reply to any possible belief $\mu_{-i} \in \Delta(A_{-i})$. That is, for all μ_{-i}, there exists some other strategy $\alpha'_i \in \Delta(A_i)$ such that:

$$u_i(\alpha'_i, \mu_{-i}) > u_i(a_i, \mu_{-i})$$

Remark 3.2.1 (Equivalence)**.** In finite games, a pure strategy is strictly dominated if and only if it is a never-best reply. This is a foundational result connecting the concepts of dominance (comparing strategies directly) and best-response (comparing a strategy to a belief). For this equivalence to hold for more than two players, we must allow beliefs to be correlated across opponents.

3.3 Weakly Dominant and Dominated Strategies

The concept of dominance can be relaxed.

Definition 3.3.1 (Weakly Dominant Strategy)**.** A strategy $a_i \in A_i$ is a **weakly dominant strategy** for player i if it weakly dominates every other pure strategy $a'_i \in A_i$. That is, for all $a'_i \neq a_i$:

- $u_i(a_i, a_{-i}) \geq u_i(a'_i, a_{-i})$ for all $a_{-i} \in A_{-i}$, and
- for each a'_i, there exists at least one a_{-i} for which $u_i(a_i, a_{-i}) > u_i(a'_i, a_{-i})$.

Definition 3.3.2 (Weakly Dominated Strategy)**.** A strategy $a_i \in A_i$ is **weakly dominated** if there exists another strategy $\alpha_i' \in \Delta(A_i)$ that weakly dominates it. Formally, there exists $\alpha_i' \in \Delta(A_i)$ such that:

- $u_i(\alpha_i', a_{-i}) \geq u_i(a_i, a_{-i})$ for all $a_{-i} \in A_{-i}$, and
- there exists at least one a_{-i} for which $u_i(\alpha_i', a_{-i}) > u_i(a_i, a_{-i})$.

The connection between weak dominance and Nash equilibrium is more subtle than in the strict case. While it seems intuitive to avoid weakly dominated strategies, they can be part of a Nash Equilibrium.

Proposition 3.3.1. If every player plays a weakly dominant strategy, the resulting profile is a Nash Equilibrium. However, a Nash Equilibrium can involve players playing weakly dominated strategies.

Example 3.3.1. Consider the following game:

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

For Player 1, U is a weakly dominant strategy because it weakly dominates D ($1 > 0$ against L, $0 = 0$ against R). For Player 2, L is a weakly dominant strategy because it weakly dominates R. The strategy profile (U,L) consists of weakly dominant strategies and is a Nash Equilibrium. However, (D,R) is also a Nash Equilibrium. In this equilibrium, Player 1 plays D (which is weakly dominated by U) and Player 2 plays R (weakly dominated by L). Neither player wants to deviate, because given the other's action, they are indifferent.

3.4 Iterative Elimination of Dominated Strategies

3.4.1 Iterative Elimination of Strictly Dominated Strategies (IESDS)

Since rational players will not play strictly dominated strategies, we can remove them from the game. This might make other strategies newly dominated in the reduced game, which we can then remove, and so on.

Definition 3.4.1 (IESDS)**.** The process of **Iterative Elimination of Strictly Dominated Strategies** proceeds as follows:

1. Let $A_i^0 = A_i$ for all $i \in N$.
2. For $k \geq 1$, define A_i^k as the set of strategies in A_i^{k-1} that are not strictly dominated with respect to the strategy sets A_{-i}^{k-1} of the other players.
3. The process stops when no more strategies can be eliminated for any player. The set of surviving strategies is $A^\infty = \times_i A_i^\infty$.

Theorem 3.4.1. For finite games, the order of elimination in IESDS does not matter. The final set of surviving strategies is always the same.

Theorem 3.4.2. The set of strategy profiles that survive IESDS contains all Nash Equilibria of the original game. If IESDS results in a unique strategy profile for each player, that profile is the unique Nash Equilibrium.

Example 3.4.1 (A Multi-Step IESDS)**.** Consider the following 3x3 game.

Player 2

Player 1		L	C	R
	U	4, 3	5, 1	6, 2
	M	5, 7	8, 4	3, 6
	D	3, 0	9, 6	2, 8

Let's apply IESDS step-by-step:

- **Step 1:** We check for dominated strategies. For Player 2, strategy C is strictly dominated by R (since $2 > 1$, $6 > 4$, and $8 > 6$). A rational Player 2 will never play C. We eliminate column C.
- **Step 2:** The reduced game is now:

  Player 2

Player 1		L	R
	U	4, 3	6, 2
	M	5, 7	3, 6
	D	3, 0	2, 8

  Now we re-examine for dominated strategies. For Player 1, strategy D is strictly dominated by U (since $4 > 3$ and $6 > 2$). We eliminate row D.
- **Step 3:** The game is reduced further to:

  Player 2

Player 1		L	R
	U	4, 3	6, 2
	M	5, 7	3, 6

  In this 2x2 game, we check again. For Player 2, L strictly dominates R (since $3 > 2$ and $7 > 6$). We eliminate column R.
- **Step 4:** The final game is:

  Player 2

Player 1		L
	U	4, 3
	M	5, 7

  Player 1 must now simply choose the best response to L. The payoff for playing M (5) is greater than for U (4). Player 1 will choose M.

The process of IESDS leads to the unique prediction that the outcome of the game will be **(M, L)**. This is the unique Nash Equilibrium of the game.

3.4.2 Iterative Elimination of Weakly Dominated Strategies (IEWDS)

We can apply the same iterative process using weak dominance. However, this procedure is less reliable.

Definition 3.4.2 (IEWDS)**.** The process of **Iterative Elimination of Weakly Dominated Strategies** is analogous to IESDS, but at each step we remove all weakly dominated strategies for a player.

IEWDS has two significant issues:

1. The final set of surviving strategies can depend on the order in which strategies are eliminated.

2. The process can eliminate some Nash Equilibria of the original game.

Example 3.4.2 (Order Matters in IEWDS)**.** Consider the following game. This game has four pure strategy Nash equilibria: (U,L), (U,R), (M,L), and (D,R).

		Player 2	
		L	R
Player 1	U	2, 1	1, 1
	M	2, 1	0, 0
	D	1, 0	1, 1

Let's trace two different paths of elimination, both starting by removing one of Player 1's weakly dominated strategies.

- **Path A (Eliminate M first):** For Player 1, strategy U weakly dominates strategy M. (Against L, payoffs are equal, $2 = 2$. Against R, U is strictly better, $1 > 0$). Let's eliminate M. The game becomes:

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

  In this reduced game, for Player 2, strategy R now weakly dominates strategy L. (Against U, $1 = 1$. Against D, $1 > 0$). We eliminate L. Player 1 now chooses between U and D, knowing Player 2 will play R. They are indifferent. No more eliminations are possible. The surviving strategy profiles are **{(U,R), (D,R)}**.

- **Path B (Eliminate D first):** In the original game, for Player 1, strategy U also weakly dominates strategy D. (Against L, $2 > 1$. Against R, $1 = 1$). Let's eliminate D this time. The game becomes:

		Player 2	
		L	R
Player 1	U	2, 1	1, 1
	M	2, 1	0, 0

  In this reduced game, for Player 2, strategy L now weakly dominates strategy R. (Against U, $1 = 1$. Against M, $1 > 0$). We eliminate R. Player 1 now chooses between U and M, knowing Player 2 will play L. They are indifferent. No more eliminations are possible. The surviving strategy profiles are **{(U,L), (M,L)}**.

This example clearly shows that the order in which weakly dominated strategies are eliminated can lead to a different set of outcomes and can eliminate valid Nash equilibria from consideration.

Chapter 4

Monopoly and Imperfect Competition: Cournot and Bertrand

4.1 A Spectrum of Market Structures

In microeconomics, we analyze firm behavior across a spectrum of market structures, defined primarily by the number of firms and the degree of competition. The two endpoints of this spectrum are perfect competition and monopoly.

4.1.1 Perfect Competition

In a perfectly competitive market, there are many small firms, each selling an identical product. No single firm has the market power to influence the market price.

- **Price Takers:** Firms are price takers, meaning they accept the market price P as given.
- **Marginal Revenue:** For a price-taking firm, the revenue from selling one additional unit is simply the market price. Thus, Marginal Revenue (MR) equals Price (P).
- **Profit Maximization:** A firm maximizes its profit by producing the quantity Q where its Marginal Cost (MC) equals Marginal Revenue. The rule is therefore $\mathbf{P} = \mathbf{MC}$.
- **Efficiency:** Perfect competition leads to an efficient allocation of resources. The total economic surplus, which is the sum of Consumer Surplus (CS) and Producer Surplus (PS), is maximized. There is no deadweight loss.

4.1.2 Monopoly

At the other extreme is a monopoly, where a single firm is the sole producer of a good with no close substitutes. This firm has maximum market power.

- **Price Setter:** The monopolist faces the entire market demand curve and can choose the price (or quantity) that maximizes its profit.
- **Marginal Revenue:** Because the monopolist must lower the price on all units to sell an additional unit, its marginal revenue is less than the price ($MR < P$).
- **Profit Maximization:** The monopolist maximizes profit by producing the quantity Q_M where $\mathbf{MR} = \mathbf{MC}$. It then charges the highest price P_M that the market will bear for that quantity, as determined by the demand curve.

- **Inefficiency:** By restricting output and raising the price, the monopolist creates a deadweight loss (DWL), which represents a loss of total economic surplus compared to the competitive outcome.

Example 4.1.1 (Monopoly Pricing)**.** Let's analyze a monopolist facing a linear demand curve and constant marginal costs.

- **Demand:** $Q = 100 - 10P$.
- **Marginal Cost:** $MC = 4$ TL.

First, we find the inverse demand curve (price as a function of quantity): $10P = 100 - Q \implies P(Q) = 10 - 0.1Q$.

Total Revenue (TR) is $TR(Q) = P(Q) \times Q = (10 - 0.1Q)Q = 10Q - 0.1Q^2$. Marginal Revenue (MR) is the derivative of TR with respect to Q: $MR(Q) = \frac{d(TR)}{dQ} = 10 - 0.2Q$.

The monopolist sets $MR = MC$: $10 - 0.2Q = 4 \implies 6 = 0.2Q \implies Q_M = 30$.

The monopoly price is found by plugging Q_M into the inverse demand curve: $P_M = 10 - 0.1(30) = 10 - 3 = 7$ TL.

The monopolist's profit is $\pi_M = (P_M - MC) \times Q_M = (7 - 4) \times 30 = 90$ TL.

Remark 4.1.1 (Second-Order Condition)**.** Please note that, in all optimization problems in these notes, we find the optimal choice by taking the first derivative of the profit function and setting it to zero. This is the First-Order Condition (FOC). This procedure correctly identifies a unique maximum because the Second-Order Condition (SOC) is satisfied in each case. Specifically, the second derivative of each profit function with respect to the firm's choice variable (quantity or price) is negative. This implies that the profit functions are strictly concave, guaranteeing that the point where the FOC is met is a unique global maximum.

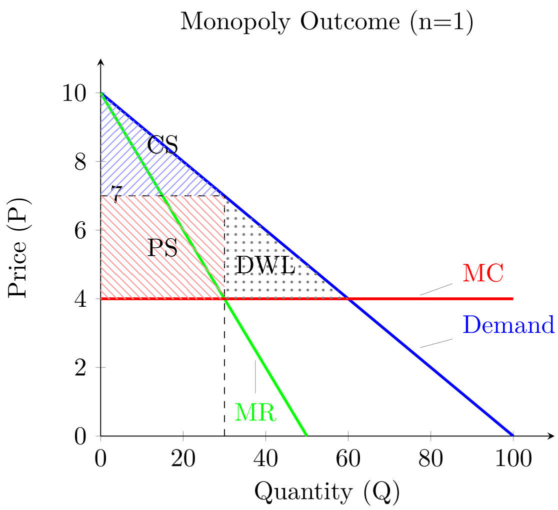

4.2 Cournot Duopoly Competition

What happens between the extremes of perfect competition and monopoly? One of the most famous models of oligopoly (a market with a few firms) is the Cournot model, where firms compete by simultaneously choosing their output quantities.

Remark 4.2.1 (Cournot Duopoly as a Normal Form Game)**.** The Cournot duopoly can be formally described as a normal form game $\langle N, (A_i)_{i \in N}, (\pi_i)_{i \in N} \rangle$:

- **Players:** $N = \{1, 2\}$.

- **Action Sets:** The action set for each firm i is the set of possible quantities it can produce. Since quantity must be non-negative, $A_i = [0, \infty)$. A strategy a_i is a choice of quantity $q_i \in A_i$.
- **Payoff Functions:** The payoff function for firm i is its profit function, $\pi_i(q_i, q_j)$, which depends on its own quantity and the quantity of its rival. For our example, this is:

$$\pi_i(q_i, q_j) = \underbrace{(10 - 0.1(q_i + q_j))}_{\text{Price}} \cdot q_i - \underbrace{4q_i}_{\text{Total Cost}}$$

Let's break our monopoly into two identical firms (a duopoly) facing the same market demand and costs.

- Firm 1 produces quantity q_1.
- Firm 2 produces quantity q_2.
- Total market quantity is $Q = q_1 + q_2$.
- The market price is determined by the total quantity: $P(Q) = 10 - 0.1(q_1 + q_2)$.
- Both firms have $MC = 4$.

Each firm's goal is to choose its quantity to maximize its own profit, taking the other firm's quantity as given. This is a classic application of Nash Equilibrium.

4.2.1 Firm 1's Best Reply Analysis (Profit Maximization)

Firm 1's profit, π_1, is:

$$\pi_1(q_1, q_2) = (P(q_1 + q_2) - MC) \times q_1 = (10 - 0.1(q_1 + q_2) - 4)q_1$$
$$\pi_1(q_1, q_2) = (6 - 0.1q_1 - 0.1q_2)q_1 = 6q_1 - 0.1q_1^2 - 0.1q_1q_2$$

To find the profit-maximizing q_1 for a given q_2, we take the partial derivative with respect to q_1 and set it to zero:

$$\frac{\partial \pi_1}{\partial q_1} = 6 - 0.2q_1 - 0.1q_2 = 0$$

Solving for q_1 gives Firm 1's **best-reply function** (or reaction function):

$$0.2q_1 = 6 - 0.1q_2 \implies BR_1^*(q_2) = 30 - 0.5q_2$$

4.2.2 Firm 2's Best Reply Analysis (Profit Maximization)

Since the firms are identical, Firm 2's problem is symmetric. Its best-reply function will be:

$$BR_2^*(q_1) = 30 - 0.5q_1$$

4.2.3 Nash-Cournot Equilibrium

The Nash-Cournot equilibrium is a pair of quantities (q_1^*, q_2^*) where each firm's quantity is a best reply to the other's. Formally, $q_1^* = BR_1(q_2^*)$ and $q_1^* = BR_2(q_1^*)$. We find this by solving the system of best-reply functions:

$$q_1^* = 30 - 0.5(30 - 0.5q_1^*)$$
$$q_1^* = 30 - 15 + 0.25q_1^*$$
$$0.75q_1^* = 15 \implies q_1^* = 20$$

Since the firms are symmetric, $q_2^* = 20$.

The Nash-Cournot equilibrium quantities are $q_1^* = 20$ and $q_2^* = 20$.

- **Total Quantity:** $Q_C = q_1^* + q_2^* = 40$.
- **Market Price:** $P_C = 10 - 0.1(40) = 6$ TL.
- **Firm Profit:** $\pi_1 = \pi_2 = (6 - 4) \times 20 = 40$ TL.
- **Total Profit:** $\Pi_C = 80$ TL.

Remark 4.2.2 (IESDS and Cournot Equilibrium)**.** For the Cournot model with linear demand, the Nash Equilibrium is the unique outcome that survives the Iterated Elimination of Strictly Dominated Strategies (IESDS). One can show that any quantity $q_i > 30$ (the monopoly quantity) is strictly dominated by $q_i = 30$. Once all such strategies are removed for both firms, the feasible range of actions shrinks. In the new, smaller game, some previously non-dominated strategies become dominated. This iterative process continues until only one strategy remains for each player: the Nash Equilibrium quantity, $q^* = 20$.

4.2.4 Comparison: Monopoly vs. Cournot Competition

Variable	Monopoly	Cournot Duopoly
Total Quantity (Q)	30	40
Market Price (P)	7 TL	6 TL
Total Profit (Π)	90 TL	80 TL

As we can see, when the monopoly is broken into a duopoly, the firms' competition leads to a higher total output and a lower market price. Consumers are better off, but total industry profits are lower.

Cournot Duopoly Outcome

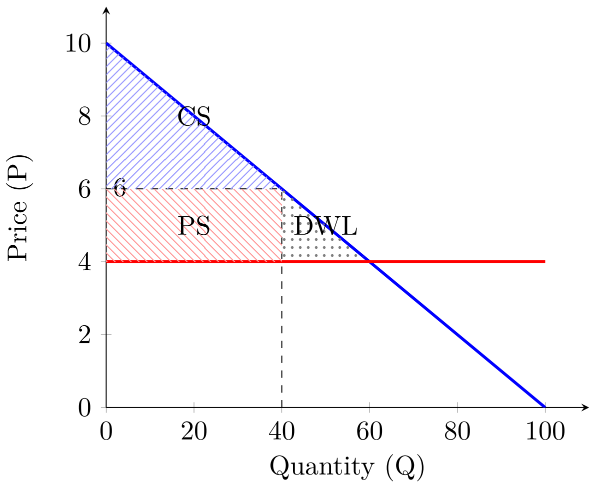

4.2.5 Generalizing the Cournot Model to n Firms

Let's now consider a market with n identical firms. Let q_i be the quantity produced by firm i. The total market quantity is $Q = \sum_{j=1}^{n} q_j$. The price is $P(Q) = 10 - 0.1Q = 10 - 0.1\sum_{j=1}^{n} q_j$.

Consider firm i's profit maximization problem. We can write $Q = q_i + Q_{-i}$, where $Q_{-i} = \sum_{j\neq i} q_j$ is the total quantity of all other firms. Firm i's profit is:

$$\pi_i = (P(Q) - MC)q_i = (10 - 0.1(q_i + Q_{-i}) - 4)q_i = 6q_i - 0.1q_i^2 - 0.1Q_{-i}q_i$$

The first-order condition for profit maximization is:

$$\frac{\partial \pi_i}{\partial q_i} = 6 - 0.2q_i - 0.1Q_{-i} = 0$$

This gives firm i's best-reply function: $q_i^*(Q_{-i}) = 30 - 0.5Q_{-i}$.

To find the symmetric Nash Equilibrium, we assume all firms produce the same quantity, $q_1 = q_2 = \cdots = q_n = q^*$. In this case, for firm i, the output of its rivals is $Q_{-i} = (n-1)q^*$. Substituting this into the best-reply function:

$$q^* = 30 - 0.5((n-1)q^*)$$

$$q^* = 30 - 0.5nq^* + 0.5q^*$$

$$0.5q^* + 0.5nq^* = 30 \implies q^*(0.5(1+n)) = 30$$

The individual firm's equilibrium quantity is:

$$q^* = \frac{60}{n+1}$$

The total market quantity is $Q^* = n \cdot q^* = \frac{60n}{n+1}$. The equilibrium market price is $P^* = 10 - 0.1Q^* = 10 - \frac{6n}{n+1} = \frac{10(n+1)-6n}{n+1} = \frac{4n+10}{n+1}$.

4.2.6 The Limit of Cournot Competition

What happens as the number of firms n gets very large?

$$\lim_{n\to\infty} Q^* = \lim_{n\to\infty} \frac{60n}{n+1} = 60$$

$$\lim_{n\to\infty} P^* = \lim_{n\to\infty} \frac{4n+10}{n+1} = 4$$

As the number of firms approaches infinity, the Cournot equilibrium converges to the perfectly competitive outcome. The market price equals the marginal cost ($P = MC = 4$), and the total quantity is the competitive quantity ($Q = 100 - 10(4) = 60$).

4.2.7 Summary: From Monopoly to Perfect Competition

Let's summarize the outcomes for different numbers of firms.

Variable	**Monopoly (n=1)**	**Duopoly (n=2)**	**Triopoly (n=3)**	**Perf. Comp. (n $\to \infty$)**
Indiv. Q (q^*)	30	20	15	$\to 0$
Total Q (Q^*)	30	40	45	$\to 60$
Price (P^*)	7 TL	6 TL	5.5 TL	$\to$ 4 TL
Indiv. Profit (π_i^*)	90	40	22.5	$\to 0$
Total Profit (Π^*)	90	80	67.5	$\to 0$

As more firms enter the market, competition drives the price down and total quantity up. Individual firm profits shrink, and the market outcome moves steadily from the monopoly point towards the perfectly competitive point.

4.3 Bertrand Duopoly Competition

An alternative model of duopoly, proposed by Joseph Bertrand in 1883, assumes that firms compete by setting prices, not quantities. Consumers observe the prices and buy from the firm with the lowest price. If prices are equal, consumers are split evenly between the two firms.

Remark 4.3.1 (Bertrand Duopoly as a Normal Form Game)**.** The Bertrand duopoly can also be formally described as a normal form game:

- **Players:** $N = \{1, 2\}$.

- **Action Sets:** The action set for each firm i is the set of possible prices it can set, $A_i = [0, \infty)$. A strategy a_i is a choice of price $p_i \in A_i$. Note that any price $p_i < MC$ or $p_i > P_{choke}$ (where demand is zero) is a weakly dominated strategy.
- **Payoff Functions:** The payoff function for firm i is its profit function, $\pi_i(p_i, p_j)$, which is discontinuous and depends on a comparison of the two prices:

$$\pi_i(p_i, p_j) = \begin{cases} (p_i - 4)(100 - 10p_i) & \text{if } p_i < p_j \\ \frac{1}{2}(p_i - 4)(100 - 10p_i) & \text{if } p_i = p_j \\ 0 & \text{if } p_i > p_j \end{cases}$$

This change in the strategic variable leads to a dramatically different outcome. Let's analyze the same market as before.

- Two identical firms, Firm 1 and Firm 2.
- Market demand: $Q = 100 - 10P$.
- Both firms have $MC = 4$. Recall also that the monopoly price for this market is $P_M = 7$.
- Firms simultaneously choose prices p_1 and p_2.

4.3.1 Finding the Nash-Bertrand Equilibrium

We look for a pair of prices (p_1^*, p_2^*) that constitutes a Nash Equilibrium. Let's analyze a firm's best response to its rival's price, p_j.

- **If $p_j > P_M = 7$:** Firm i's best response is to set the monopoly price, $p_i = 7$. It captures the entire market (since $p_i < p_j$) and earns the maximum possible profit ($\pi_M = 90$).
- **If $MC < p_j \leq P_M$ (i.e., $4 < p_j \leq 7$):** Firm i wants to set its price just below p_j. Technically, a best-reply function does not exist in this range, because for any undercut $p_i = p_j - \epsilon$, there is a better reply $p_i' = p_j - (\epsilon/2)$ that is also an undercut but yields a slightly higher profit. However, the *incentive* is always to undercut the rival.
- **If $p_j = MC = 4$:** Firm i's best response is to set any $p_i \geq 4$. If $p_i > 4$, it sells nothing and makes $\pi = 0$. If $p_i = 4$, it splits the market and makes $\pi = 0$. Thus, $p_i^* = 4$ is a best response.
- **If $p_j < MC = 4$:** Firm i would make a loss. Its best response is to set any $p_i > p_j$ (e.g., $p_i = 4$) to sell nothing and make $\pi = 0$.

Given this, no price above $MC = 4$ can be an equilibrium: if $p_1 = p_2 = p > 4$, either firm can deviate by undercutting. If $p_1 > p_2 > 4$, Firm 1 has 0 profit and can deviate by undercutting Firm 2. The only stable outcome is when both firms set their price equal to marginal cost.

The unique Nash-Bertrand Equilibrium is $p_1^* = p_2^* = MC = 4$.

4.3.2 Equilibrium Outcome

- **Market Price:** $P_B = 4$ TL.
- **Total Quantity:** $Q_B = 100 - 10(4) = 60$. This is the perfectly competitive quantity.
- **Firm Quantities:** Assuming they split the market, $q_1 = q_2 = 30$.
- **Firm Profit:** $\pi_1 = \pi_2 = (4 - 4) \times 30 = 0$ TL.

Remark 4.3.2. This result is striking. With just two firms competing on price, we get the perfectly competitive outcome: price equals marginal cost, and firms earn zero economic profit. This highlights the intensity of price competition and relies heavily on the assumptions that firms have identical costs, no capacity constraints, and sell a homogeneous product.

4.4 Final Comparison: Monopoly, Cournot, and Bertrand

The choice of strategic variable (quantity vs. price) dramatically alters the equilibrium outcome in an oligopoly.

Variable	**Monopoly**	**Cournot Duopoly**	**Bertrand Duopoly**	**Perfect Competition**
Strategic Variable	Price or Quantity	Quantity	Price	-
Total Quantity (Q)	30	40	60	60
Market Price (P)	7 TL	6 TL	4 TL	4 TL
Total Profit (Π)	90 TL	80 TL	0 TL	0 TL

This comparison shows that Cournot competition results in an outcome intermediate between monopoly and perfect competition. In contrast, Bertrand competition, under these specific assumptions, surprisingly collapses directly to the perfectly competitive outcome.

4.5 Bertrand Competition with Differentiated Products

The Bertrand-Nash equilibrium might seem like an extreme outcome. One of the ways to "soften" the intense price competition is to introduce product differentiation. When products are not perfect substitutes, a firm does not lose all of its customers by pricing slightly above its rival. This gives each firm a degree of market power.

Let's consider a duopoly where the firms' products are substitutes, but not perfect substitutes. We can model this with the following linear demand functions:

- Demand for Firm 1: $q_1 = 30 - p_1 + p_2$
- Demand for Firm 2: $q_2 = 30 - p_2 + p_1$
- Let's assume both firms have a constant marginal cost $MC = 2$.

Notice how the quantity demanded for Firm 1 decreases with its own price but increases with its rival's price.

Best Reply Functions: Each firm chooses its price to maximize profit. For Firm 1:

$$\pi_1(p_1, p_2) = (p_1 - MC)q_1 = (p_1 - 2)(30 - p_1 + p_2) = 32p_1 - p_1^2 + p_1p_2 - 60 - 2p_2$$

To find the profit-maximizing price, we take the derivative with respect to p_1 and set it to zero:

$$\frac{\partial \pi_1}{\partial p_1} = 32 - 2p_1 + p_2 = 0$$

Solving for p_1 gives Firm 1's best-reply function:

$$BR_1(p_2) = 16 + 0.5p_2$$

By symmetry, Firm 2's best-reply function is:

$$BR_2(p_1) = 16 + 0.5p_1$$

Nash-Bertrand Equilibrium with Differentiated Products: We find the equilibrium by solving the system of best-reply functions:

$$p_1^* = 16 + 0.5(16 + 0.5p_1^*)$$

$$p_1^* = 16 + 8 + 0.25p_1^*$$

$$0.75p_1^* = 24 \implies p_1^* = 32$$

Since the firms are symmetric, $p_2^* = 32$.

The equilibrium prices are $p_1^* = 32$ and $p_2^* = 32$.

- **Equilibrium Quantities:** $q_1^* = q_2^* = 30 - 32 + 32 = 30$.
- **Firm Profits:** $\pi_1^* = \pi_2^* = (p^* - MC) \times q^* = (32 - 2) \times 30 = 900$.

With differentiated products, firms are able to maintain prices above marginal cost and earn positive economic profits, resolving the extreme “paradox" of the basic Bertrand model.

Chapter 5

Extensive-Form Games

5.1 From Normal Form to Extensive Form

Thus far, we have primarily analyzed games in **normal form** (or strategic form). In this representation, we implicitly assume that players choose their strategies simultaneously, without knowing the actions of their opponents. The Cournot and Bertrand models are classic examples of this.

However, many strategic interactions are not simultaneous. One player moves, others observe that move, and then they react. The timing of actions and the flow of information are critically important.

To analyze such sequential interactions, we use the **extensive form** representation. This representation describes the game as a "game tree," capturing:

- The set of players.
- Who moves when (the order of play).
- What each player knows when they move.
- The actions available at each decision point.
- The payoffs for all players at the end of the game.

5.2 Formal Definition of an Extensive Form Game

We can formally define an extensive form game with perfect information (where all players know all past moves) as a tuple. This definition, following Osborne and Rubinstein (1994), is built around the concept of "histories."

Definition 5.2.1 (Extensive Form Game with Perfect Information)**.** An extensive form game is a tuple $\Gamma = \langle N, H, P, (u_i)_{i \in N} \rangle$ with the following components:

- $N = \{1, ..., n\}$**:** A finite set of players.
- H**:** A set of **histories** (finite sequences of actions). This set must satisfy two properties:
  1. The empty sequence, $\emptyset$ (representing the start of the game), is in H.
  2. H is **prefix-closed**: If a sequence $h = (a^1, a^2, ..., a^k)$ is in H, then any prefix of that sequence $h' = (a^1, ..., a^m)$ for $m < k$ is also in H.

  A history $h \in H$ is called **terminal** if it is not a prefix for any other history in H. We denote the set of all terminal histories as $Z \subset H$. All other histories $h \in H \setminus Z$ are non-terminal.

- P**:** The **player function**, $P : (H \setminus Z) \to N$. This function assigns a player $P(h)$ to every non-terminal history h. $P(h)$ is the player who moves next after history h has occurred.

- $(u_i)_{i \in N}$**:** A set of **payoff functions**, $u_i : Z \to \mathbb{R}$ for each player $i \in N$. This function assigns a real-valued payoff $u_i(h)$ to each player i for every terminal history $h \in Z$.

Remark 5.2.1 (Perfect Information)**.** This definition implicitly assumes **perfect information**. At any point a player is called upon to move, they know the exact history h that has led to that point. If a player did not know all past actions, we would need a more complex definition involving "information sets." We will cover that later.

5.3 Game Trees: Visualizing the Extensive Form

The formal definition is precise, but it is often easier to visualize an extensive form game using a **game tree**.

- **Nodes:** Each history $h \in H$ corresponds to a node in the tree.
- The root node is the empty history $\emptyset$.
- Non-terminal nodes $(H \setminus Z)$ are **decision nodes**.
- Terminal nodes (Z) are **leaf nodes**.
- **Branches:** The actions available to a player $P(h)$ at history h are represented by branches (edges) coming out of the node h.
- **Payoffs:** Payoffs are listed at the end of each terminal node, typically as a vector $(u_1, u_2, ..., u_n)$.

Example 5.3.1 (A Simple Entry Game)**.** Consider a game with two players: an Entrant (Player 1) and an Incumbent (Player 2).

1. Player 1 (Entrant) moves first. They can either "Enter" the market or "Stay Out."
2. If Player 1 stays out, the game ends.
3. If Player 1 enters, Player 2 (Incumbent) observes this and must choose whether to "Fight" (e.g., start a price war) or "Accommodate."

The payoffs (u_1, u_2) are as follows:

- (Stay Out): Player 1 gets 0, Player 2 gets 2 (monopoly profit).
- (Enter, Fight): Player 1 gets -1, Player 2 gets -1 (price war is costly).
- (Enter, Accommodate): Player 1 gets 1, Player 2 gets 1 (they share the market).

We can draw this as a game tree:

P1 (Entrant)
Enter ╱ ╲ Stay Out
P2 (Incumbent) (0, 2)
Fight ╱ ╲ Accommodate
(-1, -1) (1, 1)

5.4 Strategies in Extensive Form Games

A strategy in an extensive form game is more complex than in a normal form game. It must be a *complete contingent plan.*

Definition 5.4.1 (Strategy)**.** A strategy s_i for a player i is a **complete contingent plan**. Formally, it is a function that assigns a feasible action to **every** non-terminal history h at which that player is to move (i.e., for every $h \in H \setminus Z$ such that $P(h) = i$).

A **strategy profile** is a combination of strategies, one for each player, e.g., (Enter, Fight).

Example 5.4.1 (Actions vs. Strategies)**.** Let's add a move for Player 1 in the Entry Game.

1. Player 1 moves first, choosing L or R.
2. If P1 chooses R, the game ends with payoff (2, 2).
3. If P1 chooses L, P2 chooses A or B.
4. If P2 chooses A, the game ends with payoff (3, 1).
5. If P2 chooses B, P1 moves again, choosing X or Y.
6. (L, B, X) → (4, 0). (L, B, Y) → (1, 3).

```
            “P1”
           L/  \R
        “P2”    (2, 2)
       A/  \B
   (3, 1)   “P1”
           X/  \Y
       (4, 0)   (1, 3)
```

Now, let's list the *strategies* for each player.

- **Player 2:** P2 only moves at $h = (L)$. P2's actions are $\{A, B\}$. P2's strategy set is $S_2 = \{A, B\}$.
- **Player 1:** P1 moves at two different histories: $\emptyset$ and (L, B).
  - At $h = \emptyset$, P1's actions are $\{L, R\}$.
  - At $h = (L, B)$, P1's actions are $\{X, Y\}$.

  A strategy for P1 must specify an action at *both* nodes.

$$S_1 = \{(L, X), (L, Y), (R, X), (R, Y)\}$$

  (The first action is for $\emptyset$, the second is for (L, B)).

This highlights the difference: L is just an *action.* (L, X) is a *strategy.* Notice that the strategies (R, X) and (R, Y) must still specify what P1 would do at $h = (L, B)$, even though P1's own plan (to play R) means that node will never be reached. This is crucial for analyzing threats and equilibria.

5.5 Normal Form Representation and Nash Equilibrium

We can convert any extensive form game into its normal form representation by identifying the complete set of strategies for each player.

Definition 5.5.1 (Nash Equilibrium of an Extensive Form Game)**.** A **Nash Equilibrium** of an extensive form game is any Nash Equilibrium of its normal form representation.

As we will see, this solution concept is sometimes unsatisfying, as it allows for non-credible threats.

Example 5.5.1 (Normal Form of the Entry Game)**.** For our Simple Entry Game (Example 3.1):

- Player 1 (Entrant) moves only at $h = \emptyset$. Her strategy set is $S_1 = \{\text{Enter, Stay Out}\}$.
- Player 2 (Incumbent) moves only at $h = (\text{Enter})$. His strategy set is $S_2 = \{\text{Fight, Accommodate}\}$.

We can now represent this game as a 2×2 matrix. The payoffs are determined by the terminal history that results from each strategy profile.

- (Stay Out, Fight): P1 stays out. The game ends. Payoff is (0, 2).
- (Stay Out, Accommodate): P1 stays out. The game ends. Payoff is (0, 2).
- (Enter, Fight): P1 enters. P2's strategy is to fight. Payoff is (-1, -1).
- (Enter, Accommodate): P1 enters. P2's strategy is to accommodate. Payoff is (1, 1).

Normal Form of the Entry Game

	Player 2 (Incumbent)	
Player 1 (Entrant)	Fight	Accommodate
Enter	(-1, -1)	**(1, 1)**
Stay Out	**(0, 2)**	(0, 2)

By analyzing this normal form, we can find all pure-strategy Nash Equilibria (NE).

1. **(Enter, Accommodate):** If P1 enters, P2's best response is Accommodate ($1 > -1$). If P2 accommodates, P1's best response is Enter ($1 > 0$). This is a Nash Equilibrium.
2. **(Stay Out, Fight):** If P1 stays out, P2 is indifferent between Fight and Accommodate (payoff is 2 regardless). If P2 plays Fight, P1's best response is Stay Out ($0 > -1$). This is also a Nash Equilibrium.

Remark 5.5.1 (The Problem with NE in Sequential Games)**.** Both (Enter, Accommodate) and (Stay Out, Fight) are Nash Equilibria of the game. However, the NE **(Stay Out, Fight)** relies on a **non-credible threat**.

This equilibrium requires Player 2 to have the strategy "Fight". But if Player 1 were to deviate and actually "Enter", Player 2's decision node $h = (\text{Enter})$ would be reached. At that point, it would be irrational for Player 2 to "Fight" (payoff -1) when "Accommodate" gives a better payoff (1).

The simple Nash Equilibrium concept, applied to the normal form, ignores the sequential structure of the game and the credibility of threats. We need a stronger solution concept.

5.6 Subgame Perfect Nash Equilibrium (SPNE)

To solve this problem, we introduce the concept of a subgame and a refined equilibrium concept that requires strategies to be optimal in all parts of the game tree.

Definition 5.6.1 (Subgame)**.** A **subgame** of an extensive form game Γ is a part of the game that starts at a decision node h and includes all subsequent nodes and branches. (In games with perfect information, every decision node h defines the start of a new subgame).

In our Entry Game, there are two subgames:

1. The "whole game" (starting at P1's node, $h = \emptyset$).

2. The subgame starting at P2's decision node, $h =$ (Enter).

Definition 5.6.2 (Subgame Perfect Nash Equilibrium (SPNE))**.** A strategy profile s^* is a **Subgame Perfect Nash Equilibrium** if the strategy it specifies is a Nash Equilibrium in *every* subgame of the original game.

Remark 5.6.1 (SPNE and Non-Credible Threats)**.** Let's re-examine our two Nash Equilibria from the Entry Game:

1. **(Enter, Accommodate):**

   - Is it a NE in the whole game? Yes (we just showed this).
   - Is it a NE in the subgame at $h =$ (Enter)? Yes. In this subgame, P2's strategy is "Accommodate". P2's best response is to choose the action that gives the best payoff. "Accommodate" (1) is better than "Fight" (-1).

   This profile is an SPNE.

2. **(Stay Out, Fight):**

   - Is it a NE in the whole game? Yes (we just showed this).
   - Is it a NE in the subgame at $h =$ (Enter)? No. In this subgame, P2's strategy is "Fight", which gives payoff -1. P2 could deviate to "Accommodate" and get 1. Thus, "Fight" is not a NE in this subgame.

   This profile is **not** an SPNE.

The strategy "Fight" is a **non-credible threat**. P2 threatens to fight, but if P1 actually enters, it would be irrational for P2 to carry out that threat. SPNE eliminates non-credible threats by requiring all strategies to be optimal on every part of the game tree.

5.7 Backward Induction: The Algorithm for SPNE

For finite games with perfect information, the algorithm of **backward induction** is the procedure used to find all Subgame Perfect Nash Equilibria. The idea is simple: start at the end of the game and work your way back to the beginning.

5.7.1 The Algorithm

1. Start at the terminal nodes.
2. Move up to the "last" decision nodes (nodes whose branches all lead directly to terminal nodes).
3. For each such node, find the player $P(h)$'s optimal action by comparing the payoffs they receive.
4. "Prune" the tree: Assume the player will take this optimal action. Replace the decision node with the payoffs from that optimal action.
5. Repeat this process, moving up the tree until you reach the root node.

5.7.2 Solving the Entry Game via Backward Induction

Let's apply this to our Entry Game (Example 3.1).

1. **Step 1: P2's Decision.** We start at the last decision node, which belongs to Player 2 (Incumbent) at $h = (\text{Enter})$.
   - If P2 chooses "Fight," P2's payoff is -1.
   - If P2 chooses "Accommodate," P2's payoff is 1.

   Since $1 > -1$, a rational Player 2 will **always choose Accommodate**.
2. **Step 2: P1's Decision.** Now we move up to the first node ($h = \emptyset$), where Player 1 (Entrant) moves. Player 1 is rational and anticipates what Player 2 will do in Step 1.
   - If P1 chooses "Stay Out," P1's payoff is 0.
   - If P1 chooses "Enter," P1 knows that P2 *will* choose "Accommodate," leading to a payoff of 1 for P1.

   Since $1 > 0$, a rational Player 1 will **choose Enter**.

The Solution: The backward induction solution is the strategy profile **(Enter, Accommodate)**. This leads to the equilibrium path $h = (\text{Enter}, \text{Accommodate})$ and the payoffs $(1, 1)$. This is the unique SPNE of the game.

We can visualize this on the tree, with the chosen path marked.

"P1 (Entrant)"
Enter / \ Stay Out
"P2 (Incumbent)" (0, 2)
Fight / \ Accommodate
(-1, -1) (1, 1)

5.7.3 The One-Shot Deviation Principle

How do we check if a strategy profile s^* is an SPNE in a complex game? Do we have to check every possible alternative strategy in every possible subgame?

For finite horizon games, we can use a simpler method called the **One-Shot Deviation Principle**

Theorem 5.7.1 (One-Shot Deviation Principle)**.** In a finite extensive game with perfect information, a strategy profile s^* is a Subgame Perfect Nash Equilibrium if and only if no player i has a "one-shot deviation" that is profitable.

A one-shot deviation is a strategy s_i that differs from s_i^* at exactly one non-terminal history h, and agrees with s_i^* everywhere else.

Remark 5.7.1. This principle means we don't need to check complex deviations (e.g., "I'll do a instead of a^* at history h_1, and then b instead of b^* at history h_2"). We only need to check, for every single decision node h, if the player $P(h)$ can get a better payoff by *only* changing their action at h and then reverting to their original s^* plan for all future decisions.

This is the logic that backward induction implicitly uses. At each node, it assumes the player will play optimally "at that node", given the optimal play that will follow in all subsequent subgames. This principle is extremely useful for verifying SPNEs.

5.8 A Critique of SPNE: The Centipede Game

The following famous finite-horizon game illustrates the power and sometimes counter-intuitive results of subgame perfect equilibrium.

Example 5.8.1 (The Centipede Game (Rosenthal, 1981))**.** Two players, 1 and 2, take turns deciding whether to **Take** (or "Stop") a growing pot of money, or **Pass** (or "Continue") to the other player. The game ends when one player Takes or after a fixed number of rounds (e.g., 6).

- If a player Takes, they get a larger share of the current pot, but the other player gets a smaller share.
- If a player Passes, the total pot grows, and the other player gets to move.

Game Tree for (a relatively short) Centipede Game:

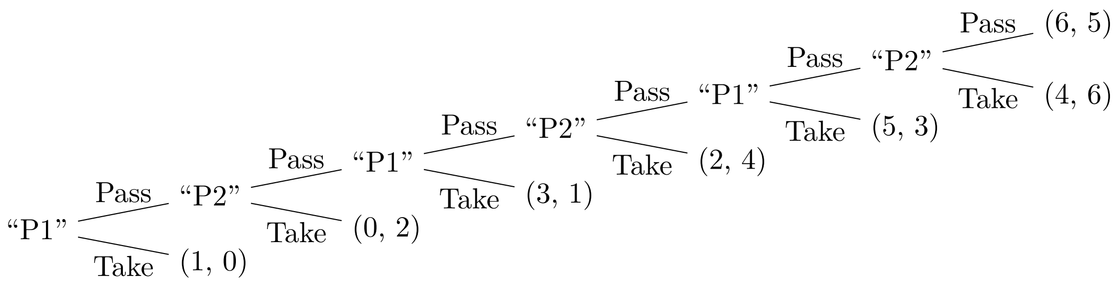

SPNE Analysis (Backward Induction):

1. **Round 6 (P2):** P2 chooses between Take (payoff 6) and Pass (payoff 5). P2 will **Take**.
2. **Round 5 (P1):** P1 chooses between Take (payoff 5) and Pass (knowing P2 will Take, payoff 4). P1 will **Take**.
3. **Round 4 (P2):** P2 chooses between Take (payoff 4) and Pass (knowing P1 will Take, payoff 3). P2 will **Take**.
4. **...By Induction:** At every node, the player whose turn it is will **Take**.

The SPNE Strategy Profile is for *both players to play "Take" at every decision node they have.*

The Equilibrium Path is simply "Take". The game ends immediately, and payoffs are $(\mathbf{1}, \mathbf{0})$.

Remark 5.8.1 (The Centipede Paradox)**.** This result is called a "paradox" because it is deeply counter-intuitive. Both players would be much better off if they could just "Pass" a few times. The path (Pass, Pass, Pass, Pass, Pass, Take) yields (4, 6), which is vastly better for both. However, the logic of backward induction is ruthless. The only credible action at the end is to Take, which makes the only credible action at the second-to-last node to Take, and this logic unravels all the way to the beginning.

5.9 Application: Stackelberg Competition (Leader-Follower Model)

We can now use our tools for extensive form games to analyze a sequential oligopoly model. The **Stackelberg Duopoly Model** is a classic "Leader-Follower" game. It provides a powerful contrast to the simultaneous-move Cournot model we analyzed in our previous lecture.

5.9.1 The Stackelberg Game

Let's use the same market parameters from our Cournot example:

- **Inverse Demand:** $P(Q) = 10 - 0.1Q$, where $Q = q_1 + q_2$.
- **Costs:** $MC_1 = MC_2 = 4$ TL/unit.

The key difference is the timing. This is an extensive form game with perfect information:

1. **Stage 1:** Firm 1 (the "Leader") chooses its quantity q_1.
2. **Stage 2:** Firm 2 (the "Follower") *observes* q_1 and then chooses its quantity q_2.

Because this is a sequential game, we must find the **Subgame Perfect Nash Equilibrium (SPNE)** using **Backward Induction**.

5.9.2 Solving for the SPNE

5.9.2.1 Step 1: The Follower's Problem (Find s_2^*)

We start at the end of the game. Firm 2 (the Follower) has observed q_1 and must choose its optimal q_2. Firm 2's strategy, s_2, will be a function of q_1.

Firm 2's profit is:

$$\pi_2(q_1, q_2) = (P(Q) - MC) \cdot q_2 = (10 - 0.1(q_1 + q_2) - 4)q_2$$

$$\pi_2(q_1, q_2) = (6 - 0.1q_1 - 0.1q_2)q_2 = 6q_2 - 0.1q_1q_2 - 0.1q_2^2$$

To maximize π_2 given q_1, Firm 2 takes the FOC with respect to q_2:

$$\frac{\partial \pi_2}{\partial q_2} = 6 - 0.1q_1 - 0.2q_2 = 0$$

Solving for q_2 gives Firm 2's **Best-Reply Function (or Reaction Function)**:

$$0.2q_2 = 6 - 0.1q_1 \implies q_2^*(q_1) = 30 - 0.5q_1$$

This function is Firm 2's SPNE strategy, $s_2^* = q_2^*(q_1)$. It is a complete contingent plan specifying Firm 2's optimal action for *every possible* history $h = (q_1)$.

5.9.2.2 Step 2: The Leader's Problem (Find s_1^*)

Firm 1 (the Leader) moves first at $h = \emptyset$. It knows that whatever quantity q_1 it chooses, Firm 2 will respond according to the plan s_2^*. Therefore, Firm 1 can substitute s_2^* directly into its own profit function.

Firm 1's profit is:

$$\pi_1(q_1, q_2) = (6 - 0.1q_1 - 0.1q_2)q_1$$

Substitute $q_2 = 30 - 0.5q_1$:

$$\pi_1(q_1) = (6 - 0.1q_1 - 0.1(30 - 0.5q_1))q_1$$

$$\pi_1(q_1) = (6 - 0.1q_1 - 3 + 0.05q_1)q_1$$

$$\pi_1(q_1) = (3 - 0.05q_1)q_1 = 3q_1 - 0.05q_1^2$$

Firm 1 chooses q_1 to maximize this profit. We take the FOC with respect to q_1:

$$\frac{d\pi_1}{dq_1} = 3 - 0.1q_1 = 0 \implies \mathbf{q_1^S = 30}$$

5.9.2.3 Step 3: The SPNE Outcome

The Leader's strategy is $s_1^* = 30$. The Follower's strategy is $s_2^* = 30 - 0.5q_1$. The equilibrium *path* (or outcome) is found by plugging q_1^S into s_2^*:

- **Leader's Quantity:** $q_1^S = 30$
- **Follower's Quantity:** $q_2^S = 30 - 0.5(30) = 15$

Now we find the market outcomes:

- **Total Quantity:** $Q_S = q_1^S + q_2^S = 30 + 15 = 45$
- **Market Price:** $P_S = 10 - 0.1(45) = 5.5$ TL
- **Profits:**
  - π_1^S (Leader) $= (P_S - MC)q_1^S = (5.5 - 4) \times 30 = \mathbf{45}$ TL
  - π_2^S (Follower) $= (P_S - MC)q_2^S = (5.5 - 4) \times 15 = \mathbf{22.5}$ TL
- **Total Producer Surplus (PS):** $\Pi_S = 45 + 22.5 = 67.5$ TL

5.9.3 Comparison: Stackelberg vs. Cournot

This sequential setup gives a very different outcome from the simultaneous-move Cournot game. In the Cournot game, we found $q_1^C = 20, q_2^C = 20$, $P_C = 6$, and $\pi_1^C = \pi_2^C = 40$.

- **First-Mover Advantage:** The Leader, by committing to a large quantity (30) first, forces the Follower to retreat (produce 15 instead of 20). This gives the Leader a higher profit (45 vs. 40) and the Follower a lower profit (22.5 vs. 40).
- **Market Outcome:** The Stackelberg outcome has a *higher* total quantity ($Q_S = 45 > Q_C = 40$) and a *lower* price ($P_S = 5.5 < P_C = 6$).

5.9.4 Welfare Analysis

Let's compare the total welfare generated by the Cournot and Stackelberg models.

1. **Perfect Competition (Benchmark for Efficiency):**
   - $P_{PC} = MC = 4$.
   - $Q_{PC} = 100 - 10(4) = 60$.
   - $PS_{PC} = (P_{PC} - MC) \times Q_{PC} = (4 - 4) \times 60 = 0$.
   - $CS_{PC} = 0.5 \times (10 - 4) \times 60 = 180$.
   - **Max Economic Surplus** $(ES_{PC}) = 180 + 0 = 180$.
2. **Cournot (Simultaneous):**
   - $P_C = 6$, $Q_C = 40$.
   - $PS_C = (P_C - MC) \times Q_C = (6 - 4) \times 40 = 80$.
   - $CS_C = 0.5 \times (10 - 6) \times 40 = 80$.
   - **Economic Surplus** $(ES_C) = 80 + 80 = 160$.
   - **Deadweight Loss** $(DWL_C) = 180 - 160 = 20$.
3. **Stackelberg (Sequential):**
   - $P_S = 5.5$, $Q_S = 45$.
   - $PS_S = (P_S - MC) \times Q_S = (5.5 - 4) \times 45 = 67.5$.
   - $CS_S = 0.5 \times (10 - 5.5) \times 45 = 0.5 \times 4.5 \times 45 = 101.25$.
   - **Economic Surplus** $(ES_S) = 67.5 + 101.25 = 168.75$.
   - **Deadweight Loss** $(DWL_S) = 180 - 168.75 = 11.25$.

Welfare Comparison Summary

Welfare Component	Cournot (Simul.)	Stackelberg (Seq.)	Perfect Comp.
Price (P^*)	6.00	5.50	4.00
Total Quantity (Q^*)	40	45	60
Consumer Surplus (CS)	80.00	101.25	180.00
Producer Surplus (PS)	80.00	67.50	0.00
Economic Surplus (ES)	**160.00**	**168.75**	**180.00**
Deadweight Loss (DWL)	20.00	11.25	0.00

Remark 5.9.1 (Conclusion on Welfare)**.** The Stackelberg model, by virtue of its sequential structure and first-mover advantage, leads to a more aggressive output level (and lower price) than the simultaneous-move Cournot model.

This has two important welfare effects:

1. **Consumers are better off** $(CS_S > CS_C)$.
2. **Total Economic Surplus is higher** $(ES_S > ES_C)$, and the Deadweight Loss is smaller.

Interestingly, the strategic advantage of the first mover benefits consumers and improves overall efficiency, even though it harms the rival firm. The total loss in producer surplus ($80 \to 67.5$) is more than offset by the gain in consumer surplus ($80 \to 101.25$).

Chapter 6

Bayesian Games

6.1 Introduction to Incomplete Information

In our study of Nash Equilibrium, we made a strong assumption: all players know the rules of the game, and in particular, they know all other players' payoff functions. This is the assumption of **complete information**.

In many real-world strategic interactions, this assumption is violated. Players may have private information about their own costs, valuations, or preferences that other players do not know.

- A firm in a duopoly may know its own marginal cost, but not the marginal cost of its rival.
- A bidder in an auction knows their own valuation for an item, but not the valuations of other bidders.
- A country in a diplomatic negotiation may know its own "red lines" (true willingness to fight), but not those of its counterpart.

Games with **incomplete information**, also known as **Bayesian games**, are designed to model these situations. The key insight, developed by John Harsanyi, is to model incomplete information by introducing **types** for each player. A player's type, t_i, is their private information.

We assume that all players know the possible types of their opponents. In Harsanyi's original formulation, there is a **common prior** probability distribution over the set of type profiles. An equivalent "reduced form" representation, which we will use, defines each player's beliefs about their opponents' types, conditional on their own type.

6.2 Formal Definition of a Bayesian Game

We can now formally define a normal form Bayesian game.

Definition 6.2.1 (Bayesian Game)**.** A **Bayesian game** is a tuple $\langle N, (A_i)_{i \in N}, (T_i)_{i \in N}, (\mu_i)_{i \in N}, (u_i)_{i \in N} \rangle$, where:

- $N = \{1, 2, \ldots, n\}$ is a finite set of players.
- $A = A_1 \times A_2 \times \cdots \times A_n$ is the set of action profiles, where A_i is the set of actions available to player i.
- $T = T_1 \times T_2 \times \cdots \times T_n$ is the set of type profiles, where T_i is the set of possible types for player i. Player i's type, $t_i \in T_i$, is known only to player i.

- $\mu_i : T_i \to \Delta(T_{-i})$ is the belief function for player i. For each $t_i \in T_i$, $\mu_i(\cdot \mid t_i)$ is a probability distribution over the set of other players' types, T_{-i}. We let $\mu_i(t_{-i} \mid t_i)$ denote the probability that i assigns to the opponents' type profile being t_{-i}, given i's own type is t_i.
- $u = (u_1, \ldots, u_n)$ is a set of payoff functions, where $u_i : A \times T \to \mathbb{R}$. The payoff for player i depends on the actions taken by all players and the true type profile of all players.

(When beliefs $(\mu_i)_{i \in N}$ are all derived from a single common prior P via Bayes' rule, we say the game has a *common prior*.)

Example 6.2.1 (Battle of the Sexes with Two-Sided Incomplete Information)**.** Let's consider a variation of the Battle of the Sexes game.

- **Players:** $N = \{\text{He (1), She (2)}\}$.
- **Actions:** $A_1 = A_2 = \{\text{Football (F), Ballet (B)}\}$.
- **Types:**
  - He can be a type who loves her, $t_1 =$ Romeo (R), or a type who hates her, $t_1 =$ Montague (M). $T_1 = \{R, M\}$.
  - She can be a type who loves him, $t_2 =$ Juliet (J), or a type who hates him, $t_2 =$ Capulet (C). $T_2 = \{J, C\}$.
- **Beliefs:** The beliefs are derived from a common prior over a state space $\Omega = \{ll, lh, hl, hh\}$, where the first letter is He's state (l=love, h=hate) and the second is She's (l=love, h=hate).
  - The states map to types as:
    - $(R, J) \iff ll$
    - $(R, C) \iff lh$
    - $(M, J) \iff hl$
    - $(M, C) \iff hh$
  - The common prior probabilities are:
    - $\mathbb{P}(ll) = 1/3$
    - $\mathbb{P}(lh) = 1/3$
    - $\mathbb{P}(hl) = 1/6$
    - $\mathbb{P}(hh) = 1/6$
- **Derived Beliefs:** We use Bayes' rule to find the players' beliefs.
  - **He (Player 1):**
    - $\mathbb{P}(t_1 = R) = \mathbb{P}(ll) + \mathbb{P}(lh) = 1/3 + 1/3 = 2/3$.
    - $\mathbb{P}(t_1 = M) = \mathbb{P}(hl) + \mathbb{P}(hh) = 1/6 + 1/6 = 1/3$.
    - $\mu_1(J \mid R) = \frac{\mathbb{P}(ll)}{\mathbb{P}(R)} = \frac{1/3}{2/3} = 1/2$.
    - $\mu_1(C \mid R) = \frac{\mathbb{P}(lh)}{\mathbb{P}(R)} = \frac{1/3}{2/3} = 1/2$.
    - $\mu_1(J \mid M) = \frac{\mathbb{P}(hl)}{\mathbb{P}(M)} = \frac{1/6}{1/3} = 1/2$.
    - $\mu_1(C \mid M) = \frac{\mathbb{P}(hh)}{\mathbb{P}(M)} = \frac{1/6}{1/3} = 1/2$.
    - (He's belief about Her is independent of his type: 50% J, 50% C).
  - **She (Player 2):**
    - $\mathbb{P}(t_2 = J) = \mathbb{P}(ll) + \mathbb{P}(hl) = 1/3 + 1/6 = 1/2$.

- $\mathbb{P}(t_2 = C) = \mathbb{P}(lh) + \mathbb{P}(hh) = 1/3 + 1/6 = 1/2$.
- $\mu_2(R \mid J) = \frac{\mathbb{P}(ll)}{\mathbb{P}(J)} = \frac{1/3}{1/2} = 2/3$.
- $\mu_2(M \mid J) = \frac{\mathbb{P}(hl)}{\mathbb{P}(J)} = \frac{1/6}{1/2} = 1/3$.
- $\mu_2(R \mid C) = \frac{\mathbb{P}(lh)}{\mathbb{P}(C)} = \frac{1/3}{1/2} = 2/3$.
- $\mu_2(M \mid C) = \frac{\mathbb{P}(hh)}{\mathbb{P}(C)} = \frac{1/6}{1/2} = 1/3$.
- (She's belief about Him is independent of her type: 2/3 R, 1/3 M).

- **Payoffs:** Payoffs depend on the state (type-profile).
  - "Loves" types (Romeo, Juliet) want to coordinate (F,F or B,B).
  - "Hates" types (Montague, Capulet) want to anti-coordinate (F,B or B,F).
  - Payoff of 3 for preferred outcome, 1 for less-preferred, 0 for worst.

This defines four "states of the world," which we can represent in a 2×2 matrix of payoff tables:

		Juliet (J) F	**Juliet (J)** B	**Capulet (C)** F	**Capulet (C)** B
Romeo (R)	F	**3, 1**	0, 0	**3, 0**	0, 3
	B	0, 0	**1, 3**	0, 1	**1, 0**
Montague (M)	F	0, 1	**3, 0**	0, 0	**3, 3**
	B	**1, 0**	0, 3	**1, 1**	0, 0

6.3 Bayesian Nash Equilibrium

How do players make decisions in such a game? A player knows their own type but not their opponent's. They must choose an action that is optimal *in expectation*, given their beliefs about the opponent's type and the opponent's strategy.

Definition 6.3.1 (Strategy in a Bayesian Game)**.** A **(pure) strategy** for player i in a Bayesian game is a function $\sigma_i : T_i \to A_i$ that specifies an action $a_i = \sigma_i(t_i)$ for *each* possible type $t_i \in T_i$ that player i might have.

A strategy profile is a set of such strategies, $\sigma = (\sigma_1(\cdot), \sigma_2(\cdot), \ldots, \sigma_n(\cdot))$.

Definition 6.3.2 (Bayesian Nash Equilibrium)**.** A pure-strategy **Bayesian Nash Equilibrium (BNE)** is a strategy profile $\sigma^* = (\sigma_1^*(\cdot), \ldots, \sigma_n^*(\cdot))$ such that for every player $i \in N$ and for every type $t_i \in T_i$, the action $a_i^* = \sigma_i^*(t_i)$ maximizes player i's expected payoff, given their beliefs $\mu_i(\cdot \mid t_i)$ and the strategies of the other players $\sigma_{-i}^*(\cdot)$.

Formally, for all $i \in N$ and all $t_i \in T_i$, $\sigma_i^*(t_i)$ must satisfy:

$$\sigma_i^*(t_i) \in \arg\max_{a_i \in A_i} \sum_{t_{-i} \in T_{-i}} \mu_i(t_{-i} \mid t_i) \cdot u_i((a_i, \sigma_{-i}^*(t_{-i})), (t_i, t_{-i}))$$

Equivalently, for all $i \in N$, all $t_i \in T_i$, and all $a_i \in A_i$:

$$\sum_{t_{-i} \in T_{-i}} \mu_i(t_{-i} \mid t_i) \cdot u_i((\sigma_i^*(t_i), \sigma_{-i}^*(t_{-i})), (t_i, t_{-i})) \geq \sum_{t_{-i} \in T_{-i}} \mu_i(t_{-i} \mid t_i) \cdot u_i((a_i, \sigma_{-i}^*(t_{-i})), (t_i, t_{-i}))$$

Remark 6.3.1. In simple terms, a BNE is a set of "contingency plans" (one for each type) such that no player, upon learning their type, has an incentive to deviate from their plan, assuming no other player deviates from theirs.

Example 6.3.1 (Deriving a BNE for the BoS)**.** Let's find all pure-strategy BNEs for our BoS example. A strategy for He is a pair of actions $\sigma_1 = (\sigma_1(R), \sigma_1(M))$. A strategy for She is a pair of actions $\sigma_2 = (\sigma_2(J), \sigma_2(C))$.

She has 4 pure strategies: $\sigma_2 \in \{FF, FB, BF, BB\}$, where FB means $\sigma_2(J) = F$ and $\sigma_2(C) = B$. He has 4 pure strategies: $\sigma_1 \in \{FF, FB, BF, BB\}$, where FB means $\sigma_1(R) = F$ and $\sigma_1(M) = B$.

We will first derive the best-reply correspondence for each type of each player.

Part 1: Best-Reply Analysis

1. **He (Player 1), Type Romeo (R):**

- Romeo's beliefs: $\mu_1(J \mid R) = 1/2$, $\mu_1(C \mid R) = 1/2$.
- His expected payoff $E[u_R(a)|\sigma_2] = \frac{1}{2}u_R(a, \sigma_2(J)) + \frac{1}{2}u_R(a, \sigma_2(C))$.
- **Table for** $BR_R(\sigma_2)$**:**

Best replies of (He, Romeo)

σ_2 :	**FF**	**FB**	**BF**	**BB**
$\mathbf{E[u_R(F)]}$	$\frac{1}{2}(3) + \frac{1}{2}(3) = 3$	$\frac{1}{2}(3) + \frac{1}{2}(0) = 1.5$	$\frac{1}{2}(0) + \frac{1}{2}(3) = 1.5$	$\frac{1}{2}(0) + \frac{1}{2}(0) = 0$
$\mathbf{E[u_R(B)]}$	$\frac{1}{2}(0) + \frac{1}{2}(0) = 0$	$\frac{1}{2}(0) + \frac{1}{2}(1) = 0.5$	$\frac{1}{2}(1) + \frac{1}{2}(0) = 0.5$	$\frac{1}{2}(1) + \frac{1}{2}(1) = 1$
$\mathbf{BR_R}$	F	F	F	B

2. **He (Player 1), Type Montague (M):**

- Montague's beliefs: $\mu_1(J \mid M) = 1/2$, $\mu_1(C \mid M) = 1/2$.
- His expected payoff $E[u_M(a)|\sigma_2] = \frac{1}{2}u_M(a, \sigma_2(J)) + \frac{1}{2}u_M(a, \sigma_2(C))$.
- **Table for** $BR_M(\sigma_2)$**:**

Best replies of (He, Montague)

σ_2 :	**FF**	**FB**	**BF**	**BB**
$\mathbf{E[u_M(F)]}$	$\frac{1}{2}(0) + \frac{1}{2}(0) = 0$	$\frac{1}{2}(0) + \frac{1}{2}(3) = 1.5$	$\frac{1}{2}(3) + \frac{1}{2}(0) = 1.5$	$\frac{1}{2}(3) + \frac{1}{2}(3) = 3$
$\mathbf{E[u_M(B)]}$	$\frac{1}{2}(1) + \frac{1}{2}(1) = 1$	$\frac{1}{2}(1) + \frac{1}{2}(0) = 0.5$	$\frac{1}{2}(0) + \frac{1}{2}(1) = 0.5$	$\frac{1}{2}(0) + \frac{1}{2}(0) = 0$
$\mathbf{BR_M}$	B	F	F	F

3. **She (Player 2), Type Juliet (J):**

- Juliet's beliefs: $\mu_2(R \mid J) = 2/3$, $\mu_2(M \mid J) = 1/3$.
- Her expected payoff $E[u_J(a)|\sigma_1] = \frac{2}{3}u_J(\sigma_1(R), a) + \frac{1}{3}u_J(\sigma_1(M), a)$.
- **Table for** $BR_J(\sigma_1)$**:**

4. **She (Player 2), Type Capulet (C):**

- Capulet's beliefs: $\mu_2(R \mid C) = 2/3$, $\mu_2(M \mid C) = 1/3$.
- Her expected payoff $E[u_C(a)|\sigma_1] = \frac{2}{3}u_C(\sigma_1(R), a) + \frac{1}{3}u_C(\sigma_1(M), a)$.
- **Table for** $BR_C(\sigma_1)$**:**

Summary of the Best Reply Analysis

Best replies of (She,Juliet)

σ_1 :	**FF**	**FB**	**BF**	**BB**
$\mathbf{E[u_J(F)]}$	$\frac{2}{3}(1)+\frac{1}{3}(1)=1$	$\frac{2}{3}(1)+\frac{1}{3}(0)=\frac{2}{3}$	$\frac{2}{3}(0)+\frac{1}{3}(1)=\frac{1}{3}$	$\frac{2}{3}(0)+\frac{1}{3}(0)=0$
$\mathbf{E[u_J(B)]}$	$\frac{2}{3}(0)+\frac{1}{3}(0)=0$	$\frac{2}{3}(0)+\frac{1}{3}(3)=1$	$\frac{2}{3}(3)+\frac{1}{3}(0)=2$	$\frac{2}{3}(3)+\frac{1}{3}(3)=3$
$\mathbf{BR_J}$	F	B	B	B

Best replies of (She,Capulet)

σ_1 :	**FF**	**FB**	**BF**	**BB**
$\mathbf{E[u_C(F)]}$	$\frac{2}{3}(0)+\frac{1}{3}(0)=0$	$\frac{2}{3}(0)+\frac{1}{3}(1)=\frac{1}{3}$	$\frac{2}{3}(1)+\frac{1}{3}(0)=\frac{2}{3}$	$\frac{2}{3}(1)+\frac{1}{3}(1)=1$
$\mathbf{E[u_C(B)]}$	$\frac{2}{3}(3)+\frac{1}{3}(3)=3$	$\frac{2}{3}(3)+\frac{1}{3}(0)=2$	$\frac{2}{3}(0)+\frac{1}{3}(3)=1$	$\frac{2}{3}(0)+\frac{1}{3}(0)=0$
$\mathbf{BR_C}$	B	B	B	F

Part 2: Finding Pure-Strategy BNEs

A BNE is a strategy profile $\sigma^* = (\sigma_1^*, \sigma_2^*)$ that is a mutual best reply: $\sigma_1^* \in BR_1(\sigma_2^*)$ and $\sigma_2^* \in BR_2(\sigma_1^*)$. We can find the BNEs by checking each of Player 1's strategies.

- **Case 1:** $\sigma_1^* = (F, F)$ **(i.e.,** FF**).**
  - We find She's best reply: $\sigma_2 = BR_2(FF) = (BR_J(FF), BR_C(FF))$.
  - From our tables, $BR_J(FF) = F$ and $BR_C(FF) = B$. So, $\sigma_2 = (F, B)$.
  - Now, we check if $\sigma_1^* = (F, F)$ is a best reply to $\sigma_2 = (F, B)$.
  - Is $(F, F) \in BR_1(FB)$?
  - From our tables, $BR_1(FB) = (BR_R(FB), BR_M(FB)) = (F, F)$.
  - Yes! Since $(F, F) \in BR_1(FB)$ and $(F, B) \in BR_2(FF)$, we have a BNE.
  - **BNE 1:** $\sigma^* = ((\sigma_1(R) = F, \sigma_1(M) = F), (\sigma_2(J) = F, \sigma_2(C) = B))$**.**
- **Case 2:** $\sigma_1^* = (F, B)$ **(i.e.,** FB**).**
  - She's best reply: $\sigma_2 = BR_2(FB) = (BR_J(FB), BR_C(FB)) = (B, B)$.
  - Now, we check if $\sigma_1^* = (F, B)$ is a best reply to $\sigma_2 = (B, B)$.
  - Is $(F, B) \in BR_1(BB)$? From our tables, $BR_1(BB) = (B, F)$.
  - No, $(F, B) \neq (B, F)$.
  - No BNE in this case.
- **Case 3:** $\sigma_1^* = (B, F)$ **(i.e.,** BF**).**
  - She's best reply: $\sigma_2 = BR_2(BF) = (BR_J(BF), BR_C(BF)) = (B, B)$.
  - Now, we check if $\sigma_1^* = (B, F)$ is a best reply to $\sigma_2 = (B, B)$.
  - Is $(B, F) \in BR_1(BB)$?
  - From our tables, $BR_1(BB) = (BR_R(BB), BR_M(BB)) = (B, F)$.
  - Yes! Since $(B, F) \in BR_1(BB)$ and $(B, B) \in BR_2(BF)$, we have a BNE.
  - **BNE 2:** $\sigma^* = ((\sigma_1(R) = B, \sigma_1(M) = F), (\sigma_2(J) = B, \sigma_2(C) = B))$**.**
- **Case 4:** $\sigma_1^* = (B, B)$ **(i.e.,** BB**).**
  - She's best reply: $\sigma_2 = BR_2(BB) = (BR_J(BB), BR_C(BB)) = (B, F)$.
  - Now, we check if $\sigma_1^* = (B, B)$ is a best reply to $\sigma_2 = (B, F)$.
  - Is $(B, B) \in BR_1(BF)$?

He (Romeo)

σ_2 :	FF	FB	BF	BB
$E[u_R(F)]$	3	1.5	1.5	0
$E[u_R(B)]$	0	0.5	0.5	1
BR_R	F	F	F	B

He (Montague)

σ_2 :	FF	FB	BF	BB
$E[u_M(F)]$	0	1.5	1.5	3
$E[u_M(B)]$	1	0.5	0.5	0
BR_M	B	F	F	F

She (Juliet)

σ_1 :	FF	FB	BF	BB
$E[u_J(F)]$	1	2/3	1/3	0
$E[u_J(B)]$	0	1	2	3
BR_J	F	B	B	B

She (Capulet)

σ_1 :	FF	FB	BF	BB
$E[u_C(F)]$	0	1/3	2/3	1
$E[u_C(B)]$	3	2	1	0
BR_C	B	B	B	F

- From our tables, $BR_1(BF) = (F, F)$.
- No, $(B, B) \neq (F, F)$. No BNE in this case.

Conclusion: After checking all 4 possible strategies for Player 1, we find there are **two** pure-strategy Bayesian Nash Equilibria:

1. $\sigma^* = ((\sigma_1(R) = F, \sigma_1(M) = F), (\sigma_2(J) = F, \sigma_2(C) = B))$
2. $\sigma^* = ((\sigma_1(R) = B, \sigma_1(M) = F), (\sigma_2(J) = B, \sigma_2(C) = B))$

6.4 Application: Cournot Duopoly with Incomplete Information

Let's analyze a setting with a continuous action space.

Example 6.4.1 (Bayesian Cournot)**.**

- **Players:** $N = \{\text{Firm 1, Firm 2}\}$.
- **Actions:** $q_1, q_2 \in [0, \infty)$ are the quantities produced.
- **Demand:** Inverse demand is $P(Q) = 10 - Q/10$, where $Q = q_1 + q_2$.
- **Types:**
  - Firm 1 has one type. Its marginal cost is $c_1 = 4$. This is common knowledge.
  - Firm 2 has two types: "Low Cost" ($t_2 = c_L = 2$) or "High Cost" ($t_2 = c_H = 6$). $T_2 = \{c_L, c_H\}$.
- **Beliefs:** Firm 1 believes $\mu_1(c_L \mid c_1) = 1/2$ and $\mu_1(c_H \mid c_1) = 1/2$. Firm 2 knows its own cost.

Finding the BNE: A BNE is a strategy profile $(q_1^*, \sigma_2^*(t_2))$ where $\sigma_2^*(t_2)$ is a pair of quantities (q_{2L}^*, q_{2H}^*).

1. Firm 2's Problem: Firm 2 knows its type and knows Firm 1 will play q_1^*. It solves a standard maximization problem for each type.

- **If Low Cost ($c_L = 2$):** $\pi_{2L} = [P(q_1^* + q_{2L}) - c_L]q_{2L} = [(10 - (q_1^* + q_{2L})/10) - 2]q_{2L} = [8 - q_1^*/10 - q_{2L}/10]q_{2L}$

$$\text{FOC: } \frac{\partial \pi_{2L}}{\partial q_{2L}} = 8 - q_1^*/10 - 2q_{2L}/10 = 0$$

$$2q_{2L} = 80 - q_1^* \implies q_{2L}^* = 40 - \frac{1}{2}q_1^* \quad \text{(BR for } c_L \text{ type)}$$

- **If High Cost ($c_H = 6$):** $\pi_{2H} = [P(q_1^* + q_{2H}) - c_H]q_{2H} = [(10 - (q_1^* + q_{2H})/10) - 6]q_{2H} = [4 - q_1^*/10 - q_{2H}/10]q_{2H}$

$$\text{FOC: } \frac{\partial \pi_{2H}}{\partial q_{2H}} = 4 - q_1^*/10 - 2q_{2H}/10 = 0$$

$$2q_{2H} = 40 - q_1^* \implies q_{2H}^* = 20 - \frac{1}{2}q_1^* \quad \text{(BR for } c_H \text{ type)}$$

As expected, the low-cost type produces more than the high-cost type.

2. Firm 1's Problem: Firm 1 does not know Firm 2's type. It must choose q_1 to maximize its *expected* profit, given its beliefs and Firm 2's strategy (q_{2L}^*, q_{2H}^*).

$$E[\pi_1] = \frac{1}{2} \cdot \pi_1(q_1, q_{2L}^*) + \frac{1}{2} \cdot \pi_1(q_1, q_{2H}^*)$$

$$E[\pi_1] = \frac{1}{2} \cdot [(10 - (q_1 + q_{2L}^*)/10) - 4]q_1 + \frac{1}{2} \cdot [(10 - (q_1 + q_{2H}^*)/10) - 4]q_1$$

$$E[\pi_1] = \frac{1}{2} \cdot [6 - q_1/10 - q_{2L}^*/10]q_1 + \frac{1}{2} \cdot [6 - q_1/10 - q_{2H}^*/10]q_1$$

This can be simplified:

$$E[\pi_1] = (6 - q_1/10)q_1 - \frac{q_1}{10}[\frac{1}{2}q_{2L}^* + \frac{1}{2}q_{2H}^*]$$

Let $E[q_2^*] = \frac{1}{2}q_{2L}^* + \frac{1}{2}q_{2H}^*$.

$$\text{FOC: } \frac{\partial E[\pi_1]}{\partial q_1} = 6 - 2q_1/10 - E[q_2^*]/10 = 0$$

$$2q_1 = 60 - E[q_2^*] \implies q_1^* = 30 - \frac{1}{2}E[q_2^*] \quad \text{(BR for Firm 1)}$$

Firm 1 best-responds to the *expected quantity* of its rival.

3. Solving the System: We have a system of 3 linear equations for $(q_1^*, q_{2L}^*, q_{2H}^*)$:

1. $q_1^* = 30 - \frac{1}{2}(\frac{1}{2}q_{2L}^* + \frac{1}{2}q_{2H}^*) = 30 - \frac{1}{4}(q_{2L}^* + q_{2H}^*)$
2. $q_{2L}^* = 40 - \frac{1}{2}q_1^*$
3. $q_{2H}^* = 20 - \frac{1}{2}q_1^*$

Substitute (2) and (3) into (1):

$$q_1^* = 30 - \frac{1}{4}\left[(40 - \frac{1}{2}q_1^*) + (20 - \frac{1}{2}q_1^*)\right]$$

$$q_1^* = 30 - \frac{1}{4}[60 - q_1^*]$$

$$q_1^* = 30 - 15 + \frac{1}{4}q_1^*$$

$$\frac{3}{4}q_1^* = 15$$

$$\mathbf{q_1^* = 20}$$

Now we can find Firm 2's equilibrium quantities:

$$\mathbf{q_{2L}^*} = 40 - \frac{1}{2}(20) = 40 - 10 = \mathbf{30}$$

$$\mathbf{q_{2H}^*} = 20 - \frac{1}{2}(20) = 20 - 10 = \mathbf{10}$$

The unique pure-strategy BNE is $(q_1^* = 20, (\sigma_2(c_L) = 30, \sigma_2(c_H) = 10))$.

6.5 Application: First-Price Sealed-Bid Auction

Bayesian games are the primary tool used to analyze auctions.

Example 6.5.1 (First-Price Auction).

- **Players:** n bidders, $N = \{1, \ldots, n\}$.
- **Types:** Each bidder i has a private valuation $v_i \in [0, 1]$. $T_i = [0, 1]$.
- **Beliefs:** Each v_i is independently drawn from a uniform distribution, $v_i \sim U[0, 1]$.
- **Actions:** Each bidder i submits a sealed bid $b_i \geq 0$. $A_i = [0, \infty)$.
- **Payoffs:** The highest bidder wins the object and pays their bid. All other bidders pay 0.

$$u_i(b_i, b_{-i}, v_i) = \begin{cases} v_i - b_i & \text{if } b_i > b_j \text{ for all } j \neq i \\ 0 & \text{otherwise} \end{cases}$$

(We assume ties are broken randomly, e.g., probability of winning is 0).

Finding the BNE: We look for a symmetric BNE where each player uses the same bidding function $\beta : [0, 1] \to \mathbb{R}^+$, such that $b_i = \beta(v_i)$. We will guess that $\beta(v)$ is strictly increasing.

Suppose player i has valuation v_i. They choose a bid b_i to maximize their expected payoff, assuming all other $n - 1$ players use the strategy $\beta(v_j)$.

$$E[\pi_i(b_i, v_i)] = \mathbb{P}(\text{Win with } b_i) \cdot (v_i - b_i)$$

For i to win, b_i must be greater than all other bids: $b_i > \beta(v_j)$ for all $j \neq i$. Since $\beta(\cdot)$ is increasing, we can invert it: $\beta^{-1}(b_i) > v_j$. Let $y = \beta^{-1}(b_i)$. This y is the "type" that player i is "pretending" to be with their bid b_i.

The probability of winning is $\mathbb{P}(v_j < y)$ for all $j \neq i$. Since $v_j \sim U[0, 1]$, $\mathbb{P}(v_j < y) = y$. Given independence, $\mathbb{P}(\text{Win}) = \prod_{j \neq i} \mathbb{P}(v_j < y) = y^{n-1}$.

So, player i (with type v_i) chooses y (which implies bid $b_i = \beta(y)$) to maximize:

$$E[\pi_i] = y^{n-1} \cdot (v_i - \beta(y))$$

To find the optimal y, we take the FOC w.r.t y:

$$\frac{\partial E[\pi_i]}{\partial y} = (n - 1)y^{n-2}(v_i - \beta(y)) + y^{n-1}(-\beta'(y)) = 0$$

In a symmetric equilibrium, the optimal choice for type v_i is to bid "truthfully" (in terms of strategy), so this FOC must be satisfied at $y = v_i$.

$$(n - 1)v_i^{n-2}(v_i - \beta(v_i)) - v_i^{n-1}\beta'(v_i) = 0$$

Dividing by v_i^{n-2} (for $v_i > 0$):

$$(n - 1)(v_i - \beta(v_i)) - v_i\beta'(v_i) = 0$$

$$(n - 1)v_i - (n - 1)\beta(v_i) = v_i\beta'(v_i)$$

This is a first-order differential equation. We can solve it, but it is easier to guess the linear solution $\beta(v) = kv$ and solve for k. If $\beta(v) = kv$, then $\beta'(v) = k$.

$$(n - 1)v_i - (n - 1)(kv_i) = v_i(k)$$

Divide by v_i:

$$(n - 1) - (n - 1)k = k$$

$$n - 1 = k + (n - 1)k = k(1 + n - 1) = kn$$

$$k = \frac{n - 1}{n}$$

We also need to check the boundary condition $\beta(0) = 0$, which kv satisfies.

Thus, the unique symmetric BNE is $\beta(v_i) = \frac{n-1}{n}v_i$.

Remark 6.5.1. This is a famous result. In a first-price auction with n bidders and $U[0,1]$ valuations, the equilibrium strategy is to bid a fraction $\frac{n-1}{n}$ of your true value.

- If $n = 2$, you bid $\frac{1}{2}v_i$ (you shade your bid by 50%).
- As $n \to \infty$, $\frac{n-1}{n} \to 1$. Bids converge to the true value, as the competitive pressure from many bidders eliminates all profits.

6.6 Application: Second-Price Sealed-Bid Auction

Also known as the Vickrey Auction, this format is central to auction theory because of its elegant equilibrium properties.

Example 6.6.1 (Second-Price Auction).

- **Players:** n bidders, $N = \{1, \dots, n\}$.
- **Types:** Each bidder i has a private valuation $v_i \in [0,1]$. $T_i = [0,1]$.
- **Beliefs:** Valuations are independent (Independent Private Values - IPV). The distribution does not need to be uniform; let it be any strictly increasing F.
- **Actions:** Each bidder i submits a sealed bid $b_i \geq 0$.
- **Rules:** The highest bidder wins but pays the **second-highest** bid.
- **Payoffs:** Let $\bar{b}_{-i} = \max_{j \neq i} b_j$ be the highest bid among i's opponents.

$$u_i(b_i, b_{-i}, v_i) = \begin{cases} v_i - \bar{b}_{-i} & \text{if } b_i > \bar{b}_{-i} \\ 0 & \text{if } b_i < \bar{b}_{-i} \end{cases}$$

(Ties can be broken randomly, payoff is 0 if winning bid equals second highest).

Finding the Symmetric BNE: We will show that bidding your true valuation, $\beta(v_i) = v_i$, is a symmetric Bayesian Nash Equilibrium. In fact, we will prove a stronger result: it is a **weakly dominant strategy**.

Proposition 6.6.1. In a second-price sealed-bid auction with private values, bidding $b_i = v_i$ is a **weakly dominant strategy** for every player i.

Proof. Consider player i with valuation v_i. Let $\bar{b}_{-i} = \max_{j \neq i} b_j$ be the highest bid submitted by opponents. Player i does not know $\bar{b}_{-i}$, but we can analyze payoffs for any realization of $\bar{b}_{-i}$.

Compare bidding truth $b_i = v_i$ versus deviating to some other bid b'.

Case 1: Underbidding ($b' < v_i$). There are five subcases depending on where $\bar{b}_{-i}$ falls relative to the bids:

- If $\bar{b}_{-i} < b' < v_i$: You win with b' and pay $\bar{b}_{-i}$. You would also win with v_i and pay $\bar{b}_{-i}$. Payoff is $v_i - \bar{b}_{-i}$ in both cases.
- If $b' < \bar{b}_{-i} < v_i$:
  - With truthful bid v_i: You win and pay $\bar{b}_{-i}$. Payoff is $v_i - \bar{b}_{-i} > 0$.
  - With deviation b': You lose. Payoff is 0.

  Truth-telling is strictly better.
- If $b' < v_i < \bar{b}_{-i}$: You lose with b' (payoff 0). You also lose with v_i (payoff 0). Same outcome.

Deviation	Case ($\bar{b}_{-i}$ location)	Outcome (v_i)	Payoff (v_i)	Outcome (b')	Payoff (b')	Comparison
Underbidding: $b' < v_i$						
1	$\bar{b}_{-i} < b'$	Win	$v_i - \bar{b}_{-i}$	Win	$v_i - \bar{b}_{-i}$	Same
2	$\bar{b}_{-i} = b'$	Win	$v_i - b'$	Tie	$\frac{1}{2}(v_i - b')$	v_i is better
3	$b' < \bar{b}_{-i} < v_i$	Win	$v_i - \bar{b}_{-i}$	Lose	0	v_i is better
4	$\bar{b}_{-i} = v_i$	Tie	0	Lose	0	Same
5	$\bar{b}_{-i} > v_i$	Lose	0	Lose	0	Same
Overbidding: $b' > v_i$						
6	$\bar{b}_{-i} < v_i$	Win	$v_i - \bar{b}_{-i}$	Win	$v_i - \bar{b}_{-i}$	Same
7	$\bar{b}_{-i} = v_i$	Tie	0	Win	$v_i - \bar{b}_{-i} = 0$	Same
8	$v_i < \bar{b}_{-i} < b'$	Lose	0	Win	$v_i - \bar{b}_{-i} < 0$	v_i is better
9	$\bar{b}_{-i} = b'$	Lose	0	Tie	$\frac{1}{2}(v_i - b') < 0$	v_i is better
10	$\bar{b}_{-i} > b'$	Lose	0	Lose	0	Same

Table 6.1: Comparison of payoffs: Truthful bidding (v_i) vs. Deviating (b')

- **Equality Case ($\bar{b}_{-i} = b'$):** With b', you tie (payoff 0 or share win). With v_i, you win strictly and pay $\bar{b}_{-i} = b' < v_i$, yielding positive profit. Truth-telling is better.
- **Equality Case ($\bar{b}_{-i} = v_i$):** With b', you lose. With v_i, you tie and pay $\bar{b}_{-i} = v_i$, yielding 0 profit. Same outcome.

Result: Underbidding never increases your payoff, and in some cases (specifically when $b' \leq \bar{b}_{-i} < v_i$), it lowers it.

Case 2: Overbidding ($b' > v_i$).

- If $\bar{b}_{-i} < v_i < b'$: You win with b' and pay $\bar{b}_{-i}$. You also win with v_i and pay $\bar{b}_{-i}$. Payoff is $v_i - \bar{b}_{-i}$ in both cases.
- If $v_i < \bar{b}_{-i} < b'$:
  - With truthful bid v_i: You lose. Payoff is 0.
  - With deviation b': You win, but you pay $\bar{b}_{-i}$. Since $\bar{b}_{-i} > v_i$, your payoff is $v_i - \bar{b}_{-i} < 0$.

  Truth-telling is strictly better (avoids loss).
- If $v_i < b' < \bar{b}_{-i}$: You lose with b' and v_i. Payoff is 0 in both cases.
- **Equality Case ($\bar{b}_{-i} = v_i$):** With v_i, you tie (payoff 0). With b', you win and pay $\bar{b}_{-i} = v_i$, yielding 0 profit. Same outcome.
- **Equality Case ($\bar{b}_{-i} = b'$):** With v_i, you lose. With b', you tie. If you win the tie, you pay $b' > v_i$, yielding negative profit. Truth-telling is better.

Result: Overbidding never increases your payoff, and in some cases, it reduces it.

Conclusion: For any realization of opponents' bids $\bar{b}_{-i}$, we have $u_i(v_i, \bar{b}_{-i}) \geq u_i(b', \bar{b}_{-i})$. Thus, bidding $b_i = v_i$ is a weakly dominant strategy.

Since it is optimal regardless of what others do, it is certainly optimal given what others are *expected* to do. Therefore, everyone bidding their true valuation is a Bayesian Nash Equilibrium. □

Chapter 7

Adverse Selection

7.1 Introduction to Asymmetric Information

In our previous discussions (e.g., standard Nash or Bayesian games), we assumed players knew the payoff functions of their opponents, or at least had common-prior beliefs over a known set of types. We now turn to a broader class of problems under the umbrella of **Asymmetric Information**.

Asymmetric information exists when one party to a transaction has more or better information than the other. This information advantage can lead to severe market inefficiencies, from market collapse to suboptimal contracts.

There are two primary categories of information asymmetry, distinguished by *when* the information is hidden:

1. **Adverse Selection (Hidden Information):** This is a *pre-contractual* problem. One party has private information about their "type" (e.g., their costs, their health status, the quality of their product) *before* the game or transaction begins. The uninformed party must design a mechanism or contract to "screen" these types or risk ending up with the "worst" types.

   - **Examples:** Used car sellers (know car quality), insurance buyers (know health risk), job applicants (know their ability).

2. **Moral Hazard (Hidden Action):** This is a *post-contractual* problem. After a contract is signed, one party takes an action that is unobservable to the other (e.g., effort level, driving care). The uninformed party must design a contract that incentivizes the desired (but unobservable) action.

   - **Examples:** An employee's effort level, an insured person's risk-taking behavior, a CEO's strategic decisions.

Today, we focus entirely on **Adverse Selection**. We will motivate this as an extensive form game with incomplete information where an uninformed player must move first (a "screening" game) or second (a "lemons" market).

7.2 The Canonical Model: Akerlof's "Market for Lemons" (1970)

The most famous example of adverse selection is George Akerlof's Nobel-winning model of the used car market. It demonstrates how asymmetric information can destroy a market, even when large gains from trade exist.

Example 7.2.1 (The Market for Lemons).
- **Setup:** A market for used cars.

- **Sellers (Informed):** Each seller knows the quality θ of their car. θ is distributed uniformly on $[0, 100]$. $\theta \sim U[0, 100]$.
- **Buyers (Uninformed):** Buyers do not know θ; they only know the distribution.
- **Valuations:**
  - Seller's (owner's) value: $V_S(\theta) = \theta$.
  - Buyer's value: $V_B(\theta) = 1.5 \cdot \theta$.
  - We assume $V_B(\theta) > V_S(\theta)$ for all $\theta > 0$. This means **gains from trade exist for every single car**. In a full-information world, all cars would trade at a price $p \in [\theta, 1.5\theta]$, and the outcome would be Pareto efficient.

Analysis with Asymmetric Information:

Buyers are rational and must form an expectation about the quality of cars available on the market at a given price p.

1. **Seller's Decision:** A seller of a car with quality θ will only sell if the market price p is greater than or equal to their value.
$$\text{Sell if } p \geq V_S(\theta) \implies p \geq \theta$$
2. **The "Pool" of Cars:** At any given price p, the only sellers who remain in the market are those with cars of quality $\theta \leq p$.
3. **Buyer's Belief:** A rational buyer anticipates this. If the market price is p, the buyer knows the quality of any car offered for sale is in the range $[0, p]$. Given the uniform distribution, the buyer's expectation of the quality is:
$$\mathbb{E}[\theta \mid \theta \leq p] = \frac{p - 0}{2} = \frac{p}{2}$$
4. **Buyer's Decision:** The buyer will only purchase a car at price p if their expected value exceeds the price:
$$V_B(\mathbb{E}[\theta \mid \theta \leq p]) \geq p$$
$$1.5 \cdot \mathbb{E}[\theta \mid \theta \leq p] \geq p$$
Substituting our result from step 3:
$$1.5 \cdot \left(\frac{p}{2}\right) \geq p$$
$$0.75p \geq p$$

This inequality, $0.75p \geq p$, is **only true if** $p \leq 0$.

Proposition 7.2.1 (Market Unraveling)**.** The only possible equilibrium in this market is at $p = 0$, where no transactions take place. The market for used cars collapses entirely.

Remark 7.2.1. This is a stunningly inefficient outcome. We have a market with 100% potential gains from trade, yet **no trade occurs**. The presence of "lemons" (low-θ cars) drives the buyers' expectations down so far that the "plums" (high-θ cars) can no longer be sold. The high-quality sellers are driven from the market, which lowers the average quality, which drives more sellers out, and so on, until only $p = 0$ is sustainable. This is the core market failure of adverse selection.

7.3 Screening: The Principal-Agent Model

Akerlof's model is a "market" model. An alternative is when the *uninformed* party (the "Principal") has market power and can design the trading mechanism. This is called a **screening game**.

The Principal's goal is to design a "menu of contracts" that induces the informed party (the "Agent") to self-select, revealing their type through their choice.

7.3.1 The General Framework (Extensive Form)

We can model this as an extensive form game with incomplete information:

1. **Nature's Move:** Nature chooses the Agent's type $\theta \in \Theta$ with a known probability distribution $\mathbb{P}(\theta)$. The Agent learns their type θ, but the Principal does not.

2. **Principal's Move:** The Principal (uninformed) offers a *menu of contracts*, $M = \{(q_i, t_i)\}_{i \in I}$, where q is a "quantity" (e.g., quality, insurance coverage, work task) and t is a "transfer" (e.g., price, premium, wage).

3. **Agent's Move:** The Agent (informed) observes the menu M and chooses one contract $(q_i, t_i) \in M$, or chooses the outside option (rejection), which gives reservation utility $\bar{u}(\theta)$.

4. **Payoffs:** Payoffs $u_P(q, t, \theta)$ and $u_A(q, t, \theta)$ are realized.

Proposition 7.3.1 (The Revelation Principle)**.** The Principal's problem is to design the optimal menu M. The **Revelation Principle** is a powerful tool that states we can, without loss of generality, restrict our attention to menus where:

1. There is one contract for each type: $M = \{(q_\theta, t_\theta)\}_{\theta \in \Theta}$.

2. The menu is "incentive compatible," meaning each type θ voluntarily chooses the contract (q_θ, t_θ) designed for them.

This simplifies the problem to finding the set of contracts $\{(q_\theta, t_\theta)\}_{\theta \in \Theta}$ that maximizes the Principal's expected utility, subject to two sets of constraints.

Definition 7.3.1 (The Principal's Problem)**.** For a discrete set of types $\Theta = \{\theta_1, \dots, \theta_n\}$, the Principal solves:

$$\max_{\{q_i, t_i\}_{i=1}^n} \sum_{i=1}^{n} \mathbb{P}(\theta_i) \cdot u_P(q_i, t_i, \theta_i)$$

Subject to:

1. **Individual Rationality (IR) (also known as Participation Constraints, PC):** Each type must prefer their contract to their outside option.

$$u_A(q_i, t_i, \theta_i) \geq \bar{u}(\theta_i) \quad \text{for all } i \in \{1, \dots, n\}$$

*Note: We will refer to these constraints as **IR** throughout these notes.*

2. **Incentive Compatibility Constraints (ICC):** Each type θ_i must prefer their "own" bundle (q_i, t_i) to any "other" bundle (q_j, t_j).

$$u_A(q_i, t_i, \theta_i) \geq u_A(q_j, t_j, \theta_i) \quad \text{for all } i, j \in \{1, \dots, n\}$$

7.4 Application: Quality Screening by a Monopolist

This is a classic application of the screening model, explaining how a monopolist (Principal) sells different "qualities" of a product to buyers (Agents) who have different, private, valuations. This is a form of second-degree price discrimination studied first Mussa and Rosen (1978).

Example 7.4.1 (Turkish Airlines: Istanbul to Los Angeles)**.** • **Principal:** Turkish Airlines (THY), a monopolist for direct IST-LAX flights.

- **Cost:** THY's cost of providing "quality" q (e.g., legroom, service, lounge access) is $C(q) = \frac{1}{2}q^2$.
- **Agents:** 500 potential consumers. They have two types, "Economy" and "Business."
  - 400 are **Economy type** ($\theta_L = 5$).
  - 100 are **Business type** ($\theta_H = 10$). Business types value quality more.
- **Probabilities:**
  - $\mathbb{P}(\theta_H) = \nu = 100/500 = 0.2$
  - $\mathbb{P}(\theta_L) = 1 - \nu = 400/500 = 0.8$
- **Utility:** A consumer of type θ who buys quality q at price p gets:

  $$u_A(q, p, \theta) = \theta q - p$$

  The reservation utility $\bar{u} = 0$ for both types.
- **Single-Crossing Property:** $\frac{\partial u_A}{\partial q} = \theta$. Since $\theta_H > \theta_L$, the indifference curves cross once; Business types value quality more at the margin.
- **Problem:** THY offers a menu of two "ticket" contracts, $\{(q_L, p_L), (q_H, p_H)\}$, to maximize expected profit.

  $$\max_{q_L, q_H, p_L, p_H} \quad 0.2\left(p_H - \frac{1}{2}q_H^2\right) + 0.8\left(p_L - \frac{1}{2}q_L^2\right)$$

First-Best Analysis (Full Information Benchmark)

Before solving the screening problem, let's analyze the **first-best** outcome. This is what THY would do if it *knew* each consumer's type (i.e., if "Business" and "Economy" were written on their foreheads).

Since THY is a monopolist, it would make a separate take-it-or-leave-it offer to each type, designed to extract all of their surplus. This means THY maximizes profit from each type separately, subject only to that type's Individual Rationality (IR) constraint.

1. **For Business types ($\theta_H = 10$):** THY solves: $\max_{q_H, p_H} \pi_H = p_H - \frac{1}{2}q_H^2$
   Subject to (IR-H): $10q_H - p_H \geq 0$.

   To maximize profit, THY sets the IR to bind: $p_H = 10q_H$. Substitute this into profit: $\max_{q_H} \pi_H = 10q_H - \frac{1}{2}q_H^2$. FOC: $\frac{\partial \pi_H}{\partial q_H} = 10 - q_H = 0 \implies \mathbf{q_H^{FB} = 10}$. The price is $p_H^{FB} = 10(10) = \mathbf{100}$.

2. **For Economy types ($\theta_L = 5$):** THY solves: $\max_{q_L, p_L} \pi_L = p_L - \frac{1}{2}q_L^2$
   Subject to (IR-L): $5q_L - p_L \geq 0$.

   Set IR to bind: $p_L = 5q_L$. Substitute into profit: $\max_{q_L} \pi_L = 5q_L - \frac{1}{2}q_L^2$. FOC: $\frac{\partial \pi_L}{\partial q_L} = 5 - q_L = 0 \implies \mathbf{q_L^{FB} = 5}$. The price is $p_L^{FB} = 5(5) = \mathbf{25}$.

Proposition 7.4.1 (First-Best Outcome)**.** With full information, THY offers two efficient contracts:

- Business Bundle: $(q_H^{FB} = 10, p_H^{FB} = 100)$.
- Economy Bundle: $(q_L^{FB} = 5, p_L^{FB} = 25)$.

In this outcome, quality is set efficiently (marginal benefit θ = marginal cost q), and THY extracts 100% of the consumer surplus.

The Asymmetric Information Problem
Now, what happens if THY *cannot* tell the types apart and offers this first-best menu?

- An Economy type looks at the menu:
  - Buy Economy: $u_L = 5(5) - 25 = 0$.
  - Buy Business: $u_L = 5(10) - 100 = -50$.
  - The Economy type happily buys the Economy bundle. (IC-L is satisfied).
- A Business type looks at the menu:
  - Buy Business: $u_H = 10(10) - 100 = 0$.
  - Buy Economy: $u_H = 10(5) - 25 = \mathbf{25}$.
  - The Business type will **not** buy the Business bundle. They will pretend to be an Economy type to get a surplus of 25.

The first-best menu is **not incentive compatible**. THY must now design a "second-best" menu that accounts for this informational constraint.
The Constraints (Revisited):

- (IR-L): $u_L(q_L, p_L) = 5q_L - p_L \geq 0$
- (IR-H): $u_H(q_H, p_H) = 10q_H - p_H \geq 0$
- (IC-L): $5q_L - p_L \geq 5q_H - p_H$ (Economy type doesn't want Business bundle)
- (IC-H): $10q_H - p_H \geq 10q_L - p_L$ (Business type doesn't want Economy bundle)

Solving the Model (Determining the Binding Constraints): Rather than guessing, let's use intuition and logic to simplify the problem.

Step 1: Does the Business type need a Participation Constraint? Consider the Incentive Compatibility constraint for the Business type (IC-H) and the Participation Constraint for the Economy type (IR-L):

$$\underbrace{10q_H - p_H}_{u_{Biz}} \geq \underbrace{10q_L - p_L}_{\text{Utility from Eco bundle}}$$

$$\underbrace{5q_L - p_L}_{u_{Eco}} \geq 0$$

Notice that the utility the Business type gets from the Economy bundle is strictly higher than what the Economy type gets, because the Business type values quality more ($10 > 5$):

$$10q_L - p_L = (5q_L - p_L) + 5q_L = u_{Eco} + 5q_L$$

Since $u_{Eco} \geq 0$ (from IR-L) and $q_L > 0$, the Business type guarantees themselves a positive surplus just by mimicking. Therefore, $u_{Biz} > 0$, making (IR-H) automatically satisfied (slack). We can ignore (IR-H).

Step 2: Can both Incentive Constraints bind? Suppose both (IC-L) and (IC-H) bind with equality.

$$5q_L - p_L = 5q_H - p_H \implies p_H - p_L = 5(q_H - q_L)$$

$$10q_H - p_H = 10q_L - p_L \implies p_H - p_L = 10(q_H - q_L)$$

This implies $5(q_H - q_L) = 10(q_H - q_L)$, which is only possible if $q_H = q_L$ (pooling). Since the monopolist wants to discriminate ($q_H \neq q_L$), these constraints cannot both bind. Since the "rich" type (Business) has the incentive to mimic the "poor" type (Economy) to get a lower price, (IC-H) is the binding constraint. (IC-L) will be slack.

Conclusion: We set (IR-L) and (IC-H) as binding:

$$5q_L - p_L = 0 \implies \mathbf{p_L = 5q_L} \tag{7.1}$$

$$10q_H - p_H = 10q_L - p_L \implies p_H = 10q_H - (10q_L - p_L) \tag{7.2}$$

Substitute $p_L = 5q_L$ into the p_H equation:

$$p_H = 10q_H - (10q_L - 5q_L)$$

$$\mathbf{p_H = 10q_H - 5q_L}$$

The term $5q_L = (\theta_H - \theta_L)q_L$ is the **information rent** left to the Business type to ensure they buy the business class bundle.

Rewriting the Principal's Problem: Substitute these prices p_L and p_H into the profit function:

$$\max_{q_L, q_H} \quad 0.2\left([10q_H - 5q_L] - \frac{1}{2}q_H^2\right) + 0.8\left(5q_L - \frac{1}{2}q_L^2\right)$$

Solving via FOCs: We can now take derivatives with respect to q_L and q_H independently.

FOC w.r.t q_H:

$$\frac{\partial \pi}{\partial q_H} = 0.2(10 - q_H) = 0 \implies \mathbf{q_H^* = 10}$$

This is **"No distortion at the top."** THY provides the efficient quality ($q_H^* = \theta_H = q_H^{FB}$) to the Business (high-value) type.

FOC w.r.t q_L:

$$\frac{\partial \pi}{\partial q_L} = 0.2(-5) + 0.8(5 - q_L) = 0$$

$$-1 + 4 - 0.8q_L = 0$$

$$3 = 0.8q_L \implies \mathbf{q_L^* = 3.75}$$

The efficient quality for the Economy type would be $q_L^{FB} = 5$. But THY offers $q_L^* = 3.75$. This is **"Downward distortion at the bottom."** THY inefficiently lowers the quality of the "Economy" ticket (compared to q_L^{FB}) to make it less attractive to Business types, which allows THY to charge a higher price for the "Business" ticket.

The Optimal "Second-Best" Menu:

- **Business Class (for Business types):**
  - Quality: $q_H^* = 10$
  - Price: $p_H^* = 10(10) - 5(3.75) = 100 - 18.75 = \mathbf{81.25}$
- **Economy Class (for Economy types):**
  - Quality: $q_L^* = 3.75$
  - Price: $p_L^* = 5(3.75) = \mathbf{18.75}$

Checking Utilities (Information Rents):

- **Economy (low) type:** Buys Economy. $u_L = \theta_L q_L - p_L = 5(3.75) - 18.75 = 0$. (IR-L binds).
- **Business (high) type:** Buys Business. $u_H = \theta_H q_H - p_H = 10(10) - 81.25 = \mathbf{18.75}$.
- *Check IC-H:* If Business type bought Economy: $u_H = \theta_H q_L - p_L = 10(3.75) - 18.75 = 37.5 - 18.75 = \mathbf{18.75}$.
- The Business type is exactly indifferent, so (IC-H) binds. Their information rent is 18.75.

This is a classic "second-best" outcome. It is not fully efficient (due to $q_L^* < q_L^{FB}$), but it is the profit-maximizing strategy for THY given the information asymmetry.

7.4.1 The Intuition: Rent vs. Efficiency Trade-off

Why does THY reduce the quality of the Economy seat below the efficient level?

- THY wants to charge a high price to the Business type (p_H).
- However, the Business type always has the option to buy the cheap Economy ticket. This option places a "ceiling" on how much THY can charge for Business class.
- To raise this ceiling (and thus p_H), THY must make the Economy option **unattractive** to the Business type.
- Business travelers are very sensitive to quality ($\theta_H = 10$ vs $\theta_L = 5$). Reducing the quality of the Economy seat hurts the Business traveler *more* than it hurts the Economy traveler.
- Therefore, by distorting q_L downwards, THY can drastically reduce the value of the Economy ticket to the Business traveler, allowing them to charge a significantly higher price for the Business ticket.
- **The Trade-off:** THY sacrifices some profit from the Economy class (by selling an inefficiently bad seat) in order to gain a larger profit from the Business class (by extracting more rent).

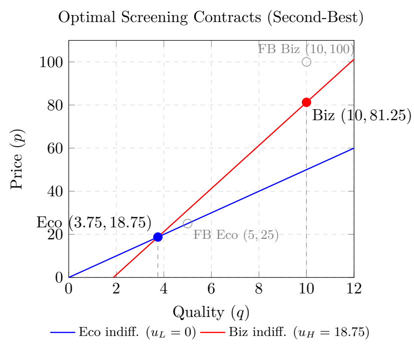

Figure 7.1: Visualization of the Screening Solution. The blue line is the Economy type's participation constraint (where they get 0 utility). The red line is the Business type's indifference curve (where they get 18.75 utility). Note clearly: The Business contract $(10, 81.25)$ is on the same red indifference curve as the Economy contract $(3.75, 18.75)$. This confirms the Incentive Constraint is binding. Also note the downward distortion: the Economy q^* (3.75) is lower than the efficient FB level (5).

Key Takeaways

- **Adverse Selection** is a pre-contractual information asymmetry (hidden type).
- **Market Unraveling:** If the uninformed party cannot screen, the market may collapse (Akerlof's Lemons), leaving only the lowest quality goods.
- **Screening:** An uninformed Principal can design a menu of contracts to separate types.
- **Revelation Principle:** We can restrict attention to direct mechanisms where agents truthfully reveal their type.
- **Second-Best Properties:**
  1. **No Distortion at the Top:** The highest type gets the efficient quality ($q_H^* = q_H^{FB}$).
  2. **Downward Distortion at the Bottom:** Lower types get inefficiently low quality ($q_L^* < q_L^{FB}$) to make their bundle less attractive to high types.
  3. **Information Rents:** High types get positive surplus (utility > 0) because of their informational advantage. Low types get zero surplus.

Chapter 8

Moral Hazard

8.1 Introduction to Moral Hazard

In the previous lecture, we studied **Adverse Selection**, a problem of *hidden information* that occurs *before* a contract is signed. Today, we turn to the second pillar of asymmetric information: **Moral Hazard**.

Moral Hazard describes a situation of **hidden action** that occurs *after* the contract is signed. It arises when one party (the Agent) takes actions that affect the payoff of another party (the Principal), but these actions are not directly observable or verifiable by the Principal.

Because the Principal cannot observe the action directly, they cannot write a contract based on the action itself (e.g., "I will pay you \$100 if you work hard"). Instead, they must contract on the observable *outcome* (e.g., "I will pay you \$100 if sales are high").

The Core Tension:

- To motivate the Agent to take a costly action (effort), the Principal must make the Agent's pay sensitive to the outcome (incentives).
- However, outcomes are often noisy (stochastic). Making pay sensitive to noise exposes the Agent to risk.
- If the Agent is strictly risk-averse, they dislike this risk.
- **Trade-off:** The optimal contract balances **Risk Sharing** (insuring the agent) against **Incentive Provision** (motivating the agent).

8.2 The Principal-Agent Model (Hidden Action)

We consider a standard setup involving a risk-neutral Principal and a strictly risk-averse Agent.

8.2.1 The Setup

1. **Players:**
    - **Principal (P):** Owner of a project. Risk-neutral. Maximize expected profit: $\mathbb{E}[\text{Output} - \text{Wage}]$.
    - **Agent (A):** Hired to work. Strictly risk-averse. Maximizes expected utility: $\mathbb{E}[u(\text{Wage})] - \text{Cost of Effort}$.
2. **Effort (e):** The Agent chooses an effort level $e \in \{0, 1\}$. This effort is costly, denoted by $c(e)$. We assume that high effort is more costly than low effort:

$$c(1) > c(0)$$

3. **Output (y):** The outcome of the project, y, is stochastic but correlated with effort.
   - Let $y \in \{y_H, y_L\}$ be the set of possible outputs (Success or Failure).
   - The probability of success depends on effort: $\mathbb{P}(y_H \mid e) = p(e)$.
   - We assume $p(1) > p(0)$, meaning higher effort increases the chance of success. We denote these probabilities as $p_H = p(1)$ and $p_L = p(0)$.
4. **Information Structure:**
   - The Principal observes the outcome y.
   - The Principal **cannot** observe the effort e.
5. **Contract:** The contract specifies a wage w contingent on the observable outcome y.

$$w(y) = \begin{cases} w_H & \text{if } y = y_H \\ w_L & \text{if } y = y_L \end{cases}$$

8.2.2 The Principal's Problem

The Principal wants to induce a specific effort level e^* that maximizes their profit, subject to the Agent accepting the job and actually choosing that effort.

The optimization problem is:

$$\max_{w_H, w_L, e} \quad \underbrace{p(e)(y_H - w_H) + (1 - p(e))(y_L - w_L)}_{\text{Expected Profit}}$$

Subject to:

1. **Individual Rationality (IR) / Participation Constraint:** The Agent must get at least their outside utility $\underline{u}$.

$$p(e)u(w_H) + (1 - p(e))u(w_L) - c(e) \geq \underline{u}$$

2. **Incentive Compatibility (IC):** The chosen effort e must be optimal for the Agent, given the wages.

$$e \in \arg\max_{e' \in \{0,1\}} [p(e')u(w_H) + (1 - p(e'))u(w_L) - c(e')]$$

8.3 Theoretical Derivations via KKT

Before we solve a numerical example, let's derive the properties of the optimal contract mathematically using the Karush-Kuhn-Tucker (KKT) conditions.

8.3.1 First-Best: Observable Effort

Suppose the Principal can observe effort. The Incentive Compatibility (IC) constraint disappears because the contract can be directly contingent on effort (e.g., "I pay you w if $e = e^*$, else 0"). The Principal simply minimizes wage cost subject to the Participation Constraint (IR).

Let $e_H = 1$ be the target effort. The Lagrangian for maximizing profit is:

$$\mathcal{L}^{FB}(w_H, w_L, \lambda) = [p_H(y_H - w_H) + (1 - p_H)(y_L - w_L)] + \lambda[p_H u(w_H) + (1 - p_H)u(w_L) - c_H - \underline{u}]$$

First-Order Conditions w.r.t w_H and w_L:

$$\frac{\partial \mathcal{L}^{FB}}{\partial w_H} = -p_H + \lambda p_H u'(w_H) = 0 \implies \frac{1}{u'(w_H)} = \lambda \tag{8.1}$$
$$\frac{\partial \mathcal{L}^{FB}}{\partial w_L} = -(1-p_H) + \lambda(1-p_H)u'(w_L) = 0 \implies \frac{1}{u'(w_L)} = \lambda \tag{8.2}$$

Result: Since $u'(w)$ is strictly decreasing (strict risk aversion), $u'(w_H) = u'(w_L)$ implies $\mathbf{w_H = w_L}$.

Proposition 8.3.1 (Full Insurance)**.** In the First-Best contract, the risk-neutral Principal bears all the risk. The strictly risk-averse Agent receives a **fixed wage** regardless of the outcome.

8.3.2 Second-Best: Unobservable Effort

Now, effort is hidden. We must respect the IC constraint. To induce high effort e_H over e_L, the IC constraint is:

$$[p_H u(w_H) + (1-p_H)u(w_L) - c_H] \geq [p_L u(w_H) + (1-p_L)u(w_L) - c_L]$$

Rearranging: $(p_H - p_L)(u(w_H) - u(w_L)) \geq (c_H - c_L)$.

The Lagrangian for the Second-Best problem is:

$$\begin{aligned}\mathcal{L}(w_H, w_L, \lambda, \mu) = {} & [p_H(y_H - w_H) + (1-p_H)(y_L - w_L)] \\ & + \lambda\,[p_H u(w_H) + (1-p_H)u(w_L) - c_H - \underline{u}] \\ & + \mu\,[(p_H - p_L)(u(w_H) - u(w_L)) - (c_H - c_L)]\end{aligned}$$

First-Order Conditions:

$$\frac{\partial \mathcal{L}}{\partial w_H} = -p_H + \lambda p_H u'(w_H) + \mu(p_H - p_L)u'(w_H) = 0 \tag{8.3}$$
$$\frac{\partial \mathcal{L}}{\partial w_L} = -(1-p_H) + \lambda(1-p_H)u'(w_L) + \mu(p_L - p_H)u'(w_L) = 0 \tag{8.4}$$

Solving for the inverse marginal utility (the "price" of providing utility):

$$\frac{1}{u'(w_H)} = \lambda + \mu\frac{p_H - p_L}{p_H} \tag{8.5}$$
$$\frac{1}{u'(w_L)} = \lambda - \mu\frac{p_H - p_L}{1 - p_H} \tag{8.6}$$

Analysis:

- If $\mu = 0$ (IC doesn't bind), we return to the First-Best ($\frac{1}{u'(w_H)} = \frac{1}{u'(w_L)} = \lambda$), which means fixed wages. But with fixed wages, the agent shirks. Thus, we must have $\mu > \mathbf{0}$.

- Since $\mu > 0$ and $p_H > p_L$, it follows that:

$$\frac{1}{u'(w_H)} > \lambda > \frac{1}{u'(w_L)}$$

  Since $1/u'(w)$ is increasing in w (for strictly concave u), this implies $\mathbf{w_H > w_L}$.

Proposition 8.3.2. In the Second-Best contract, the Agent must bear some risk ($w_H > w_L$) to be incentivized. The wage spread is determined by the likelihood ratio $\frac{p_H - p_L}{p_H}$.

8.4 Numerical Application: The Tech Sales Manager

Let's make this concrete with a clean numerical example.

Example 8.4.1 (Startup Sales)**.**
- **Scenario:** A risk-neutral Tech Founder (Principal) hires a strictly risk-averse Sales Manager (Agent).

- **Output:** Sales can be High ($y_H = \$1000$) or Low ($y_L = \$0$).

- **Effort:** The Manager can work Hard ($e = 1$) or Shirk ($e = 0$).

- **Cost of Effort:**

  – Cost of working Hard: $c(1) = 10$.

  – Cost of Shirking: $c(0) = 0$.

- **Probabilities:**

  – If Hard ($e = 1$): Probability of High Sales is $p_1 = 0.8$.

  – If Shirk ($e = 0$): Probability of High Sales is $p_0 = 0.4$.

- **Agent's Preferences:** The Manager has utility function $u(w) = \sqrt{w}$ and reservation utility $\underline{u} = 10$.

Step 0: Efficiency Check (Which effort is optimal?) We compare the total surplus (Expected Output - Wage Cost needed to achieve $\underline{u}$) for both actions.

- **Hard Effort:** $\mathbb{E}[y|e = 1] = 0.8(1000) = 800$. To pay the agent $\underline{u} = 10$ plus compensate for cost $c(1) = 10$, they need utility 20. Since $u(w) = \sqrt{w}$, the wage cost is $20^2 = 400$. Total Surplus $= 800 - 400 =$ **400**.

- **Shirking:** $\mathbb{E}[y|e = 0] = 0.4(1000) = 400$. To pay the agent $\underline{u} = 10$ (cost is 0), they need utility 10. Wage cost is $10^2 = 100$. Total Surplus $= 400 - 100 =$ **300**.

Since $400 > 300$, **Hard Effort is efficient**. The Principal wants to implement $e = 1$.

8.4.1 Benchmark: First-Best (Symmetric Information)

Suppose the Founder observes effort directly.

Contract: "I pay you fixed wage w if you Work Hard. If you Shirk, you get nothing." Since the Principal bears the risk, the Agent gets a fixed wage.

Constraint (IR):

$$\sqrt{w} - c(1) = \underline{u} \implies \sqrt{w} - 10 = 10 \implies \sqrt{w} = 20 \implies w^* = 400$$

Result (First-Best):

- Wage: $w^* = 400$ (Fixed).

- Principal Profit: $\mathbb{E}[y|e = 1] - w^* = 800 - 400 =$ **400**.

8.4.2 The Second-Best (Hidden Action)

Now, effort is unobservable. A fixed wage of 400 would lead to Shirking (Payoff: $\sqrt{400} - 0 > \sqrt{400} - 10$). To motivate $e = 1$, the Founder must offer a bonus contract (w_H, w_L).

The Optimization Problem: The Principal maximizes expected profit (Expected Revenue minus Expected Wage Cost).

$$\max_{w_H, w_L} \quad \underbrace{0.8(1000 - w_H) + 0.2(0 - w_L)}_{\text{Expected Profit}}$$

Subject to:

1. **(IR) Constraint:**

$$0.8\sqrt{w_H} + 0.2\sqrt{w_L} - 10 \geq 10 \implies 0.8u_H + 0.2u_L \geq 20$$

2. **(IC) Constraint:**

$$[0.8u_H + 0.2u_L - 10] \geq [0.4u_H + 0.6u_L - 0]$$

$$(0.8 - 0.4)u_H + (0.2 - 0.6)u_L \geq 10$$

$$0.4u_H - 0.4u_L \geq 10 \implies u_H - u_L \geq 25$$

Solving the System: Both constraints bind (as shown in Section 3.2). 1) From (IC): $u_H = u_L + 25$. 2) Substitute into (IR):

$$0.8(u_L + 25) + 0.2u_L = 20$$

$$u_L + 20 = 20 \implies \mathbf{u_L = 0}$$

$$u_H = 0 + 25 \implies \mathbf{u_H = 25}$$

Wages:

$$w_L = u_L^2 = 0^2 = \mathbf{0}$$

$$w_H = u_H^2 = 25^2 = \mathbf{625}$$

Analysis of the Second-Best Solution:

- **Expected Wage:** $\mathbb{E}[w] = 0.8(625) + 0.2(0) = 500$.
- **Agency Cost:** First-Best wage cost was 400. Second-Best expected wage cost is 500. The difference $(500 - 400 = 100)$ is the **Risk Premium** the Principal pays to motivate the Agent.
- **Profit:** $800 - 500 = \mathbf{300}$. (Profit is lower than First-Best, but still positive).

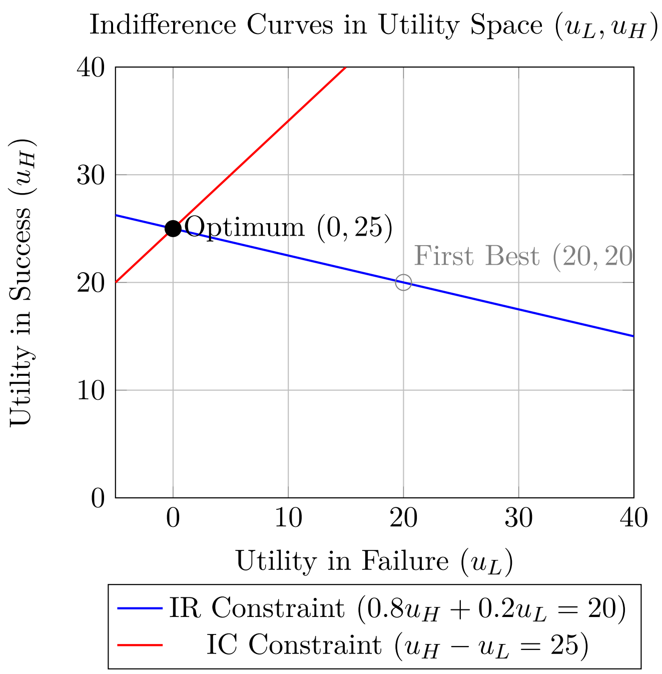

Figure 8.1: Visualization of the Optimal Contract. The blue line is the Participation Constraint (IR). The red line is the Incentive Constraint (IC). The First-Best point is $(20, 20)$ (full insurance). The Second-Best solution must lie on the intersection of the constraints at $(0, 25)$, imposing risk on the agent $(u_H > u_L)$.

Key Takeaways

- **Moral Hazard** is a post-contractual information problem (hidden action).
- **Incentive Compatibility:** Because effort is unobservable, the contract must depend on the noisy outcome. High effort is induced only if the Agent faces a "wedge" between the reward for success and failure $(w_H > w_L)$.
- **The Fundamental Trade-off:**
  - **Risk Sharing:** Efficiency requires the risk-neutral Principal to insure the risk-averse Agent (constant wage).
  - **Incentives:** Efficiency requires the Agent to bear risk to be motivated.
- **Agency Cost:** The Second-Best contract is more expensive for the Principal than the First-Best because they must pay a **Risk Premium** to the Agent.

8.5 Appendix A: The Karush-Kuhn-Tucker (KKT) Theorem

The KKT theorem generalizes the method of Lagrange multipliers to optimization problems with inequality constraints.

Problem Setup: Consider the problem of maximizing an objective function $f(x)$ subject to m inequality constraints $g_j(x) \leq b_j$:

$$\max_{x \in \mathbb{R}^n} f(x) \quad \text{subject to} \quad g_j(x) \leq b_j \quad \text{for } j = 1, \ldots, m$$

(Note: Constraints of the form $h(x) \geq c$ can be rewritten as $-h(x) \leq -c$).

The Lagrangian: We define the Lagrangian function as:

$$\mathcal{L}(x, \lambda) = f(x) - \sum_{j=1}^{m} \lambda_j (g_j(x) - b_j)$$

where λ_j are the KKT multipliers (or shadow prices).

Necessary Conditions (KKT Conditions): If x^* is a local maximum for the constrained problem, there exist unique constants $\lambda^* = (\lambda_1^*, \ldots, \lambda_m^*)$ such that the following conditions hold:

1. **Stationarity:** The gradient of the Lagrangian with respect to x is zero.

$$\nabla f(x^*) - \sum_{j=1}^{m} \lambda_j^* \nabla g_j(x^*) = 0$$

2. **Primal Feasibility:** The constraints must be satisfied.

$$g_j(x^*) \leq b_j \quad \text{for all } j$$

3. **Dual Feasibility:** The multipliers must be non-negative.

$$\lambda_j^* \geq 0 \quad \text{for all } j$$

(Note: For equality constraints $h(x) = c$, the multiplier sign is unrestricted).

4. **Complementary Slackness:**

$$\lambda_j^* (g_j(x^*) - b_j) = 0 \quad \text{for all } j$$

This condition implies that for every constraint j, either the constraint is binding ($g_j(x^*) = b_j$) or its multiplier is zero ($\lambda_j^* = 0$).

Chapter 9

Signalling

9.1 Introduction: The Failure of Pooling

In previous lectures, we established that **Adverse Selection** (hidden information) can lead to market unraveling. Akerlof's "Lemons Market" showed that if quality is unobservable, high-quality goods (or agents) may exit the market.

Today we study **Signalling**, a mechanism where informed agents take costly actions to credibly reveal their private information to uninformed principals. We ground this in the context of the Labor Market.

9.1.1 The Competitive Setup (No Signals)

Consider a labor market with the following characteristics:

1. **Agents (Workers):** There are two types of workers, $\theta \in \{G, B\}$ (Good and Bad).
   - Proportion of Good types: $\lambda \in (0, 1)$.
   - Productivity: A Good worker produces y_G, a Bad worker produces y_B, with $y_G > y_B$.
   - Reservation Utility: Normalized to 0 for simplicity.
2. **Principals (Firms):** There are N firms and M workers, with $N > M$.
   - Firms are risk-neutral profit maximizers.
   - Firms behave competitively (Bertrand competition for talent).
3. **Information:** Worker type θ is private information. Firms only know the distribution λ.

9.1.1.1 The Zero Profit Condition

Because there are more firms than workers ($N > M$), firms bid up wages until their expected profits are zero.

If firms cannot distinguish between types, they must offer a single wage w_{pool} based on the **expected productivity** of the pool:

$$\mathbb{E}[y] = \lambda y_G + (1 - \lambda) y_B$$

Proposition 9.1.1 (Pooling Wage)**.** In a competitive equilibrium without signalling, the wage is equal to the average productivity:

$$w_{pool} = \lambda y_G + (1 - \lambda) y_B$$

The firms earn zero profit: $\Pi = \mathbb{E}[y] - w_{pool} = 0$.

The Problem: Note that $w_{pool} < y_G$. The Good type is cross-subsidizing the Bad type. The Good type is not receiving their full marginal product. This creates an incentive for the Good type to find a way to distinguish themselves—to "prove" they are Good to capture the rent $(y_G - w_{pool})$.

9.2 The Spence Job Market Model

Michael Spence (1973) introduced the idea that agents can engage in an activity—say, **Education** (e)—that has no intrinsic effect on productivity but serves as a signal of type.

9.2.1 The Setup with Signalling

- **Timing:**
  1. Nature determines worker type $\theta \in \{G, B\}$.
  2. Worker chooses education level $e \geq 0$. This is observable.
  3. Firms observe e and simultaneously make wage offers $w(e)$.
  4. Worker accepts the highest wage offer.
- **Payoffs:**
  - **Firm Profit:** $\Pi = y_\theta - w$. (Note: Education does *not* increase y. It is purely a signal).
  - **Worker Utility:** $U(w, e, \theta) = w - c(e, \theta)$.
- **Cost of Education:** Education is costly (effort, tuition, time). Crucially, we assume education is *more costly* for the Bad type than for the Good type.

$$c(e, B) > c(e, G) \quad \text{for all } e > 0$$

$$c'(e, B) > c'(e, G) > 0$$

Definition 9.2.1 (Single Crossing Property)**.** The marginal cost of the signal is strictly higher for the Bad type. In terms of indifference curves in (e, w) space, the Bad type's indifference curve is steeper than the Good type's at any point.

$$\frac{\partial^2 U}{\partial e \partial \theta} > 0 \quad (\text{equivalently } -\frac{\partial^2 c}{\partial e \partial \theta} > 0)$$

9.3 Separating Equilibrium

We look for a **Perfect Bayesian Equilibrium (PBE)** where types separate.

- **Separating Strategy:** Good types choose education e^*, Bad types choose $e = 0$.
- **Beliefs:** Firms believe anyone with $e < e^*$ is type B, and anyone with $e \geq e^*$ is type G.
- **Competition (Zero Profit):**
  - If believed to be Bad, wage $w(0) = y_B$.
  - If believed to be Good, wage $w(e^*) = y_G$.

The Good types "burn money" on education to prove their type. They do this to unlock the higher wage y_G.

9.3.1 Incentive Compatibility (IC) Constraints

For a separating equilibrium to exist, two conditions must hold:

1. **IC for the Bad Type (No Mimicry):** The Bad type must prefer being identified as Bad (getting y_B with no education) rather than mimicking the Good type (getting y_G with high education).
$$\underbrace{y_B - c(0, B)}_{\text{Payoff as Bad}} \geq \underbrace{y_G - c(e^*, B)}_{\text{Payoff mimicking Good}}$$
Since $c(0, B) = 0$, this simplifies to:
$$c(e^*, B) \geq y_G - y_B$$
The cost of the signal must be high enough to deter the Bad type.

2. **IC for the Good Type (Participation):** The Good type must prefer sending the signal and getting the high wage rather than being pooled with the low types.
$$\underbrace{y_G - c(e^*, G)}_{\text{Payoff as Good}} \geq \underbrace{y_B - c(0, G)}_{\text{Payoff if pooled/mistaken}}$$
Simplifying:
$$y_G - y_B \geq c(e^*, G)$$
The wage premium must be enough to cover the signaling cost for the Good type.

9.3.2 Existence Condition

Combining the two inequalities:

$$c(e^*, G) \leq y_G - y_B \leq c(e^*, B)$$

A separating equilibrium exists if there is an education level e^* such that the wage premium $(y_G - y_B)$ is *greater* than the Good type's cost but *less* than the Bad type's cost.

Theorem 9.3.1 (Existence of Separating Equilibrium)**.** A separating equilibrium exists if and only if the Single Crossing Property holds. Under this condition, there is a range of e^* that supports separation. The Good type captures the rent y_G but dissipates some of it through the cost of signaling $c(e^*, G)$.

9.3.3 The Least-Cost Separating Equilibrium

There are many e^* that satisfy the condition. However, in standard refinements (like the Intuitive Criterion), the Good type will choose the **minimum** education necessary to separate.

The Good type sets e^* such that the Bad type is *just indifferent* between mimicking and not mimicking:

$$c(e^*, B) = y_G - y_B$$

At this level, the Bad type does not envy the Good type, and the Good type maximizes their net utility subject to separation.

Key Takeaways

- **Competition for Talent:** Risk-neutral principals compete for agents, driving expected profits to zero. In a separating equilibrium, wages equal the actual productivity of each type ($w_L = y_B, w_H = y_G$).
- **Signalling Value:** Education has no social value in this model (it does not increase y). However, it has *private* value to the Good type because it allows them to escape the pool of "lemons."
- **Burning Money:** The Good agents willingly incur the cost of education ("burn money") to prove their type. They capture the rent created by their higher productivity, net of the signaling cost.
- **Single Crossing Condition:** Signalling works only if the signal is *differentially costly.* If it were equally easy for Bad types to get a degree, they would mimic the Good types, and the equilibrium would collapse.

9.4 Appendix: Single Crossing Property Definition

The Single Crossing Property (SCP), also known as the Spence-Mirrlees condition, is the mathematical foundation of sorting in signaling models.

Let $u(w, e, \theta)$ be the utility function of the agent. In the (e, w) plane, the slope of the indifference curve is given by the marginal rate of substitution between wage and education:

$$\text{MRS}_{e,w} = -\frac{\partial u/\partial e}{\partial u/\partial w}$$

Assuming $u(w, e, \theta) = w - c(e, \theta)$, we have:

$$\text{MRS}_{e,w} = c_e(e, \theta)$$

The Single Crossing Property states that this slope is monotonic in type θ. Specifically, if Bad types have higher marginal costs:

$$\frac{\partial}{\partial \theta}\left(\text{MRS}_{e,w}\right) < 0$$

(Assuming $\theta_G > \theta_B$ implies lower cost).

Geometrically, this ensures that the indifference curves of different types cross exactly once. This enables the existence of a screen (a specific (w, e) contract) that one type accepts and the other rejects.